



Fabric Engine v9.4 User Guide

Features and Configuration for Universal Platforms

9039513-00 Rev AB
April 2026



Copyright © 2026 Extreme Networks, Inc.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: <https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks>

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Abstract

This user guide for Fabric Engine version 9.4 provides instructions for configuring and managing Extreme Networks' switches, including the 4220, 5320, 5420, 5520, 5720, 7520, 7720, and 7830 Series. It covers various features such as Zero Touch Capabilities, Auto-sense functionality, Fabric Attach (FA), Fabric Extend (FE), and Extensible Authentication Protocol (EAP). The guide details command usage for switch configuration, including setting up VLANs, managing software upgrades, configuring network protocols like IS-IS and RADIUS, and implementing security features like MACsec. Additionally, it addresses advanced configurations, including Auto-sense for Fabric Attach Proxy and handling Time Domain Reflectometer (TDR) diagnostics.



Preface

[Purpose](#) on page 4
[Conventions](#) on page 4
[Documentation and Training](#) on page 7
[Help and Support](#) on page 7
[Send Feedback](#) on page 8

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Purpose

This document provides information on features in Extreme Networks Fabric Engine™. This release of Fabric Engine runs on the following product families:

- ExtremeSwitching 4220 Series
- ExtremeSwitching 5320 Series
- ExtremeSwitching 5420 Series
- ExtremeSwitching 5520 Series
- ExtremeSwitching 5720 Series
- ExtremeSwitching 7520 Series
- ExtremeSwitching 7720 Series
- ExtremeSwitching 7830 Series



Note

Prior to Fabric Engine 8.6, 5520 Series and 5420 Series platforms ran VOSS. VOSS support ends for these platforms with VOSS 8.5.x.

Conventions

To help you better understand the information presented in this guide, the following topics describe the formatting conventions used for notes, text, and other elements.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notes and warnings






Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product.
	Note	Useful information or instructions.
	Important	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.

Table 2: Text conventions

Convention	Description
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.

Table 3: Command syntax (continued)

Convention	Description
	If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code> , you can enter <code>cfm maintenance-domain maintenance-level 4</code> .
Bold text	Bold text indicates the GUI object name you must act upon. Examples: <ul style="list-style-type: none"> Select OK. On the Tools menu, choose Options.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. For example, if the command syntax is <code>ip address {A.B.C.D}</code> , you must enter the IP address in dotted, decimal notation.
Brackets ([])	Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is <code>show clock [detail]</code> , you can enter either <code>show clock</code> or <code>show clock detail</code> .
Ellipses (...)	An ellipsis (...) indicates that you repeat the last element of the command as needed. For example, if the command syntax is <code>ethernet/2/1 [<parameter> <value>] ...</code> , you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.
<i>Italic Text</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
Plain Courier Text	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: <ul style="list-style-type: none"> <code>show ip route</code> <code>Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]</code>

Table 3: Command syntax (continued)

Convention	Description
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation pane, expand Configuration > Edit .
Vertical Line ()	A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is <code>access-policy by-mac action { allow deny }</code> , you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code> , but not both.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and Software Compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other Resources](#) such as articles, white papers, and case studies

Open Source Declarations

Some software files have been licensed under certain open source licenses. Information is available on the [Open Source Declaration](#) page.

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit the [Extreme Networks Training](#) page.

Help and Support

If you require assistance, contact Extreme Networks using one of the following methods:

[Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

[The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2800. For the support phone number in your country, visit www.extremenetworks.com/support/contact.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

Send Feedback

The User Enablement team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, email us at Product-Documentation@extremenetworks.com.

Provide as much detail as possible including the publication title, topic heading, and page number (if applicable), along with your comments and suggestions for improvement.



Document Revision Changes

The following table identifies changes between revisions of the same release document.

Table 4: 9.4 User Guide revision changes

Revision	Change
AA	Initial revision for new release, see New in this Document on page 11
AB	Updated the following topics: <ul style="list-style-type: none">• Enable Factory Default Behavior on page 363• Third Party Virtual Machine on page 1101• Premier License on page 2414• Extreme-Dynamic-Client-Assignments on page 3128



New in this Document

[Notice about Feature Support](#) on page 19

The following sections detail what is new in this document.

Fabric Enhancements

The software supports the following Fabric enhancements:

- Auto-sense guest I-SID control—In this release, you can now disable the use of the onboarding I-SID as a guest I-SID.



Note

This functionality applies to Auto-sense ports in UNI-ONBOARDING, FA, and VOICE states. It does not apply to FA-PROXY, FA-PROXY-NOAUTH, FA-PROXY-RING, or NNI states.

For more information, see the following sections:

- [Auto-sense Guest I-SID Control](#) on page 24
- [Disable Auto-sense Guest I-SID](#) on page 44 (CLI)
- [Disable Auto-sense Guest I-SID](#) on page 60 (EDM)
- Auto-sense Link Debounce— You can use Auto-sense to configure Link Debounce on all Auto-sense UNI ports or only those that connect to Preboot Execution Environment (PXE) devices (by configuring it to *auto*). This enhancement addresses the following issue: PXE devices make a port bounce after the initial DHCP Discovery. This port bounce restarts the Auto-sense state machine, the port exits the UNI state, the MAC is de-authenticated, and the wait-timeout period restarts. This process repeats in an endless loop, preventing the PXE device from receiving and applying the configuration. After the port bounce, the PXE device sends three DHCP Discovery messages during the first 6 seconds. The Link Debounce *auto* configuration enables the Link Debounce option after a well-known PXE packet is recognized on a port, which avoids a link bounce and allows the DHCP discovery

packets to pass through successfully and ensure proper onboarding of PXE booted devices on Auto-sense ports.



Note

This enhancement does not apply to 7830 Series.

With this introduction in 9.4, you can no longer use the interface-level **link-debounce** command on Auto-sense ports. Use the new **auto-sense link-debounce** command instead.

For more information, see [Configure Link Debounce for Auto-sense Ports](#) on page 56 (CLI) or [Configure Link Debounce for Auto-sense Ports](#) on page 65 (EDM).

- Auto-sense SD-WAN VRF configuration enhancements—This release introduces the following enhancements to SD-WAN dynamic VRF and local breakout (LBO) BGP configuration:
 - You can now convert a dynamic SD-WAN VRF to a static VRF, which is necessary for SD-WAN high availability (HA) branch deployments.
 - The BGP configuration on the SD-WAN assigned VRF (or GRT) now includes a BGP origin field to track whether the configuration is dynamically created by Auto-sense or is statically configured.
 - You can now convert the dynamic LBO BGP configuration on the SD-WAN VRF to a static configuration, which is necessary for SD-WAN HA branch deployments.
 - You can now convert the dynamic SD-WAN LBO BGP configuration on the GRT to a static configuration for SD-WAN branches that use 5320 Series or 4220 Series as Layer 3 branch routers.



Note

Exception: does not apply to 5320-48P-8XE and 5320-48T-8XE.

- The SD-WAN local breakout configuration on a user VRF now persists even when the SD-WAN Auto-sense port goes down, which is necessary for SD-WAN HA branch deployments.

For more information, see [SD-WAN Local Breakout](#) on page 3306.

- IPv6 discard routes—This release supports IPv6 discard static routes. A discard static route is a route with an invalid next hop that results in traffic matching the route (and not matching the more specific routes) being discarded.

For more information, see the following sections:

- [IPv6 Discard Static Routes](#) on page 2187
- [Create IPv6 Static Routes](#) on page 2207 (using CLI)
- [Create IPv6 Static Routes](#) on page 2241 (using EDM)
- Segmented VRF (and thus Segmented Layer 3 VSN)—Use a Segmented VRF to control, isolate, and secure traffic flows across the network by creating isolated routing domains within a single, traditional VRF or Layer 3 VSN. The switch separates the VRF into three access areas, or trust levels: trusted, unrestricted, and untrusted. The switch classifies the traffic into those segments by the ingress VLAN. An untrusted VLAN, and thus I-SID, can only reach unrestricted segments. Clients on trusted VLANs can reach trusted VLANs as well as unrestricted VLANs.

Typically, position a firewall on an unrestricted VLAN, which means it can respond to untrusted as well as trusted VLANs. Position unsecure IoT devices on untrusted VLANs, which means they are restricted to only reach a subset (unrestricted) destinations.

Devices on an untrusted VLAN or I-SID can communicate with each other and can form an isolated communication zone or group that can only reach unrestricted (firewall) destinations. If devices must only talk to unrestricted (firewall) destinations, use a Private VLAN (PVLAN) to further isolate the devices from each other and direct them all to an unrestricted firewall interface.

You can use a Segmented VRF to segment customers, internal services, functions, or security zones.

**Note**

Segmented VRF is not supported on 7830 Series or switch models that support a single active VRF.

This feature requires a subscription or Premier license if used with Layer 3 VSN.

For more information, see [Segmented VRF](#) on page 3446.

Operational Enhancements

The software supports the following operational enhancements:

- CDP and LLDP on the same port—You can now configure a port to send and receive both CDP and LLDP packets. In previous releases, you could not enable CDP and LLDP on the same port.

For more information, see [Enable CDP Mode on a Port](#) on page 2473 (CLI) and [Enable CDP Mode on a Port](#) on page 2497 (EDM).

- DHCP Option 150 for DHCP Server—Configure a list of up to eight TFTP server IP addresses. Clients, such as IP Phones, can request Option 150 from a DHCP Server to obtain configuration files or other information from a TFTP server.

For more information, see the following tasks:

- [Configure DHCP Server General Parameters](#) on page 833 (CLI)
- [Configure a DHCP Subnet](#) on page 836 (CLI)
- [Configure a Global TFTP Server List for the DHCP Server](#) on page 874 (EDM)
- [Configure a TFTP Server List for the DHCP Server Subnet](#) on page 879 (EDM)
- EDM changes—This release introduces the following changes:
 - On the 7830 Series, the LEDs on the management ports for 1Gbps and 10Gbps links now blink to indicate link activity. Previously, when the management ports operated at either 1Gbps or 10Gbps, the LEDs displayed solid green during link activity.
 - EDM now displays clear status notifications for both successful and failed save operations.
 - The EDM navigation pane now includes a button to close all open tabs.

- EDM access level passwords now support up to 80 characters. In previous releases, access level passwords could not exceed 32 characters.
- The **MACsec KA Key** tab located in **Configuration > Security > Data Path > MACsec** now includes the **KeyStatus** field, which provides MKA Key status for Valid, Expired, and In-Use.
- **PortMembers** fields are added to the **Fabric > IS-IS > Protocol Summary** tab to show MLT port members for IS-IS Interfaces and IS-IS Adjacency View.
- Security and Qos Group fields located in **Security > Data Path > Advanced Filters (ACE/ACLs) ACL** tab are now renamed Primary and Secondary banks.
- The **Security > Control Path > General > Web** tab adds a **InUseCertType** field.
- The **SrcVrfID** parameter on **IP > <Protocol> > Redistribute** tabs was read-only in previous release. You can now configure this value.
- VLAN lists now provide check boxes to select multiple VLANs simultaneously. Previously, you had to press and hold the **Ctrl** key when selecting multiple VLANs.
- In previous releases, when you configure the date on the switch, the maximum configurable year was 2038. In this release, support extends to year 2100. This change is implemented in CLI, EDM, and SNMP.
- IPFIX configuration per port—On IPFIX-supporting platforms, you can enable or disable IPFIX on all NNI ports or on individual UNI ports.

For more information, see [IPFIX Configuration Using CLI](#) on page 442 and [IPFIX Configuration Using EDM](#) on page 461 .

- Logging of **show khi performance-scaling** watermarks reached—After a resource monitored and displayed in the **show khi resource-scaling** command output reaches 80%, 90%, and 100%, the switch logs a WARNING message, sends an SNMP trap, and sets an alarm. If you use Extreme Cloud management applications, the switch also sends a message to that application.

For more information, see [Display Fabric KHI Resource Scaling](#) on page 2320.

- Multiple Spanning Tree Protocol Restricted Role (Root Guard) and Restricted TCN—This release introduces MSTP Restricted Role and Restricted TCNs. MSTP Restricted role prevents a port from accepting superior BPDUs from non-root bridges. When triggered, it puts the port in a root-inconsistent state. Restricted TCN prevents a switch port from propagating topology-change messages to other ports. When enabled, the port blocks and ignores all received TCNs.

For more information, see the following sections:

- [Multiple Spanning Tree Protocol Restricted Role and Restricted TCN](#) on page 3504
- [Examples](#) on page 3505
- [Configure MSTP Restricted Role](#) on page 3521
- [Configure Restricted Topology Change Notification](#) on page 3521
- [Display MSTP Port Configuration Information](#) on page 3528
- [Configure CIST Ports for MSTP](#) on page 3548

- OpenAPI Enhancements—The following OpenAPI Enhancements are available as of Release 9.3.1:
 - **openapi local-mgmt ttl <60-86400>** command—Use this command to configure the time-to-live (ttl) value for the authentication token.
 - **openapi local-mgmt minimum-tls {1.2 | 1.3}** command—Use this command to configure the minimum TLS version.
 - **show application openapi log** command—This command now includes the *reverse* parameter to display Open API logs entries in chronological order.

For more information, see [Configure OpenAPI Local Port Access](#) on page 2742 and [Display OpenAPI Log File Information](#) on page 2743.

- **show khi resource-scaling** command—This command output now provides a clearer view of how switches allocate and share hardware resources.

For more information, see [Display Fabric KHI Resource Scaling](#) on page 2320.

- **show license** command—This command output is enhanced with extra information to explicitly indicate when Premier features are included with Extreme Platform ONE Networking licenses.

For more information, see [Show a License File](#) on page 2422.

- TFTP Block Number Rollover—In this release, the software uses block number rollover functionality to transfer files larger than 32 MB using TFTP. With this functionality, when the block number reaches 65,535, it resets to 0 while it maintains the standard block size of 512 bytes and allows the transfer to continue seamlessly. In earlier releases, TFTP data blocks were numbered sequentially, which caused the transfer to stop when the block number reached 65,535 and limited the maximum file size.
- TTL handling of bridged traffic with routed SPB IP Multicast—This release adds Layer 2 VSN-based forwarding to IP Multicast over Fabric Connect. When a multicast sender and receiver reside in the same Layer 2 VSN, the Fabric bridges the multicast traffic instead of routing it. The switch preserves the IP TTL and the source C-MAC address, which prevents TTL-related packet drops for low-TTL protocols such as PTPv1. The ingress BEB advertises TLV 188 to identify the sender and Layer 2 VSN, and receiving BEBs use this information to select bridged or routed forwarding for each multicast stream. This feature operates across all Multi-area Fabric Connect topologies.

For more information, see [TTL Handling for Bridged Traffic with Routed SPB IP Multicast](#) on page 1980, [Enable Routed SPB IP Multicast TTL Bridging](#) on page 1983 (CLI), and [Enable TTL Handling for Bridged Traffic with Routed SPB IP Multicast](#) on page 1994 (EDM).

- VRF name autocompletion—You can use the CLI command completion features for VRF names in show and configuration commands.

For more information, see [Command Completion](#) on page 186.

Platform Enhancements

The software supports the following platform enhancements:

- 7830 Series supports the following new Versatile Interface Modules (VIMs) with MACsec-capable ports:
 - 7830-VIM-24CE—24 x 10Gb/25Gb/100Gb SFP56-DD ports
 - 7830-VIM-8DE—8 x QSFP56-DD 400Gb ports; when channelized - 4x100Gb, 4x10Gb, and 4x25Gb or 40Gb and 100Gb with a single QSFP+ or QSFP28 transceiver
- 7830 Series supports 100G SFP56-DD transceivers on the 7830-VIM-24CE VIM.

For optics compatibility, see the [Extreme Optics](#) website.

- This release provides hot-swap support for 7830 Series VIMs.

For more information, see the following sections:

- [Hot Swap a VIM on 7830 Series](#) on page 712
- [Swap a VIM](#) on page 733 using CLI
- [Swap a VIM](#) on page 797 using EDM
- On the 7830 Series, the system log message now includes both the port and the group associated with the transceiver.

For more information, see [Auto-speed](#) on page 677.

- This release supports the following software features on 7830 Series:
 - ACL inVLAN
 - Bridge Protocol Data Unit (BPDU) Guard
 - Dynamic Nickname Server
 - E-tree
 - E-tree support with Auto-sense
 - IGMP Snooping
 - IPFIX
 - IP Multicast over Fabric Connect
 - IPv4 IS-IS accept policies
 - IPv4 Inter-VRF Routing (RIP, OSPF, BGP)
 - IPv4 RSMLT
 - IPv6 BGP (GlobalRouter and VRF)
 - IPv6 DHCP Relay
 - IPv6 ECMP
 - IPv6 Inter I-SID Routing
 - IPv6 Management CLIP and VLAN
 - IPv6 Neighbor discovery and Routing
 - IPv6 OSPF (GlobalRouter and VRF)
 - IPv6 RIP (GlobalRouter and VRF)
 - IPv6 Routing Layer 2 and Layer 3 VSN

- IPv6 RSMLT
- IPv6 Shortcuts
- IPv6 Shortcuts with ECMP
- IPv6 VRFs
- IPv6 VRRP
- Private VLAN
- SLPP Guard
- SSH to (BMAC) IS-IS Host-name or System-ID
- Domain resolution test for a specific DNS—When you enable dynamic IP configuration, the Network Service Probe interface uses the DNS server information it receives from the DHCP server to test DNS resolution. The interface obtains its IP address, default gateway, and up to three DHCP-advertised DNS servers. You can query any of these DNS servers—primary, secondary, or tertiary—to verify connectivity.

For more information, see [Network Service Probe](#) on page 3869 and [Query Domain Name Server using Service Probe](#) on page 3917.

- Secure log file transfer—You can now configure log file transfer, with the **logging transferFile {1-10}** command, to use Secure Copy (SCP) rather than TFTP or FTP.

For more information, see [Log File Transfer](#) on page 2545.

- VLAN-based Transparent Clock for PTPv2 —This release introduces Precision Time Protocol version 2 (PTPv2) Transparent Clock with VLANs. This feature improves time-synchronization accuracy across the network by compensating for switch latency in PTP timing messages. The switch measures the residence time of PTP packets and updates the Correction Field as packets traverse the network, which maintains accurate end-to-end timing between the timeTransmitter Clock and timeReceiver Clock within a VLAN.



Note

The IEEE 1588 Working Group updated its terminology, replacing master and slave with timeTransmitter and timeReceiver. This document follows that convention.

This feature requires a Premier or subscription license.



Note

This feature does not apply to 4220 Series and 7830 Series.

For more information, see [PTPv2 Transparent Clock with VLANs](#) on page 710 , [Enable PTPv2](#) on page 732, and [Configure System Flags](#) on page 753.

- 5720 Series minimum fan speed—Use the new **sys fan set-min-speed <20-100>** command to configure the minimum fan speed as a percentage between the minimum supported speed and the maximum supported speed and adjust the thermal operations of the switch. This value is the minimum speed at which the fan operates; the switch increases the fan speed when necessary.

For more information, see [Configure the Minimum Fan Speed](#) on page 742 (CLI) and [Configure the Minimum Fan Speed](#) on page 791 (EDM).

Security Enhancements

The software supports the following security enhancements:

- Enhanced Secure Mode (ESM)—If the switch operates in Enhanced Secure Mode, this release introduces the following changes:



Note

The switch generates audit logs if you try to enable unapproved algorithms or key exchange methods in ESM.

- The default minimum TLS version for Syslog is TLS 1.2.

For more information, see [Logs and Traps](#) on page 2535.

- RSA SHA224 and ECDSA SHA224, SHA256, SHA384, and SHA512 are disabled during the SSL handshake.
- For the web server, DHE ciphers are disabled during SSL handshake.
- The following SSH key exchange methods are disabled by default:
 - diffie-hellman-group-exchange-sha256
 - diffie-hellman-group14-sha1

For more information, see [Secure Shell](#) on page 3321.

- The following additional SSH encryption types are disabled by default:
 - aes192-cbc
 - aes192-ctr
 - rijndael128-cbc
 - rijndael192-cbc
- SSH host key algorithm x509v3-ssh-rsa is no longer allowed.

For more information, see [Secure Shell](#) on page 3321.

- MACsec Enhancements—This release supports MACsec Key Agreement (MKA) and MKA Keychains on 5320 Series and 5420 Series.



Note

This feature does not apply to 5320-16P-2MXT-2X and 5320-24T-4X-XT.

For more information, see [MACsec Key Agreement Protocol](#) on page 2607.

- RADIUS VSA Enhancements—The following RADIUS VSA Enhancements are available as of Release 9.3.1:
 - Additional parameters for Extreme-Dynamic-Client-Assignments Vendor Specific Attribute (VSA) used in RADIUS for dynamic VLAN and PVLAN assignment:
 - none—Use an existing VLAN or PVLAN instead of creating a new one.

- `igmpqaddr=<IPv4 address>`—Configure the IGMP Querier address for traffic within the VLAN either for IGMP Snooping or for Multicast Lite or Routed Multicast.

The updated string format to create a dynamic VLAN is as follows:

```
create=vlan|pvlan|none, pv=Primary VLANID, sv=secondary VLANID,  
vni=L2-ISID, ev=EGRESS-VLAN-tag, vn=vlan-name, vnin=isid-name,  
mvni=MVPN-ISID, igmpqaddr=<IPv4 address>
```

- Additional IGMP features for Extreme-Dynamic-Config Vendor Specific Attribute (VSA) used in RADIUS:
 - IGMP version 3 (IGMPV3)
 - IGMP Fast Leave (IGMPFAST)

For more information, see [Extreme-Dynamic-Client-Assignments](#) on page 3128 and [Extreme-Dynamic-Config](#) on page 3116.

- SSH packet size—The maximum SSH packet size is 35840 bytes.

Documentation Changes

[1Gb SFP Auto-Negotiation in SFP+ and SFP28 Ports](#) on page 697 is added to the document.

This release restructures the following sections:

- [Licensing](#) on page 2413
- [OSPF](#) on page 2748

[Configure System Message Control](#) on page 378 is added to the document.

[Interaction Between Anycast IP Gateway, Distributed Virtual Routing and SLPP Guard](#) on page 4006 is added to the document.

[Traps Reference](#) on page 2579 is moved to the Logs and Traps chapter.

[VLAN IDs](#) on page 4005 is added to the document.

Notice about Feature Support

This document includes content for multiple hardware platforms across different software releases. As a result, the content can include features not supported by your hardware in the current software release.

If a documented command, parameter, tab, or field does not display on your hardware, it is not supported. For information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

For information about physical hardware restrictions, see the hardware documentation.



Zero Touch Capabilities

[Auto-sense](#) on page 21

[Auto-sense Logical Flowcharts](#) on page 72

[IP Phone Support](#) on page 81

[Zero Touch Deployment](#) on page 91

[Zero Touch Provisioning Plus](#) on page 96

[Zero Touch Fabric Configuration](#) on page 101

[Example: Create an IS-IS Adjacency between Auto-sense and non-Auto-sense VOSS Switches](#) on page 104

The switch supports the following zero touch capabilities:

- Auto-sense support for the following features:
 - Fabric Attach (FA)
 - Fabric Extend (FE) in conjunction with ExtremeCloud SD-WAN Appliance
 - Extensible Authentication Protocol (EAP) and non-EAPoL (NEAP)
 - IP Phones
- Zero Touch Deployment
- Zero Touch Provisioning Plus (ZTP+)
- Zero Touch Fabric Configuration including Auto-sense for network-to-network interface (NNI)



Note

For bridged or routed reachability of the management servers (DHCP, RADIUS, ExtremeCloud IQ Site Engine, or ExtremeCloud IQ) the onboarding I-SID must be manually mapped to the management segment on at least one Backbone Edge Bridge (BEB) in the network prior to zero touch deployments of new switches. For more information, see *Fabric Engine Release Notes*.

Auto-sense

Table 5: Auto-sense product support

Feature	Product	Release introduced
Auto-sense	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3
Auto-sense can apply Fabric Attach (FA)-specific configuration	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4.2
	5520 Series	VOSS 8.4.2
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.4.2
	VSP 7400 Series	VOSS 8.4.2
Auto-sense can apply Fabric Extend-specific configuration	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.10.1
	5420 Series	Fabric Engine 8.10.1
	5520 Series	Fabric Engine 8.10.1
	5720 Series	Fabric Engine 8.10.1
	7520 Series	Fabric Engine 8.10.1
	7720 Series	Fabric Engine 8.10.1
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.10.1
	VSP 7400 Series	VOSS 8.10.1

Table 5: Auto-sense product support (continued)

Feature	Product	Release introduced
Auto-Sense MIB	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

Auto-sense is a port-based functionality that supports zero touch capabilities on the switch. Auto-sense dynamically configures the port to act as an IS-IS network-to-network interface (NNI), Fabric UNI (Flex-UNI), Fabric Attach (FA), Fabric Extend, or voice (IP phone) interface, based on the Link Layer Discovery Protocol (LLDP) events.

Auto-sense provides global configuration options for IS-IS authentication, FA authentication, Fabric Extend tunnel creation, and voice configuration for IP phones on the switch. For more information about IP Phone Support, see [IP Phone Support](#) on page 81.

When a switch boots in Zero Touch Fabric Configuration mode, all ports on the switch automatically operate in Auto-sense mode, unless you manually change the port configuration. For more information on Zero Touch Fabric Configuration, see [Zero Touch Fabric Configuration](#) on page 101.

With Auto-sense functionality, ports on a switch can detect whether they connect to a Shortest Path Bridging (SPB) device, an FA client, FA Proxy, Fabric Extend tunnels, Voice IP devices, or an undefined host:

- If a port connects to an SPB device or an FA client, then the system establishes Fabric architecture.
- If a port connects to any undefined host, then the system moves all untagged traffic on the port to an onboarding service network, also known as the onboarding I-SID.
- If a port operates in Auto-sense mode, Extensible Authentication Protocol (EAP) is enabled globally with a RADIUS configuration, and the Auto-sense port does not detect an SPB or Fabric Attach proxy neighbor, then the system automatically activates EAP and Non-EAP (NEAP) authentication on them, for untagged traffic.
- A single port in Auto-sense mode can learn Fabric Extend tunnels, tunnel state, and MTU from the SD-WAN Appliance through LLDP. For more information, see [Auto-sense Port States](#) on page 28.

When you manually disable Auto-sense on a specific port, the switch removes the dynamic configuration on that port unless you use an optional parameter to convert the dynamic configuration to a manual configuration. If you do not use the optional parameter, the software removes all Auto-sense state configuration and reverts the port to the default configuration.

If you enable Auto-sense on a port with a conflicting feature configuration, the software automatically deletes the conflicting configuration from the port. Conflicting configurations include the following commands or features:

- **access-diffserv** command
- **flex-uni enable** command
- **mac-security limit-learning** command
- **qos 802.1p-override enable** command
- Other feature configurations on the port:
 - brouter port
 - port tagging (encapsulation) - If a port has encapsulation enabled and you enable Auto-sense, the port remains with encapsulation enabled. Disabling Auto-sense transitions the encapsulation value to disabled. If a port has encapsulation disabled and you enable Auto-sense, port encapsulation is enabled.
 - Extensible Authentication Protocol over LAN (EAPoL)
 - FA
 - IS-IS interface
 - Link Aggregation Control Protocol (LACP) and Virtual Link Aggregation Control Protocol (VLACP)
 - LLDP disable (Auto-sense enables LLDP on the port but does not affect CDP dual mode)
 - LLDP MED network policies
 - private VLAN
 - MLT member
 - Switched UNI (S-UNI) or Transparent Port UNI (T-UNI) interface
 - VLAN member

**Note**

You cannot enable Auto-Sense if the `spbm-config-mode` flag is false.

Implementation on Upgraded Switches with Existing Configuration

If the switch does not boot in Zero Touch Fabric Configuration mode and you want to use Auto-sense functionality with an existing switch configuration, you must:

- Enable Auto-sense on the applicable port or ports.
- Create a new VLAN 4048. If the existing configuration uses VLAN 4048, you must configure a new VLAN for those original purposes.
- Configure I-SID 159999999 as the Auto-sense onboarding I-SID.
- Assign onboarding I-SID 159999999 to private VLAN 4048.

Implementation on 7830 Series

This section documents the differences in Auto-sense support on the 7830 Series from other switches.

- Auto-sense VOICE is not supported—Auto-sense port transitions to UNI state when it detects an LLDP packet from a phone.
- DHCP Server is not supported.
- EAP is not supported—Auto-sense ports in UNI and FA states do not have EAP and guest I-SID enabled; Auto-sense EAPOL max-mac commands are not available.
- Fabric Extend is not supported—Auto-sense port transitions to SD-WAN-PENDING state when it receives SD-WAN TLV in LLDP and creates the log "Auto-sense SD-WAN feature is not supported on 7830 Series".

For information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Auto-sense Data I-SID

Auto-sense supports global configuration of a data I-SID on the switch, which applies to all Auto-sense enabled ports. You can also configure Auto-sense data I-SID on each port to separate the data traffic into individual port specific data I-SIDs. For example, if device A and device B connect to different Auto-sense enabled ports and you configure an Auto-sense data I-SID on each port, the switch separates the data traffic of device A from the data traffic of device B. A port-level data I-SID and the global data I-SID can use the same value. The system prioritizes the I-SIDs in the following order:

1. Untagged I-SID assigned per client by EAP/NEAP MHMV
2. Untagged voice I-SID
3. Port data I-SID
4. Global data I-SID
5. Onboarding I-SID

The **show running-config** output includes the configured Auto-sense data I-SID for the port module only if you enable Auto-sense on the port. If you disable Auto-sense on the port, the configuration remains on the switch even though the command output does not include it. If you disable Auto-sense on the port and use the **convert-to-config** parameter, the port remains in the I-SID until you manually remove the data I-SID configuration from the port. If you re-enable Auto-sense on the port, you must reconfigure the data I-SID on the port.

If you remove the Auto-sense data I-SID from a port, then the port uses either the global Auto-sense data I-SID, if one exists, or the Auto-sense onboarding I-SID.

NEW! Auto-sense Guest I-SID Control

With this functionality, you can control how Auto-sense with Network Access Control (NAC) manages unauthenticated or rejected devices during onboarding. You can

disable the use of the onboarding I-SID as a guest I-SID. If you disable Guest-ISID control, the switch drops the traffic instead of sending it into the onboarding I-SID.

**Note**

This functionality applies to Auto-sense ports in UNI-ONBOARDING, FA, and VOICE states. It does not apply to FA-PROXY, FA-PROXY-NOAUTH, FA-PROXY-RING, or NNI states.

For information, see [Disable Auto-sense Guest I-SID](#) on page 44 using CLI or [Disable Auto-sense Guest I-SID](#) on page 60 using EDM.

IS-IS Authentication

Auto-sense supports global configuration of IS-IS authentication key on the switch. All ports operating in Auto-sense mode and transitioned to the NNI state, use the global IS-IS authentication key that you configure using the **auto-sense isis hello-auth type** command. For more information, see [Configure Auto-sense IS-IS Authentication](#) on page 39.

FA Configuration

Depending on the device that the Auto-sense port detects, the software can apply different FA-specific configurations that you define:

- You can configure an I-SID for FA clients such as FA wap-type 1, FA camera, and FA open-virtual-switch (OVS). The software prefers the FA I-SID over the onboarding I-SID.
- You can configure a specific I-SID and customer VLAN ID to use as the management I-SID when the port is in the Auto-sense FA PROXY or in the FA PROXY RING state.

If you do not configure a management I-SID when the port is in the FA PROXY state, the port uses the onboarding I-SID.

If you do not configure a management I-SID when the port is in the FA PROXY RING state, the port prioritizes the proxy I-SID, if configured, over the onboarding I-SID.

- You can disable EAPoL authentication requirements for specific FA client types (wap-type1, camera, and ovs).

FA Authentication

Auto-sense supports FA message authentication on switches. You can enable FA message authentication globally on a switch.

All ports operating in Auto-sense mode use the global authentication key. A preconfigured authentication key exists on the switch, by default, which you can change. For more information, see [Configure Auto-sense Fabric Attach \(FA\) Authentication](#) on page 46.

For additional security, you can enable Auto-sense to force FA authentication for attached devices in the FA PROXY NOAUTH port state. For more information, see [Force Fabric Attach Authentication](#) on page 47 .

Auto-sense Voice Capabilities

Auto-sense voice capabilities are based on the events when the switch detects an IP phone in the network. For more information, see [Auto-sense Voice](#) on page 83.

Loop Prevention

Auto-sense ports between two switches that have transitioned to NNI state are not prone to loops. Any connection can be wired and SPB establishes the shortest path connections. On Auto-sense NNI links BVID information, as well as IS-IS area information, is exchanged enabling Zero Touch Fabric functionality.

Auto-sense ports that connect to non-SPB switches operate in UNI mode, or FA Proxy mode in the case of ERS, EXOS, and Switch Engine switches. In UNI mode, Fabric Engine devices send Spanning Tree BPDU packets on auto-sense ports, emulating root bridge behavior ensuring that any potential UNI loop is broken by the attached spanning tree enabled devices. For universal hardware switches that transition from Switch Engine to Fabric Engine, there can be scenarios where certain links are spanning tree blocked.

For more information on the port states, see [Auto-sense Port States](#) on page 28.

Automatic Link Aggregation on NNI Ports

Auto-sense can support automatic link aggregation (Auto-MLT) for NNI links. When enabled, the switch advertises this ability to other Auto-sense devices in the Fabric Connect TLV of LLDP packets.

Table 6: Auto-MLT product support

Feature	Product	Release introduced
Auto-MLT on NNI ports	4220 Series	Fabric Engine 9.3
	5320 Series	Fabric Engine 9.3
	5420 Series	Fabric Engine 9.3
	5520 Series	Fabric Engine 9.3
	5720 Series	Fabric Engine 9.3
	7520 Series	Fabric Engine 9.3
	7720 Series	Fabric Engine 9.3
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.3
	VSP 7400 Series	VOSS 9.3

If you enable the Auto-MLT feature, Auto-sense ports automatically aggregate over a dynamic MLT if all of the following conditions are true:

- Two devices have at least two parallel links connected using Auto-sense ports.
- The physical links use the same speed.
- Either all or none of the links use IS-IS authentication; there is no mixed support.

**Note**

If you configure LACP on one side of the link but not the other, Auto-MLT for NNI does not work.

The dynamic MLT uses the highest MLT ID available. You cannot perform the following operations on the dynamic MLT:

- `interface mlt <ID>`
- `mlt <ID> member <slot/port>`
- `mlt <ID> name, private-vlan, or encapsulation`
- `no mlt <ID>`

If an Auto-MLT already exists for the neighbor, new ports are added to the MLT only if all ports use the same port speed. If you change the speed of a port in the Auto-MLT, the switch moves the port to a different MLT, associated with the new speed.

For information about Auto-sense port state transitions, see [Auto-sense Port States](#) on page 28.

**Note**

If you convert the dynamic Auto-sense configuration of one port in a dynamic MLT to a static configuration, all ports in that MLT are also converted to a static configuration and the MLT origin changes to CONFIG.

DHCP Port Snooping

For zero touch onboarding, if a port answers Dynamic Host Configuration Protocol (DHCP) requests sent by the switch and the port is in the Auto-sense UNI state, the system automatically changes the port's Private VLAN configuration from isolated mode to promiscuous mode.

Without this port type change on the Private VLAN, the other devices in the network cannot receive an IP address through the DHCP server if they are in the Zero Touch Fabric Configuration mode unless you disable Auto-sense on the port and manually change the port from isolated mode to promiscuous mode.

Running Configuration

If you view the running configuration, the global Auto-sense configuration displays under the port module. Use the command **`show running-config module port`**.

Auto-sense Port States

The system uses a per-interface state to adapt to all Auto-sense events. Each state transition determines background configuration on the port. The system does not display these configurations in the output of the **show running-config** command or in the saved configuration file but if you disable Auto-sense on the port and use the *convert-to-config* parameter, the dynamic configuration becomes a manual configuration and is visible in the **show running-config** output. Use **show auto-sense** commands to monitor the running states of each port.

For flowcharts that describe the system logic for Auto-sense port state detection, see [Auto-sense Logical Flowcharts](#) on page 72.

Port Down State

If you run the **auto-sense enable** command on a port that is disabled or has an inactive link, the port transitions to the Auto-sense Port Down state. This state transitions to the Auto-sense Wait state after the port becomes operational or the link becomes active.

Wait State

In the WAIT state, the port modifies outgoing LLDP packets to represent the enhanced properties of the port and analyzes incoming LLDP packets for possible transitions to advanced states like network-to-network interface (NNI), Fabric Attach (FA), Fabric Extend, or VOICE.



Note

If the port is in the WAIT state because you enabled **auto-sense no-nni**, the port ignores Fabric Connect TLVs received in LLDP packets.

If the port does not receive LLDP packets, the port transitions to the UNI state.

UNI State

This state grants onboarding and data connectivity to the port if you configure the onboarding I-SID, or a data I-SID in the global Auto-sense configuration or at the port level. If you do not use a guest I-SID, you can configure the switch so that Auto-sense does not create it. For information about applying this guest I-SID behavior, see [Disable Auto-sense Guest I-SID](#) on page 44 using CLI or [Disable Auto-sense Guest I-SID](#) on page 60 using EDM.

The system also applies the trusted and untrusted Auto-sense global configuration. As with the Wait state, the port continues to monitor received LLDP packets for transitions to other states.



Note

If the port is in the UNI state because you enabled **auto-sense no-nni**, the port ignores Fabric Connect TLVs received in LLDP packets.

Network Access Control (NAC) support, through EAP/NEAP, is enabled by default on each Auto-sense port, but disabled globally. If you require EAP/NEAP operation on Auto-sense ports, you must globally enable EAP and configure a RADIUS server.

The system performs the following background configurations.

```
flex-uni enable
eapol status auto
eapol multihost radius-non-eap-enable
eapol multihost eap-oper-mode mhm
[qos 802.1p-override enable]
[access-diffserv enable]
on port X interface, if onboarding I-SID Y is configured without data I-SID:
auto-sense guest-isid enable
eapol guest i-sid Y
on onboarding I-SID interface, if it is configured without data I-SID:
untagged-traffic port X
on data I-SID interface, if it is configured:
untagged-traffic port X
```

An Auto-sense port in the UNI state remains in PVLAN isolated mode when an additional untagged I-SID is applied to the port. Auto-sense ports support multiple VLAN/I-SIDs and PVLAN/I-SIDs on the same port concurrently. Typically, this operational mode is required when you configure NAC support with Multiple Host Multiple VLAN (MHMV). The software then assigns clients to their VLAN/I-SIDs based on their NAC authentication results.

Silent devices survive reboot as the switch stores authenticated client information such as MAC and IP address when the switch reboots. After the switch reboots, it sends regular ARP and ping messages to keep the device authenticated.

NNI States

The NNI states are as follows:

- NNI
- NNI-IS-IS-UP
- NNI-AUTH-FAIL
- NNI-MLT
- NNI-MLT-ISIS-UP



Note

If you enable **auto-sense no-nni** on the port, all NNI port states transition to the Wait state immediately, and after the timer expires, to the UNI state. The port ignores Fabric Connect TLVs received in LLDP packets and does not send Fabric Connect TLVs.

NNI and NNI-IS-IS-UP

If, while in the Wait state, the port receives a Fabric Connect LLDP packet, and the port is the first, and only, port to detect this neighbor, the port transitions to the NNI state and forms the adjacency as a single-link NNI.

The switch adds the IS-IS SPBM instance on the interface. The system tries to establish an IS-IS adjacency and, if successful, transitions the port to the NNI-IS-IS-UP state. The

port remains in the NNI-IS-IS-UP state until the adjacency fails, at which time it returns to the NNI state.

The system performs the following background configurations on port x:

```
isis
isis spbm 1
isis enable
[isis hello-auth ...] inherited from global configuration
```

If the port is the second port to detect this neighbor, see [NNI-MLT and NNI-MLT-ISIS-UP](#) on page 30.



Note

If you enable multiple ports for the same neighbor simultaneously, like with the same **no shutdown** command, the system moves directly to create the dynamic MLT and does not first form a single-link NNI.

NNI-AUTH-FAIL

If an Auto-sense port on a switch without an IS-IS Hello Authentication key connects to an Auto-sense port on another switch that does have an IS-IS Hello Authentication key, both ports transition to the NNI-AUTH-FAIL state. In the NNI-AUTH-FAIL state, the port has the following dynamic configuration:

- STP enabled
- spb-multihoming enabled
- port member in the onboarding VLAN (normal VLAN membership, no S-UNI)

Ports in the NNI-AUTH-FAIL state send an `onboarding_bit` flag in the Fabric Connect LLDP TLV to signal to a neighbor that the onboarding process started.

NNI-MLT and NNI-MLT-ISIS-UP

If you enable the Auto-MLT feature, ports transition to the NNI-MLT state after the FC TLV is received on two links. After the switch detects the parallel links, it takes the following actions:

- creates an MLT with origin AUTO-SENSE
- adds the ports to the MLT
- creates and enables the IS-IS interface on the MLT

After the IS-IS adjacency is formed on the MLT, the port state transitions to NNI-MLT-ISIS-UP.



Note

To successfully complete the Auto-MLT configuration, the ports must meet specific conditions. For more information, see [Automatic Link Aggregation on NNI Ports](#) on page 26.

If an Auto-sense port is a member of a statically-created MLT with the neighbor, that port also transitions to the NNI-MLT-ISIS-UP state. Ports that are members of a static MLT send this information in a `static_mlt` flag inside the Fabric Connect LLDP TLV.

If you disable the Auto-MLT feature, all ports return to the NNI state with single-link IS-IS adjacencies.

Fabric Attach (FA) States

The FA states are as follows:

- FA - this state is used for FA capable wireless access points, Camera, or OVS devices
- FA PROXY - this state is used for interaction with ERS and third-party switches, which are capable of FA proxy function and support authentication by default
- FA PROXY NOAUTH - this state is used for interaction with ERS, EXOS, and Switch Engine switches, which are capable of FA proxy function
- FA PROXY RING - this state is used for interactions with ISW-Series Managed Industrial Ethernet Switch (ISW-Series) switches with ring topologies, which are capable of FA proxy function and support authentication by default

LLDP uses the FA TLV to detect FA-capable neighbors.

When a port is in the FA state, the system uses the following priority for untagged traffic:

1. EAP/NEAP assigned I-SID
2. WAP, camera, or open virtual switch (OVS) I-SID
3. Onboarding I-SID
4. Drop

Depending on the device that the Auto-sense port detects, the switch can apply different FA-specific configurations that you define. For more information, see [Auto-sense](#) on page 21.

FA

The port enters the FA state after LLDP detects an access point, an FA client that is not another switch.

The system performs the following background configurations on port x:

```
flex-uni enable
eapol status auto
eapol multihost radius-non-eap-enable
eapol multihost eap-oper-mode mhm
eapol guest i-sid X
fa enable
on onboarding i-sid interface, if it exists:
untagged-traffic port X
```

FA PROXY

If LLDP detects an FA proxy switch such as an ERS, EXOS, or Switch Engine switch that uses FA message authentication, the port transitions to the FA PROXY state.

The system performs the following background configurations on port x:

```
flex-uni enable
fa enable
```

```
fa message-authentication
fa management-isid
```

**Note**

By default, the FA PROXY state uses the onboarding I-SID as the management I-SID but you can override this with a specific I-SID and customer VLAN ID combination.

FA PROXY NOAUTH

If the FA proxy switch does not use FA message authentication, the port transitions to the FA PROXY NOAUTH state.

The system performs the following background configurations on port x:

```
flex-uni enable
fa enable
on onboarding i-sid interface, if it exists:
untagged-traffic port X
```

Optionally, to secure the connection, you can enable Auto-sense to force FA authentication for devices that attach in this port state. After you enable this feature, the port transitions to the FA PROXY state. FA authentication fails until you enable it on the attached device. After you enable FA authentication on the attached device, the system dynamically configures the FA PROXY management I-SID and tags traffic in the associated VLAN. For more information, see [FA PROXY](#) on page 31.

FA PROXY RING

If LLDP detects an (ISW-Series) switch with ring topologies that uses FA message authentication, the port transitions to the FA PROXY RING state. As a result, FA and FA Topology Change Notification (TCN) can process TCN BPDUs received from the ISW switch. By default, the FA PROXY RING state uses the onboarding I-SID as the management I-SID but you can override this with a specific I-SID and customer VLAN ID combination.

The system performs the following background configurations on port x:

```
flex-uni enable
fa enable
fa authentication-key
fa message-authentication
fa management-isid x c-vid y
```

Fabric Extend (FE) States

When Auto-sense is enabled, LLDP uses the FE TLV to create Fabric Extend tunnels between two Fabric switches that connect over the Internet through the SD-WAN Appliance. This functionality is supported on a single port of the switch. For more information, see [SD-WAN](#) on page 3304.

The FE states are as follows:

- SD-WAN
- SD-WAN-PENDING

SD-WAN

After the first Auto-sense port receives an FE-TLV, the port transitions to the SD-WAN state. All other Auto-sense ports transition to SD-WAN-PENDING state and remain unconfigured. When the first port transitions to the SD-WAN state, the switch verifies that VLAN 4047, VRF, and IS-IS logical interface configurations do not exist, and dynamically configures the following connectivity parameters:

- SD-WAN as the VLAN name associated with VLAN 4047 with origin ZTF
- sd-wan as the VRF name associated with the IP tunnel with origin DYNAMIC



Note

On switch models that support a single active VRF, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists but you can manually specify the VRF name that Auto-sense uses for the SD-WAN configuration on these models. For more information, see [SD-WAN](#) on page 3304.

- SD-WAN-<ifidx> as the tunnel name
- SD-WAN Tunnel SrcIP as the name associated with the Fabric Extend underlay IP
- IPv4 address for VLAN 4047
- Static host route (/32) for each Fabric Extend tunnel destination to reach the next-hop SD-WAN Appliance
- Fabric Extend tunnels with origin ZTF for IS-IS logical interfaces
- VLAN 4047 port membership
- Link Debounce timer of 8000 milliseconds on the switch port that connects to SD-WAN Appliance, if a timer configuration does not already exist
- Mgmt-sdwan interface created and enabled on VLAN 4047, with the subnet route to the ExtremeCloud SD-WAN Appliance installed for the interface

The management interface uses the same IP address that is assigned to VLAN 4047. To confirm the management IP address of the connected SD-WAN Appliance, use the **show lldp neighbor** command. Management applications, such as the SSH client, can use this interface to reach the SD-WAN Appliance. The switch deletes this management interface and removes the IP address after any of the following occur:

- Disable the port either administratively or operationally.
- Disable Auto-sense.
- LLDP timeout.

SD-WAN-PENDING

In the following cases, the port transitions to the SD-WAN-PENDING state:

- A secondary Auto-sense port receives an FE-TLV.
- The switch configuration includes the dynamic connectivity parameters, such as VLAN 4047, VRF, and IS-IS logical interfaces with the specified source IP address regardless of origin.

- You enable **auto-sense no-nni** on the port.
- An SD-WAN LLDP is received on a switch with either a global or VRF BGP configuration and with 4-byte AS numbers disabled.

Voice State

If the port detects an LLDP packet from a phone, the port transitions to the VOICE state. A global Auto-sense voice configuration is not required to transition to the VOICE state except a specific voice VLAN shall be signaled to the phone.

For more information on Auto-sense voice, see [Auto-sense Voice](#) on page 83.

Auto-sense Configuration using CLI

To change the Auto-sense configuration on a port using EDM, see [Configure Basic Port Parameters](#) on page 755.



Note

Auto-Sense cannot be enabled when the spbm-config-mode flag is false.

Enable Auto-sense on Port(s)

About This Task

Perform this procedure to manually enable Auto-sense on a specific port.



Note

After a switch boots without a configuration file, Auto-sense is enabled on all ports, by default.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
`enable`

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable Auto-sense on the port:
`auto-sense enable`

Example

Enabling Auto-sense on port 1/2:

```
Switch:1>enable  
Switch:1#configure terminal
```



```
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#auto-sense enable
Warning: Enabling Auto-Sense will default port configurations.
```

Disable Auto-sense on Port(s)

About This Task

Perform this procedure to disable Auto-sense on a specific port. You can disable Auto-sense on the port but retain the configuration that the system applied dynamically.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable Auto-sense on the port:

```
no auto-sense enable [convert-to-config]
```

Example

Disable Auto-sense on port 1/2 but retain the configuration. The dynamic configuration becomes a manual configuration and is visible in the **show running-config** output and can be saved to the configuration file using the **save config** command.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface interface gigabitEthernet 1/2
Switch:1(config-if)#no auto-sense enable convert-to-config
Switch:1(config-if)#save config
```

Variable Definitions

The following table defines parameters for the **no auto-sense enable** command.

Variable	Value
<i>convert-to-config</i>	Retains the Auto-sense configuration that the system applies dynamically on the specific port. The dynamic configuration becomes a manual configuration and is visible in the show running-config output. If you run the no auto-sense enable command without the <i>convert-to-config</i> parameter, then the configuration will be removed from the port and the port returns to the default state where VLAN 1 is assigned.

Configure Auto-sense IS-IS Level 1 Metric

Perform this procedure to configure a global IS-IS Level 1 metric for ports that operate in Auto-sense mode.

About This Task

For the command to take effect, at least one adjacency must be formed through Auto-sense.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the configure the Level 1 metric on Auto-sense ports:

```
auto-sense isis ll-metric {<1-16777215>|auto}
```

Variable Definitions

The following table defines parameters for the **auto-sense isis ll-metric** command.

Variable	Value
<code><1-16777215></code>	Configures a static Level 1 metric for the Auto-sense ports. The higher the value, then the lower the cost for a route. The default value is 10.
<code>auto</code>	Enable auto if you want the network route determined by summing the lowest value metrics, which are inversely proportional to port speed. The faster port speed is the lowest cost and the preferred route for all Auto-sense ports.

Configure the Auto-sense Wait Interval Globally

Perform this task to configure the time, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the Auto-sense wait interval:

```
auto-sense wait-interval <2-120>
```
3. Verify the Auto-sense wait interval information:

```
show auto-sense wait-interval
```

Examples

Configure the Auto-sense wait interval as 50 seconds:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense wait-interval 50
```

Verify the Auto-sense wait-interval information:

```
Switch:1>show auto-sense wait-interval
=====
                        AUTO-SENSE GLOBAL Config
=====
WAIT
INTERVAL
-----
50
-----

0 out of 0 Total Num of AUTO-SENSE entries displayed
=====
```

Variable Definitions

The following table defines parameters for the **auto-sense wait-interval** command.

Variable	Value
<2-120>	Specifies the global Auto-sense wait interval, in seconds. The global configuration applies to all Auto-sense ports unless you configure a port specific wait-interval. The default value is 35.

Configure the Auto-sense Wait Interval for a Port

Perform this task to configure the time, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state.

About This Task

The per-port wait-interval configuration takes precedence over the globally configured Auto-sense wait-interval. If you change the port configuration and the port is already in the wait state, the wait-interval adjusts automatically.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the wait interval on a port:

```
auto-sense wait-interval <10-120>
```

3. Verify the Auto-sense wait interval information:

```
show interfaces gigabitethernet auto-sense {slot/port[-slot/port]
[,...][slot/all][all]}
```

Examples

Configure the Auto-sense wait interval on port 1/1 to 20 seconds:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#interface GigabitEthernet 1/1
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense wait-interval 20
```

Verify the Auto-sense wait-interval information for a range of ports:

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet auto-sense 1/1-1/5
=====
Port Auto-sense
=====
```

PORT NUM	AUTO-SENSE STATUS	AUTO-SENSE STATE	AUTO-SENSE PORT-DATA-ISID	AUTO-SENSE PORT-WAIT-INTERVAL	AUTO-SENSE NO-NNI
1/1	Enable	FA-PROXY-RING	--	20	Disable
1/2	Enable	SD-WAN	--	10	Disable
1/3	Enable	DOWN	--	10	Disable
1/4	Enable	DOWN	--	10	Disable
1/5	Disable	OFF	--	10	Disable

Variable Definitions

The following table defines parameters for the **auto-sense wait-interval** command.

Variable	Value
<10-120>	Specifies the Auto-sense wait interval, in seconds, for a specific port. The port-specific configuration takes precedence over the global wait-interval configuration. The default value is 35.

Configure Auto-sense IS-IS Authentication

Before You Begin

Enable IS-IS globally.

About This Task

Perform this procedure to configure a global IS-IS authentication key for ports that are operating in Auto-sense mode.



Note

If the IS-IS authentication keys on auto-sense ports between two switches do not match, then the auto-sense port state will be auto-sense UNI onboarding, until the keys are matching, then an IS-IS adjacency will be established.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the authentication type for IS-IS hello packets on Auto-sense ports:

```
auto-sense isis hello-auth type {none|simple|hmac-md5|hmac-sha-256}  
[key WORD<1-16>] [key-id <1-255>]
```

Example

Configuring simple authentication for IS-IS hello packets on Auto-sense ports:

```
Switch:1>enable  
Switch:1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#auto-sense isis hello-auth type simple key Secure
```

Variable Definitions

The following table defines parameters for the **auto-sense isis hello-auth type** command.

Variable	Value
<code>{none simple hmac-md5 hmac-sha-256}</code>	<p>Specifies the authentication type for IS-IS hello packets on Auto-sense ports:</p> <ul style="list-style-type: none"> • none • simple - simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5 - MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. • hmac-sha-256 - with SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default authentication type is none.</p>
<code>key WORD<1-16></code>	Specifies the authentication key (password) used by the receiving router to verify the packet.
<code>key-id <1-255></code>	Specifies the key ID.

Configure Auto-sense Access Ports

About This Task

Perform this procedure to configure ports operating in Auto-sense mode to determine the Layer 3 Quality of Service (QoS) actions the switch performs. The Auto-sense access ports override the Differentiated Services Code Point (DSCP) markings.

Procedure

1. Enter Global Configuration mode:


```
enable
```

```
configure terminal
```
2. Configure Auto-sense access ports:


```
auto-sense access-diffserv [enable]
```

Example

Configure the Auto-sense access ports:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#auto-sense access-diffserv enable
```

Variable Definitions

The following table defines parameters for the **auto-sense access-diffserv** command.

Variable	Value
<i>enable</i>	Configures the ports operating in Auto-sense mode to determine the Layer 3 Quality of Service (QoS) actions the switch performs. The Auto-sense access ports override the Differentiated Services Code Point (DSCP) markings. The default configuration is enabled.

Disable Auto-sense DHCP Server Detection

About This Task

Perform this procedure to disable Dynamic Host configuration Protocol (DHCP) server detection in Auto-sense mode.



Note

By default Auto-sense DHCP server detection is enabled. This ensures automatic detection of the DHCP uplink ports in Zero Touch Deployment.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Disable DHCP server detection:
`no auto-sense dhcp-detection`

Example

Enable DHCP server detection:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no auto-sense dhcp-detection
```

*Configure the Auto-sense Onboarding I-SID on the Switch***About This Task**

Perform this procedure to configure the onboarding I-SID for ports that are operating in Auto-sense mode. The onboarding I-SID is typically used to onboard networking devices such as switches and non FA capable access points. By default, the onboarding I-SID provides automatic reachability when switches are booted from factory without a configuration file. For security reasons, the onboarding I-SID forms an isolated PVLAN/ETREE to block any unwanted port to port cross talk.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the onboarding I-SID:

```
auto-sense onboarding i-sid <1-159999999>
```

Example

Configuring the Auto-sense onboarding I-SID:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense onboarding i-sid 15000000
```

Variable Definitions

The following table defines parameters for the **auto-sense onboarding** command.

Variable	Value
<i>i-sid<1-159999999></i>	Specifies the service instance identifier (I-SID). The default onboarding I-SID value is 159999999.

*Configure a Auto-sense Global Data I-SID***Before You Begin**

- Enable Auto-sense on the port.
- Associate a VLAN with the I-SID before you configure it as the global data I-SID.

About This Task

Perform this task to configure Auto-sense data traffic information for ports that are operating in Auto-sense mode.

**Note**

This option applies to the auto-sense UNI and voice states only, it replaces the onboarding I-SID and places an (untagged) client device into a pre-defined global data I-SID.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the data service instance identifier (I-SID):
`auto-sense data i-sid <1-15999999>`

Example

Configuring the Auto-sense data I-SID:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense data i-sid 1000
```

Variable Definitions

The following table defines parameters for the **auto-sense data** command.

Variable	Value
<code>i-sid <1-15999999></code>	Specifies the service instance identifier (I-SID).

Configure an Auto-sense Port Data I-SID

Before You Begin

- Enable Auto-sense on the port.
- Associate a VLAN with the I-SID before you configure it as the data I-SID on the port. This does not apply to a DvR leaf.

About This Task

Perform this procedure to configure a data I-SID on a port.



Note

This option applies to the Auto-sense UNI and voice states only, it replaces the onboarding I-SID and places an (untagged) client device into a pre-defined port specific data I-SID.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the Auto-sense port data I-SID:

```
auto-sense data i-sid <1-15999999>
```

Example

Configuring the Auto-sense data I-SID for ports 1/1 to 1/5:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/1-1/5  
Switch:1(config-if)#auto-sense data i-sid 15000000
```

Variable Definitions

The following table defines parameters for the **auto-sense data** command.

Variable	Value
<i>i-sid</i> <1-15999999>	Specifies the service instance identifier (I-SID).

NEW! *Disable Auto-sense Guest I-SID*

About This Task

Perform this procedure to disable guest I-SID detection in Auto-sense mode when there is no response from the RADIUS server.



Note

If the switch boots in Zero Touch Deployment mode, the Auto-sense guest I-SID is disabled. If the switch boots with the factorydefaults boot flag enabled, the Auto-sense guest I-SID is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable the Auto-sense guest I-SID:
`no auto-sense guest-isid enable`

Configure Layer 2 Trusted Auto-sense Ports

About This Task

Perform this procedure to override incoming 802.1p bits on ports that operate in Auto-sense UNI or voice mode.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure Auto-sense ports as Layer 2 trusted:
`auto-sense qos 802.1p-override`

Configure EAPoL Authentication Requirements for Auto-sense Fabric Attach Clients

You can disable EAPoL authentication for specific Fabric Attach (FA) client types.

About This Task

By default, authentication is required before the connection is authorized.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure EAPoL authentication requirements from the following choices:
 - For auto-sensed cameras: `auto-sense fa camera eapol status {authorized | auto}`
 - For auto-sensed virtual switches: `auto-sense fa ovs eapol status {authorized | auto}`
 - For auto-sensed wireless access points (WAP): `auto-sense fa wap-type1 eapol status {authorized | auto}`

Variable Definitions

The following table defines parameters for the **auto-sense** commands related to EAPoL authentication for Fabric Attach (FA).

Variable	Value
<code>{authorized auto}</code>	<p>Configures the EAPoL authentication requirement for the specific client type. Choose from the following options:</p> <ul style="list-style-type: none"> <code>authorized</code>— the port skips EAPoL authentication and authorizes the connection. <code>auto</code>— authorization depends on the result of EAPoL authentication. <p>By default, authentication is required before the connection is authorized.</p>

Configure Auto-sense Fabric Attach (FA) Authentication

About This Task

Perform this procedure to configure FA authentication for ports that are operating in Auto-sense mode.



Important

To increase security for devices that attach in the FA PROXY NOAUTH port state, you can force FA authentication. For more information, see [Force Fabric Attach Authentication](#) on page 47 .

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the FA authentication key:

```
auto-sense fa authentication-key WORD<0-32>
```
3. Enable FA message authentication:

```
auto-sense fa message-authentication
```

Example

Configuring FA message authentication globally:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense fa message-authentication
```

Variable Definitions

The following table defines parameters for the **auto-sense fa** command.

Variable	Value
<i>authentication-key</i> <i>WORD<0-32></i>	Specifies the authentication key value.
<i>message-authentication</i>	Enables Fabric Attach (FA) message authentication globally, for ports that operate in Auto-sense mode.

Force Fabric Attach Authentication

About This Task

To secure the FA connection, enable Auto-sense to force FA authentication for devices that attach in the FA PROXY NOAUTH port state. If you enable force-authentication, the Auto-sense port transitions to the FA-PROXY state, even if it receives FA PROXY NOAUTH.

By default, forcing authentication is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Enable Auto-sense to force authentication:

```
auto-sense fa proxy-no-auth force-authentication
```

Configure an I-SID for Auto-sense Fabric Attach Clients

For Zero Touch Deployment and assignments of dedicated I-SIDs for FA capable cameras, Wireless Access Points, FA proxy switches and Open Virtual Switches (OVS), configure a specific I-SID to use instead of the onboarding I-SID when a port is in an Auto-sense Fabric Attach (FA) state and detects an FA client.

Before You Begin

- Create the I-SID.
- Associate the I-SID with either a platform or private VLAN; this association is not required on a DvR Leaf.

About This Task

You can create only one I-SID of each type.

The FA I-SID can be the same as the voice I-SID because they are used by different Auto-sense port states.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the FA I-SID from the following choices:

- For auto-sensed cameras: `auto-sense fa camera i-sid <1-159999999>`
- For auto-sensed FA client switches that do not use FA message authentication, like EXOS or Switch Engine: `auto-sense fa proxy-no-auth i-sid <1-159999999>`



Important

To increase security, rather than allowing ports to remain in the FA PROXY NOAUTH port state for these devices, you can force FA authentication. For more information, see [Force Fabric Attach Authentication](#) on page 47.

- For auto-sensed virtual switches: `auto-sense fa ovs i-sid <1-159999999>`
- For auto-sensed wireless access points (WAP): `auto-sense fa wap-type1 i-sid <1-159999999>`

Variable Definitions

The following table defines parameters for the **auto-sense** commands related to Fabric Attach (FA) I-SIDs.

Variable	Value
<code>i-sid <1-159999999></code>	Specifies the service instance identifier (I-SID).

Configure a Management I-SID for Auto-sense Fabric Attach Proxy Switches

About This Task

Perform the following procedure to configure a specific I-SID and customer VLAN ID to use as the management I-SID when a port is in the Auto-sense FA PROXY state.

The switch creates this I-SID dynamically and uses it instead of the onboarding I-SID.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the management I-SID:

```
auto-sense fa proxy management i-sid <1-159999999> c-vid <1-4094>
```

Variable Definitions

The following table defines parameters for the **auto-sense fa proxy management** command.

Variable	Value
<i>c-vid</i> <1-4094>	Specifies the customer VLAN ID.
<i>i-sid</i> <1-15999999>	Specifies the service instance identifier (I-SID).

Configure a Management I-SID for Auto-sense Fabric Attach Proxy Ring Switches

About This Task

Perform the following procedure to configure a specific I-SID and customer VLAN ID to use as the management I-SID when a port is in the Auto-sense FA PROXY RING state.

The switch creates this I-SID dynamically and uses it instead of the proxy I-SID or the onboarding I-SID.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the management I-SID:

```
auto-sense fa proxy-ring management i-sid <1-15999999> c-vid <1-4094>
```

Variable Definitions

The following table defines parameters for the **auto-sense fa proxy management** command.

Variable	Value
<i>c-vid</i> <1-4094>	Specifies the customer VLAN ID.
<i>i-sid</i> <1-15999999>	Specifies the service instance identifier (I-SID).

Display Auto-sense Configuration on the Switch

About This Task

Perform this procedure to display the Auto-sense configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the Auto-sense configuration:

```
show auto-sense [access-differv] [auto-mlt] [data] [dhcp-detection]
[eapol] [fa] [guest-isid] [isis] [onboarding] [qos] [sd-wan] [voice]
[wait-interval]
```

3. Display the Auto-sense status and state on a port:

```
show interfaces gigabitEthernet auto-sense [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

Examples

Display the Auto-sense status for Fabric Attach (FA):

```
Switch:1>show auto-sense fa
=====
                        AUTO-SENSE FA Config
=====
MSG-AUTH                MSG-AUTH-KEY                PROXY-FORCE-AUTH
-----
enabled                 ****                 disabled
=====

                        AUTO-SENSE FA Client specific config
=====
TYPE                    EAPOL STATUS    I-SID    VLANID    C-VID    MGMT I-SID MGMT C-VID
-----
camera                  Auto            100      100       untag    -        -
wap-type1               Auto            200      200       untag    -        -
open-virtual-switch     Auto            -        -         -        -        -
proxy-no-auth           Auth            300      300       untag    -        -
proxy                   Auth            400      n/a       400      400      400
=====
6 out of 6 Total Num of AUTO-SENSE entries displayed
=====
```

Display the Auto-sense configuration related to ExtremeCloud SD-WAN:

```
Switch:1>show auto-sense sd-wan
=====
                        AUTO-SENSE GLOBAL Config
=====
SDWAN      SDWAN
MULTI-AREA VRF
=====
HOME
=====

                        AUTO-SENSE SD-WAN Logical Interfaces Multi-Area Config
=====
DEST-IP      MULTI-AREA
=====
10.10.10.10   REMOTE
20.20.20.20   HOME
30.30.30.30   REMOTE
=====
4 out of 4 Total Num of AUTO-SENSE entries displayed
=====
```

Display the Auto-sense configuration related to voice:

```
Switch:1>show auto-sense voice
=====
                        AUTO-SENSE VOICE Config
=====
TYPE  LDDP-AUTH ENABLE I-SID    C-VID    DSCP    PRIORITY
```



```

-----
phone  FALSE           2000      2000      46        6
-----
1 out of 1 Total Num of AUTO-SENSE entries displayed
-----

```

Display the global Auto-sense wait-interval information:

```

Switch:1>show auto-sense wait-interval
=====
                        AUTO-SENSE GLOBAL Config
=====
WAIT
INTERVAL
-----
50
-----

0 out of 0 Total Num of AUTO-SENSE entries displayed
-----

```

Display the Auto-sense status and state on a range of ports:

```

Switch:1>enable
Switch:1#show interfaces gigabitethernet auto-sense 1/1-1/5
=====
                        Port Auto-sense
=====
-----
PORT      AUTO-SENSE  AUTO-SENSE  AUTO-SENSE  AUTO-SENSE  AUTO-SENSE
NUM       STATUS    STATE      PORT-DATA-ISID  PORT-WAIT-INTERVAL NO-NNI
-----
1/1       Enable     FA-PROXY-RING  --           20           Disable
1/2       Enable     SD-WAN        --           10           Disable
1/3       Enable     DOWN          --           10           Disable
1/4       Enable     DOWN          --           10           Disable
1/5       Disable    OFF           --           10           Disable

```

Display the Auto-sense configuration related to the guest I-SID:

```

Switch:1#show auto-sense guest-isid
=====
                        GUEST-ISID
=====
ENABLED
-----
TRUE
-----

0 out of 0 Total Num of AUTO-SENSE entries displayed
-----

```

*Configure Maximum EAP Clients on Auto-sense Ports***About This Task**

Use this procedure to configure the maximum EAP clients allowed on the Auto-sense enabled port at one time.

**Note**

If you manually configure values on a specific port(s), then these values will take precedence over the Auto-sense global values.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the maximum EAP clients on Auto-sense enabled ports:

```
auto-sense eapol multihost eap-mac-max <0-32>
```

**Note**

Eapol multihost eap-mac-max also provides the EAP and NEAP separation functionality. By default, the system enables EAP clients on each port with a maximum limit of 2 EAP clients. If you configure the maximum limit to 0 then the system disables the EAP client authentication.

Example

Configure a maximum of 10 EAP clients allowed on the Auto-sense enabled port at one time.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense eapol multihost eap-mac-max 10
```

Variable Definitions

The following table defines parameters for the **auto-sense eapol multihost eap-mac-max** command.

Variable	Value
<0-32>	Specifies the maximum EAP clients allowed on the Auto-sense enabled port at one time. The default is 4.

Configure Maximum MAC Clients on Auto-sense Ports

About This Task

Use this procedure to configure the maximum EAP and NEAP MAC clients supported on Auto-sense enabled ports.



Note

If you manually configure values on a specific port(s), then these values will take precedence over the Auto-sense global values.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the maximum limit of allowed EAP and NEAP MAC clients supported on the Auto-sense enabled port:
`auto-sense eapol multihost mac-max <1-8192>`

Example

Configure a maximum of 10 EAP and NEAP MAC clients supported on Auto-sense enabled ports.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense eapol multihost mac-max 10
```

Variable definitions

The following table defines parameters for the **auto-sense eapol multihost mac-max** command.

Variable	Value
<code>mac-max <1-8192></code>	Specifies the maximum number of EAP and NEAP MAC addresses allowed on the Auto-sense enabled port. The default is 4.

Configure Maximum NEAP Clients on Auto-sense Ports

About This Task

Use this procedure to configure the maximum NEAP clients allowed on the Auto-sense enabled port at one time.



Note

If you manually configure values on a specific port(s), then these values will take precedence over the Auto-sense global values.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the maximum NEAP clients allowed on Auto-sense enabled ports:

```
auto-sense eapol multihost non-eap-mac-max <0-8192>
```



Note

auto-sense eapol multihost eap-mac-max also provides the EAP and NEAP separation functionality. By default, the system enables NEAP clients on each port with a maximum limit of 2 NEAP clients. If you configure the maximum limit to 0, then the system disables the NEAP client authentication.

Example

Configure a maximum of 10 NEAP clients allowed on the Auto-sense enabled port at one time.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense eapol multihost non-eap-mac-max 10
```

Variable Definitions

The following table defines parameters for the **auto-sense eapol multihost non-eap-mac-max** command.

Variable	Value
<0-8192>	Specifies the maximum NEAP clients allowed on the Auto-sense enabled port at one time. The default is 4.

Display Maximum EAP and NEAP Clients Configured on Auto-sense Enabled Ports

About This Task

Perform this procedure to display maximum number of EAP and NEAP clients configured on Auto-sense enabled ports.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display maximum number of EAP and NEAP clients supported on Auto-sense enabled ports:

```
show auto-sense eapol multihost
```

Example

Display maximum number of EAP and NEAP clients configured on Auto-sense enabled ports.

```
Switch:1>show auto-sense eapol multihost
=====
                        AUTO-SENSE EAPOL MULTIHOST Config
=====
MAC-MAX          EAP-MAC-MAX      NEAP-MAC-MAX
-----
4                 4                4
-----

1 out of 1 Total Num of AUTO-SENSE entries displayed
=====
```

Prevent a Port From Entering SD-WAN and NNI Port States

About This Task

To increase security when these Auto-sense port states are not required, you can block the port from sending Fabric Connect LLDP TLVs and prevent the port from transitioning to the SD-WAN and NNI states.

By default, this option is disabled.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Block the port from sending Fabric Connect LLDP TLVs:

```
auto-sense no-nni
```

Enable Auto-MLT on NNI Ports

About This Task

Enable automatic link aggregation (Auto-MLT) to automatically aggregate NNI links on Auto-sense ports into an MLT to form IS-IS adjacencies. The default is enabled in Zero Touch Deployment mode, when the switch has no configuration file. Otherwise, the default is disabled.

The switch aggregates the NNI links into an MLT if a minimum of two links with same speed are present between two devices. If only one link exists, the switch creates a single-port NNI on the Auto-sense ports.

If you enable this feature, the switch advertises this ability inside the Fabric Connect LLDP TLV using an `auto-lag-nni-capability` flag.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable Auto-MLT:

```
auto-sense auto-mlt enable
```

3. Verify the configuration:

```
show auto-sense auto-mlt
```

NEW! *Configure Link Debounce for Auto-sense Ports*



Note

This procedure does not apply to 7830 Series.

About This Task

Perform this procedure to configure Link Debounce on Auto-sense UNI ports. Link Debounce maintains a port link even if it is physically down. The port link remains up for the timer you configure. If the link is not restored within the timer period, the switch announces the port as down.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure Link Debounce:

```
auto-sense link-debounce {auto <0-300000> | off | on <0-300000>}
```

3. Verify the configuration:

```
show auto-sense link-debounce
```

NEW! Variable Definitions

The following table defines parameters for the **auto-sense link-debounce** command.

Variable	Value
auto <0-300000>	Enables Link Debounce only on Auto-sense UNI ports that connect to Preboot Execution Environment (PXE) devices. The auto parameter enables Link Debounce if the port receives a DHCP DISCOVERY message with an Option 60 vendor class ID string of PXEClient:Arch:xxxx:UNDI:yyyyzzz. The auto parameter is the default configuration on Auto-sense ports. <0-300000> specifies the port timer in milliseconds. The default timer is 3,000 milliseconds.
off	Disables Link Debounce on Auto-sense UNI ports.
on <0-300000>	Enables Link Debounce on all Auto-sense UNI ports. <0-300000> specifies the port timer in milliseconds. The default timer is 3,000 milliseconds.

Auto-sense Global Configuration using EDM

Perform the procedures in this section to configure Auto-sense globally using EDM. You can perform many tasks on the same EDM tab. For a quick look at field descriptions for that tab, see [Globals Field Descriptions](#) on page 65.



Note

Auto-Sense cannot be enabled when the spbm-config-mode flag is false.

Enable LLDP Authentication of IP Phones

Before You Begin

You must enable EAPoL globally.

About This Task

Perform this procedure to enable Link Layer Discovery Protocol (LLDP) authentication of IP phones. The switch authenticates the phone after it receives LLDP packets from the phone.

Auto-sense LLDP authentication applies to Auto-sense ports in the VOICE state. Auto-sense LLDP authentication does not require a global Auto-sense voice configuration.

The system removes the LLDP session for the following reasons:

- You disable EAPoL globally.
- You disable Auto-sense on the port.
- The LLDP neighbor is removed.

If the LLDP authentication configuration exists and one of the following situations occur, the LLDP session is recreated:

- You enable EAPoL globally.
- You enable Auto-sense on the port.
- The LLDP neighbor is recreated.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **EapolVoiceLldpAuthEnable**, to enable the EAPoL LLDP authorization for voice Auto-sense ports.
5. Select **Apply**.

Configure Auto-sense Voice Information for IP Phones

The switch applies the Auto-sense voice configuration on specific port(s), after it discovers IP phones on the port through LLDP packets.

Before You Begin

If you boot the switch with a configuration file, and not through Zero Touch Fabric Configuration, you must manually enable Auto-sense on specific port(s).

About This Task

Perform this procedure to configure Auto-sense voice information for IP phones. A global Auto-sense voice configuration does not require Auto-sense LLDP authentication.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. For **VoiceIid**, type the I-SID value.
5. For **VoiceCvid**, type the CVID value associated with the voice I-SID.
6. Select **Apply**.

Disable Auto-sense DHCP Server Detection

About This Task

Perform this procedure to disable DHCP server detection in Auto-sense mode.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **DhcpDetection** to disable DHCP detection.
5. Select **Apply**.

*Configure Auto-sense Onboarding I-SID Globally***About This Task**

Perform this procedure to configure the onboarding I-SID for ports that are operating in Auto-sense mode.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. For **OnboardingIsid**, type I-SID value for the Auto-sense ports.
5. Select **Apply**.

Configure the Auto-sense Wait Interval Globally

Perform this task to configure the time, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Type a value for **WaitInterval** to configure the wait time interval for all Auto-sense enabled ports.
5. Select **Apply**.

*Configure Auto-sense Data I-SID Globally***Before You Begin**

- Enable Auto-sense on the port.
- Associate a VLAN with the I-SID before you configure it as the global data I-SID.

About This Task

Perform this task to configure Auto-sense data traffic information for ports that are operating in Auto-sense mode.



Note

This option applies to the Auto-sense UNI and voice states only; it replaces the onboarding I-SID and places an (untagged) client device into a pre-defined global data I-SID.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. For **Datalsid**, type the data I-SID value used by the Auto-sense ports.
5. Select **Apply**.

NEW!

Disable Auto-sense Guest I-SID

About This Task

Perform this procedure to disable guest I-SID detection in Auto-sense mode when there is no response from the RADIUS server.



Note

If the switch boots in Zero Touch Deployment mode, the Auto-sense guest I-SID is disabled. If the switch boots with the factorydefaults boot flag enabled, the Auto-sense guest I-SID is enabled.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Clear **GuestIsidEnable**.
5. Select **Apply**.

Configure Layer 2 Trusted Auto-sense Ports

About This Task

Perform this procedure to override incoming 802.1p bits on ports that operate in Auto-sense mode.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.

4. Select **Qos8021pOverrideEnable** to override incoming 802.1p bits on ports that operate in Auto-sense mode.
5. Select **Apply**.

Configure Auto-sense IS-IS Authentication

About This Task

Perform this procedure to configure a global IS-IS authentication key for ports that are operating in Auto-sense mode.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. For **IsisHelloAuthType**, select a type of IS-IS hello authentication.
5. For **IsisHelloAuthKeyId**, type the key ID for IS-IS authentication for the Auto-sense ports.
6. For **IsisHelloAuthKey**, type the key for IS-IS authentication for the Auto-sense ports.
7. Select **Apply**.

Configure Auto-sense Access Ports

About This Task

Perform this procedure to configure ports operating in Auto-sense mode to determine the Layer 3 QoS actions the switch performs. The Auto-sense access ports override the Differentiated Services Code Point (DSCP) markings.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **AccessDiffservEnable** to enable differentiated serve type as access for Auto-sense ports.
5. Select **Apply**.

Configure Auto-sense for Fabric Attach

Perform this procedure for the following purposes:

- Configure Fabric Attach (FA) authentication for ports that are operating in Auto-sense mode.
- For Zero Touch Deployment and assignments of dedicated I-SIDs for FA capable cameras, Wireless Access Points, FA proxy switches and Open Virtual Switches (OVS), you can configure a specific I-SID to use instead of the onboarding I-SID when a port is in an Auto-sense Fabric Attach (FA) state and detects an FA client.
- Configure a specific I-SID and customer VLAN ID to use as the management I-SID when a port is in the Auto-sense FA PROXY or in the FA PROXY RING state.

Before You Begin

- Create the I-SID.
- Associate the I-SID with either a platform or private VLAN; this association is not required on a DvR Leaf.

About This Task

You can create only one I-SID of each type.

The FA I-SID can be the same as the voice I-SID because they are used by different Auto-sense port states.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Configure Fabric Attach authentication:
 - a. Select **FaMsgAuthEnable**, to enable FA message authentication.
 - b. For **FaAuthenticationKey**, type the key for FA authentication for the Auto-sense ports.
 - c. (Optional) To force FA authentication for devices that attach to ports in the FA PROXY NOAUTH port state, select **FaProxyNoAuthForceAuth**.
5. Configure a specific I-SID to use instead of the onboarding I-SID:
 - a. For auto-sensed cameras, type the I-SID in **FaCamerasId**.
 - b. For auto-sensed FA client switches that do not use FA message authentication, like EXOS or Switch Engine, type the I-SID in **FaProxyNoAuthIsid**.
 - c. For auto-sensed virtual switches, type the I-SID in **FaVirtualSwitchIsid**.
 - d. For auto-sensed wireless access points (WAP), type the I-SID in **FaWapTypeIsid**.
6. (Optional) Configure a specific I-SID and customer VLAN ID to use as the management I-SID for proxy switches:
 - a. In **FaProxyMgmtIsid**, type the I-SID.
 - b. In **FaProxyMgmtCvid**, type the customer VLAN ID.
7. (Optional) Configure a specific I-SID and customer VLAN ID to use as the management I-SID for proxy ring switches:
 - a. In **FaProxyRingMgmtIsid**, type the I-SID.
 - b. In **FaProxyRingMgmtCvid**, type the customer VLAN ID.
8. Select **Apply**.

*Configure Maximum MAC Clients on Auto-sense Ports using EDM***About This Task**

Use this procedure to configure the maximum EAP and NEAP MAC clients supported on Auto-sense enabled ports.

**Note**

If you manually configure values on a specific port(s), then these values will take precedence over the Auto-sense global values.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **Eapol multihost mac-max** to configure the maximum EAP and NEAP MAC clients for Auto-sense enabled ports.
5. Select **Apply**.

*Configure Maximum EAP Clients on Auto-sense Ports using EDM***About This Task**

Use this procedure to configure the maximum EAP clients allowed on Auto-sense enabled ports.

**Note**

If you manually configure values on a specific port(s), then these values will take precedence over the Auto-sense global values.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **Eapol multihost eap-mac-max** to configure the maximum EAP clients for Auto-sense enabled ports.

**Note**

Eapol multihost eap-mac-max also provides the EAP and NEAP separation functionality. By default, the system enables EAP clients on each port with a maximum limit of 2 EAP clients. If you configure the maximum limit to 0 then the system disables the EAP client authentication.

5. Select **Apply**.

Configure Maximum NEAP Clients on Auto-sense Ports using EDM

About This Task

Use this procedure to configure the maximum NEAP clients allowed on Auto-sense enabled ports.



Note

If you manually configure values on a specific port(s), then these values will take precedence over the Auto-sense global values.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **Eapol multihost non-eap-mac-max** to configure the maximum EAP and NEAP MAC clients for Auto-sense enabled ports.



Note

Eapol multihost non-eap-mac-max also provides the EAP and NEAP separation functionality. By default, the system enables NEAP clients on each port with a maximum limit of 2 NEAP clients. If you configure the maximum limit to 0 then the system disables the NEAP client authentication.

5. Select **Apply**.

Configure Level 1 Metric for Auto-sense

About This Task

Configures the Level 1 metric for Auto-sense.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Configure your preferred method:
 - **IsisL1Metric**—Manually enter a value between 1-16777215.
 - **IsisL1MetricAuto**—Select **auto** and the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed.
5. Select **Apply**.

Enable Auto-MLT on NNI Ports

About This Task

Enable automatic link aggregation (Auto-MLT) to automatically aggregate NNI links on Auto-sense ports into an MLT to form IS-IS adjacencies. The default is enabled in Zero

Touch Deployment mode, when the switch has no configuration file. Otherwise, the default is disabled.

The switch aggregates the NNI links into an MLT if a minimum of two links with same speed are present between two devices. If only one link exists, the switch creates a single-port NNI on the Auto-sense ports.

If you enable this feature, the switch advertises this ability inside the Fabric Connect LLDP TLV using an `auto-lag-nni-capability` flag.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. Select **AutoMltEnable**.
5. Select **Apply**.

NEW! *Configure Link Debounce for Auto-sense Ports*



Note

This procedure does not apply to 7830 Series.

About This Task

Perform this procedure to configure Link Debounce on Auto-sense UNI ports. Link Debounce maintains a port link even if it is physically down. The port link remains up for the timer you configure. If the link is not restored within the timer period, the switch announces the port as down.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. In **LinkDebounceType**, select if the feature is enabled for all UNI ports, only for certain ports, or disabled for all UNI ports.
5. In **LinkDebounceTimeout**, configure the port timer.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AccessDiffservEnable	Enables or disables the differentiated service type as access for Auto-sense ports. The default is enabled.
Datalsid	Specifies the data I-SID used by the Auto-sense ports.
EapolVoiceLldpAuthEnable	Enables the EAPoL LLDP authentication for Auto-sense voice ports. The default is disabled.

Name	Description
FaMsgAuthEnable	Enables or disables the FA message authentication for Auto-sense ports. The default is enabled.
FaAuthenticationKey	Specifies the FA authentication key for Auto-sense ports.
IsisHelloAuthType	<p>Specifies the authentication type for IS-IS hello packets on Auto-sense ports:</p> <ul style="list-style-type: none"> • None • simple - simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5 - MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. • hmac-sha256 - with SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default authentication type is none.</p>
IsisHelloAuthKeyId	Specifies the IS-IS hello authentication number key id for the Auto-sense ports.
IsisHelloAuthKey	Specifies the IS-IS hello authentication number key for the Auto-sense ports. You must configure the IS-IS hello authentication key along with the IS-IS hello authentication type.
OnboardingIsid	Specifies the onboarding I-SID used by the Auto-sense ports.
Qos8021pOverrideEnable	Overrides the incoming 802.1p bits on ports that operate in Auto-sense mode. The default is enabled.
VoiceIsid	Specifies the voice I-SID used by Auto-sense ports.
VoiceCvid	Specifies the customer VLAN ID associated with the voice I-SID used by Auto-sense ports. Voice C-Vid is configured for tagged voice traffic only. You must configure the Auto-sense voice customer VLAN ID along with the Auto-sense voice I-SID.
DhcpDetection	Enables or disables the DHCP detection in Auto-sense mode. The default is enabled.

Name	Description
FaCameraIsid	Specifies the FA camera I-SID used by Auto-sense ports.
FaProxyMgmtIsid	Specifies the FA proxy management I-SID used by Auto-sense ports.
FaProxyMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
FaProxyRingMgmtIsid	Specifies the FA proxy ring management I-SID used by Auto-sense ports.
FaProxyRingMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
FaProxyNoAuthIsid	Specifies the FA proxy no-auth I-SID used by Auto-sense ports.
FaProxyNoAuthForceAuth	Configures Auto-sense to force FA authentication for devices that attach to ports in the FA PROXY NOAUTH port state. The default is disabled.
FaVirtualSwitchIsid	Specifies the FA virtual-switch I-SID used by Auto-sense ports.
FaWapType1Isid	Specifies the FA WAP type-1 I-SID used by Auto-sense ports.
FaCameraEapolStatus	Specifies the FA EAPoL status for Camera I-SID used by Auto-sense ports.
FaEapolOVSStatus	Specifies the FA EAPoL status for OVS (Open-Virtual-Switch) I-SID used by Auto-sense ports.
FaEapolWap1Status	Specifies the FA EAPoL status for Wap-type-1 I-SID used by Auto-sense ports.
WaitInterval	Specifies the wait interval, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state. This configuration is a global configuration that applies to all Auto-sense ports. The default value is 35.
MultihostMacMax	Specifies the maximum number of EAPoL and non-EAPoL authentication MAC addresses allowed on this port. The default value is 4.
MultihostEapMacMax	Specifies the maximum number of EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port. The default value is 4.
MultihostNonEapMacMax	Specifies the maximum number of non-EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port. The default value is 4.
IsisL1Metric	Manually configure a value for the Level 1 metric. A higher number represents a higher cost and the least preferred route. The default value for L1 metric is 10 for any link, despite the port speed.

Name	Description
IsisLlMetricAuto	Enable the Level 1 metric as automatic. By enabling Level 1 metric as auto, the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed. This ensures that the fastest port speed determines the network route. The default is disabled.
IsisStpMultiHomingOverwrite	Displays if Auto-sense overwrites MSTP-Fabric Connect Multi Homing.
AutoMltEnable	Enables automatic link aggregation (Auto-MLT) to automatically aggregate NNI links on Auto-sense ports into an MLT to form IS-IS adjacencies. The default is enabled in Zero Touch Deployment mode, when the switch has no configuration file. Otherwise, the default is disabled.
GuestIsidEnable	Enables or disables Auto-sense guest I-SID in the UNI, FA, and VOICE states. If the switch boots in Zero Touch Deployment mode, the Auto-sense guest I-SID is disabled. If the switch boots with the factorydefaults boot flag enabled, the Auto-sense guest I-SID is enabled.
LinkDebounceType	Configures Link Debounce for Auto-sense ports: <ul style="list-style-type: none"> off—Disables Link Debounce on all Auto-sense UNI ports on—Enables Link Debounce on all Auto-sense UNI ports. auto—Enables Link Debounce only on Auto-sense UNI ports that connect to Preboot Execution Environment (PXE) devices. The auto parameter enables Link Debounce if the port receives a DHCP DISCOVERY message with an Option 60 vendor class ID string of PXEClient:Arch:xxxxx:UNDI:yyyzzz. The default is auto.
LinkDebounceTimeout	Specifies the port timer in milliseconds. The default timer is 3,000 milliseconds.

Auto-sense Port Configuration using EDM

Perform the procedures in this section to configure Auto-sense on specific ports using EDM.

Port Configuration Using the Fabric Menu

The topics in this section include instructions to configure an Auto-sense port using the **Fabric > AutoSense** menu in EDM. You can perform many tasks on the same EDM tab. For a quick look at field descriptions for that tab, see [Port AutoSense Field Descriptions](#) on page 70.

View or Modify the Auto-sense Port Configuration

Perform this procedure to view or modify the Auto-sense configuration on one or more ports. You also have the option to disable Auto-sense on the port but retain the configuration that the system applied dynamically. The dynamic configuration becomes a manual configuration and is visible in the **show running-config** output.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Port AutoSense** tab.
4. To change the Auto-sense state on a port, select a new value for **AutoSense** in the appropriate row for the port.
5. (Optional) If you disable Auto-sense on a port, select **true** for **AutoSenseKeepAutoConfig** to retain the configuration that the system applies dynamically.
6. (Optional) Type a value in **AutoSenseDatalsid** to configure a specific I-SID on the port.
7. (Optional) Type a value in **AutoSenseWaitInterval** to configure a specific interval on the port.
8. Select **Apply**.

Prevent a Port From Entering SD-WAN and NNI Port States

About This Task

To increase security when these Auto-sense port states are not required, you can block the port from sending Fabric Connect LLDP TLVs and prevent the port from transitioning to the SD-WAN and NNI states.

By default, this option is disabled.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Port AutoSense** tab.
4. In the port row, for **AutoSenseNoNni**, select **enable**.
5. Select **Apply**.

Port AutoSense Field Descriptions

Use the data in the following table to use the **Port AutoSense** tab.

Name	Description
Index	Displays the index of the port, written in the slot/port[/sub-port] format.
AutoSense	Enables or disables Auto-sense on the specific port. The default is disabled for existing configurations but enabled for new Zero Touch Fabric Configuration deployments.
AutoSenseKeepAutoConfig	Retains the Auto-sense configuration if you disable Auto-sense on the port. The dynamic configuration becomes a manual configuration and is visible in the show running-config output.
AutoSenseState	Displays the Auto-sense port state.
AutoSenseDataIsid	Specifies the Auto-sense data I-SID per port. The port configuration takes precedence over the global data I-SID configuration.
AutoSenseWaitInterval	Specifies the Auto-sense wait time on a port. The port configuration takes precedence over the global wait-interval configuration. The default value is 35.
AutoSenseNoNni	Blocks the port from sending Fabric Connect LLDP TLVs and prevents the port from transitioning to the SD-WAN and NNI states. The default is disabled.

Port Configuration Using the Port Menu

The topics in this section include instructions to configure an Auto-sense port using the **Edit > Port** menu in EDM.

Enable Auto-sense on Port(s)

Perform this procedure to enable Auto-sense on one or more ports.

About This Task

If you select more than one port, the format of the tab changes to a table-based tab.



Note

- After a switch boots without a configuration file, Auto-sense is enabled on all ports, by default.
- Auto-sense is disabled by default for existing configurations but enabled for new Zero Touch Fabric Configuration deployments.
- Auto-Sense cannot be enabled when the spbm-config-mode flag is false.

You can also complete this task on the **Configuration > Fabric > AutoSense > Port AutoSense** tab. For more information, see [View or Modify the Auto-sense Port Configuration](#) on page 69.

Procedure

1. In the **Device Physical View** tab, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. For **AutoSense**, select enable.
6. Select **Apply**.

Disable Auto-sense on Port(s)

Perform this procedure to disable Auto-sense on one or more ports. You also have the option to disable Auto-sense on the port but retain the configuration that the system applied dynamically. The dynamic configuration becomes a manual configuration and is visible in the **show running-config** output.

About This Task

If you select more than one port, the format of the tab changes to a table-based tab.

You can also complete this task on the **Configuration > Fabric > AutoSense > Port AutoSense** tab. For more information, see [View or Modify the Auto-sense Port Configuration](#) on page 69.

Procedure

1. In the **Device Physical View** tab, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. For **AutoSense**, select disable.
6. (Optional) Select **AutoSenseKeepAutoConfig** to retain the configuration that the system applies dynamically.
7. Select **Apply**.

Configure an Auto-sense Data I-SID on a Port

Perform this procedure to configure an Auto-sense Data I-SID on a port.

About This Task

You can also complete this task on the **Configuration > Fabric > AutoSense > Port AutoSense** tab. For more information, see [View or Modify the Auto-sense Port Configuration](#) on page 69.

Procedure

1. In the **Device Physical View** tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.

5. For **AutoSenseDataIsid**, type the data I-SID value.
6. Select **Apply**.

Configure the Auto-sense Wait Interval for a Port

Perform this task to configure the time, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state.

About This Task

You can also complete this task on the **Configuration > Fabric > AutoSense > Port AutoSense** tab. For more information, see [View or Modify the Auto-sense Port Configuration](#) on page 69.

Procedure

1. In the **Device Physical View** tab, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. Type a value for **AutoSenseWaitInterval**.
6. Select **Apply**.

Auto-sense Logical Flowcharts

The system uses a per-interface state to adapt to all Auto-sense events. Each state transition determines background configuration on the port. The system does not display Auto-sense port configurations in the **show running-config** command or in the saved configuration file; instead use the **show auto-sense** commands to show global and port specific Auto-sense information..

The following flowcharts describe the system logic for Auto-sense port state detection, how the system configurations change the logic path, and the Auto-sense configuration results.



Note

The **vlan create** CLI command examples do not apply to DvR leaf switch configurations. DvR leaf switches create VLANs automatically.

Auto-sense Fabric NNI

Detection of a Fabric network-to-network interface (NNI) results in tagged Backbone VLAN IDs (B-VID) with IS-IS NNI enabled. In the following image, NNI-MLT refers to the port state [NNI-MLT](#) and [NNI-MLT-ISIS-UP](#) on page 30.



Note

If you enable **auto-sense no-nni** on the port, all NNI port states transition to the Wait state immediately, and after the timer expires, to the UNI port state. For more information, see [Wait State](#) on page 28 and [UNI State](#) on page 28.

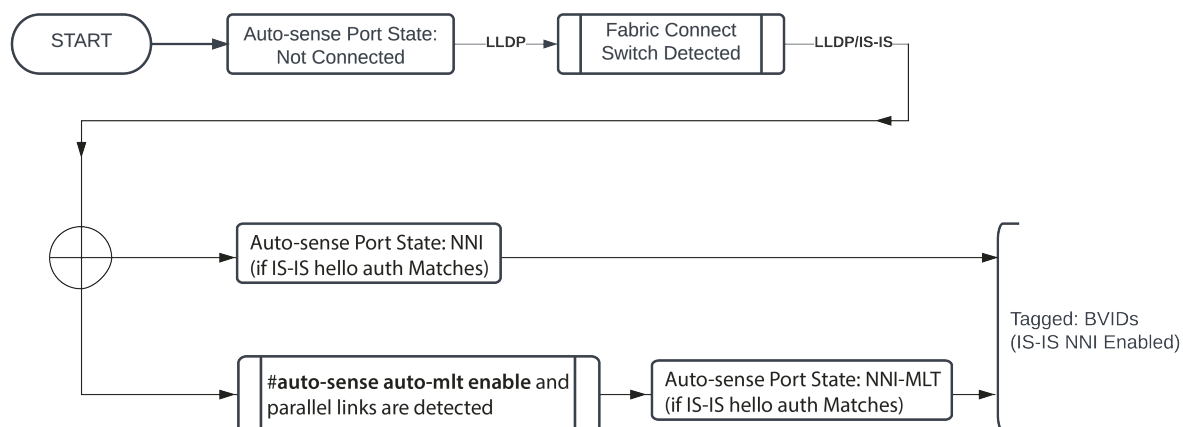


Figure 1: Auto-sense Fabric NNI

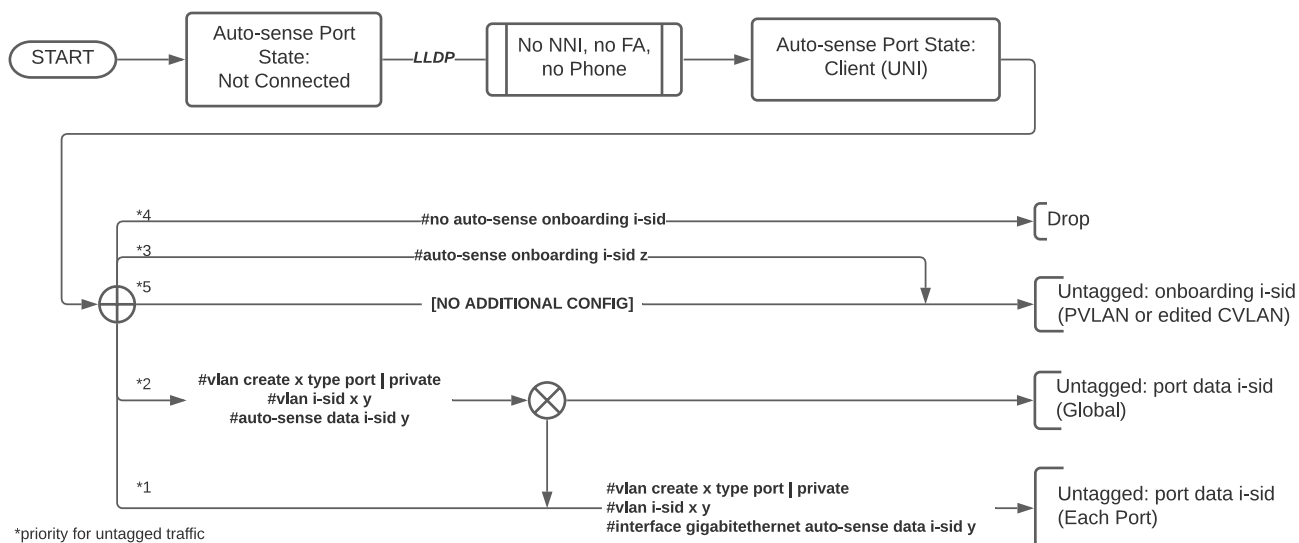


Note

Specific conditions must be met to form a dynamic MLT. For more information, see [Automatic Link Aggregation on NNI Ports](#) on page 26.

Auto-sense UNI Client without NAC

Detection of a user-to-network interface (UNI) client without network access control (NAC) results in untagged data configuration based on system configuration.



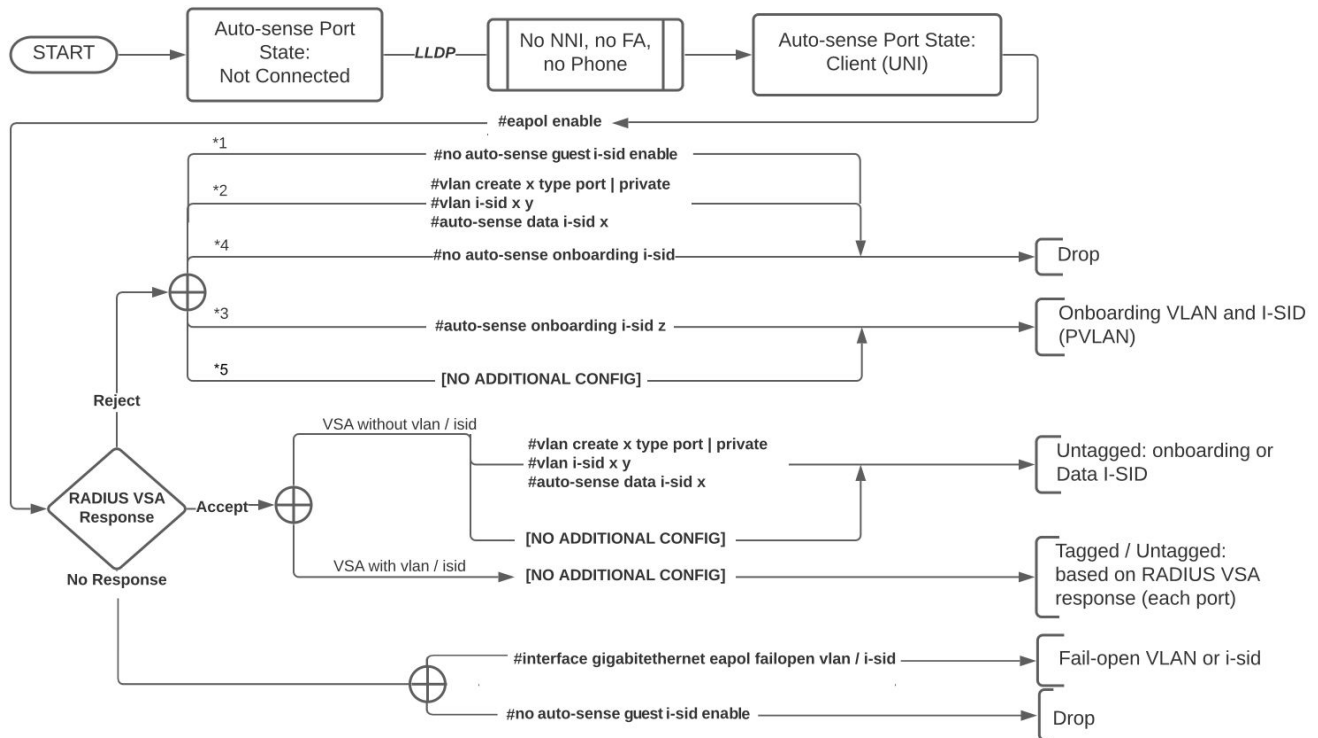
*priority for untagged traffic

Figure 2: Auto-sense UNI client without NAC

Auto-sense UNI Client with NAC

Detection of a user-to-network interface (UNI) client with network access control (NAC) results in untagged, tagged, or dropped data configuration based on system configuration.

Auto-sense UNI client with NAC



Auto-sense Voice without NAC

Detection of voice without network access control (NAC) results in untagged, tagged, or dropped data based on system configuration.

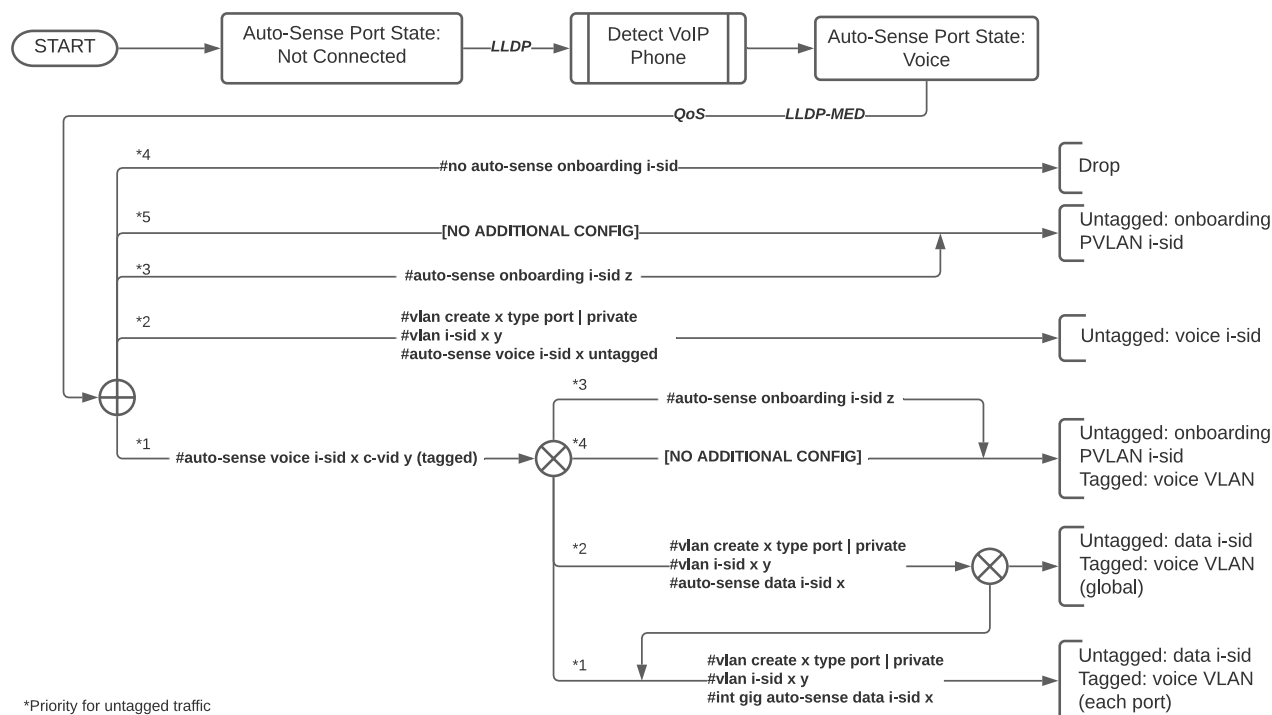
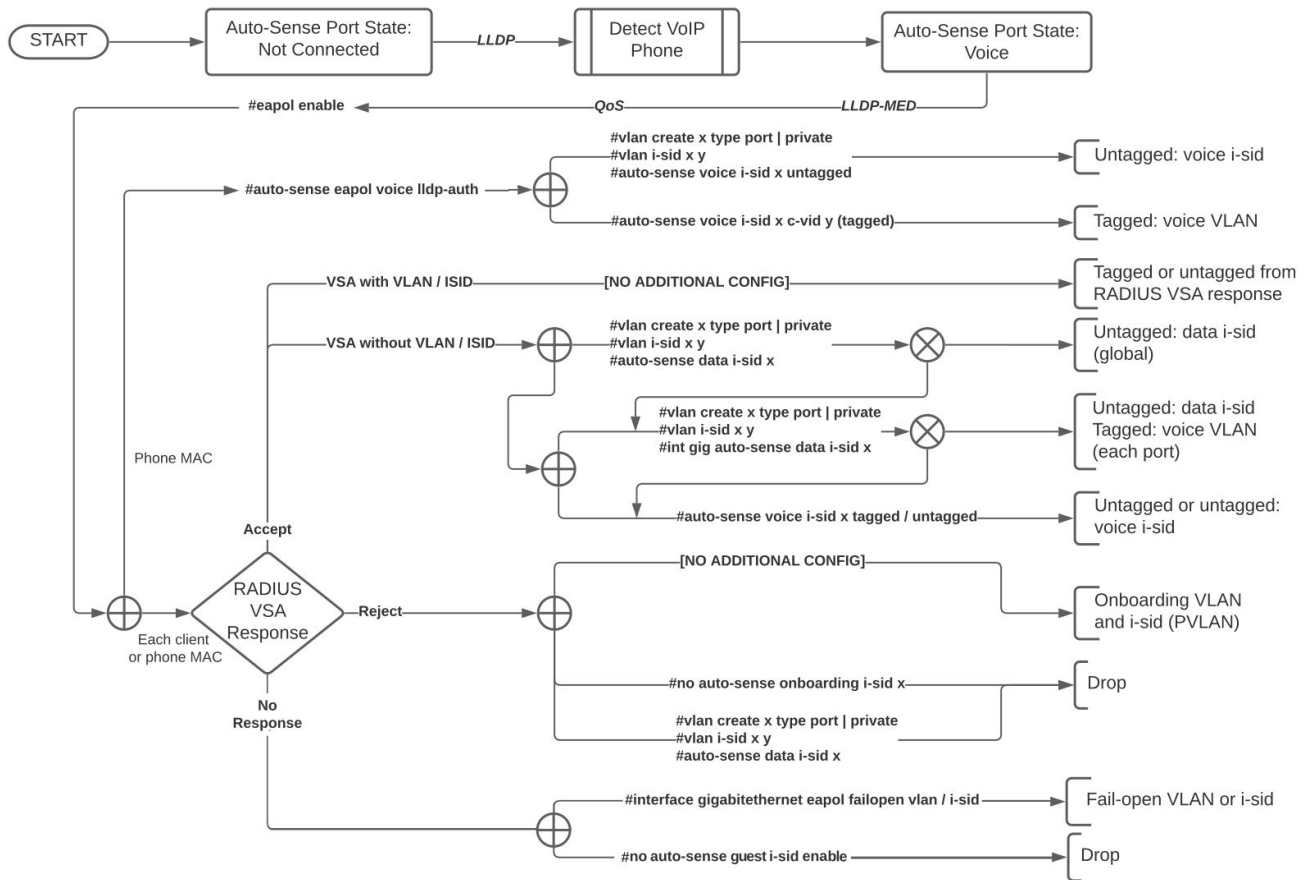


Figure 3: Auto-sense voice without NAC

Auto-sense Voice with NAC

Detection of voice with network access control (NAC) results in untagged, tagged, or dropped data based on system configuration.

Auto-sense voice with NAC**Auto-sense FA Proxy Switch**

Detection of a Fabric Attach (FA) proxy switch results in untagged or tagged data based on system configuration.

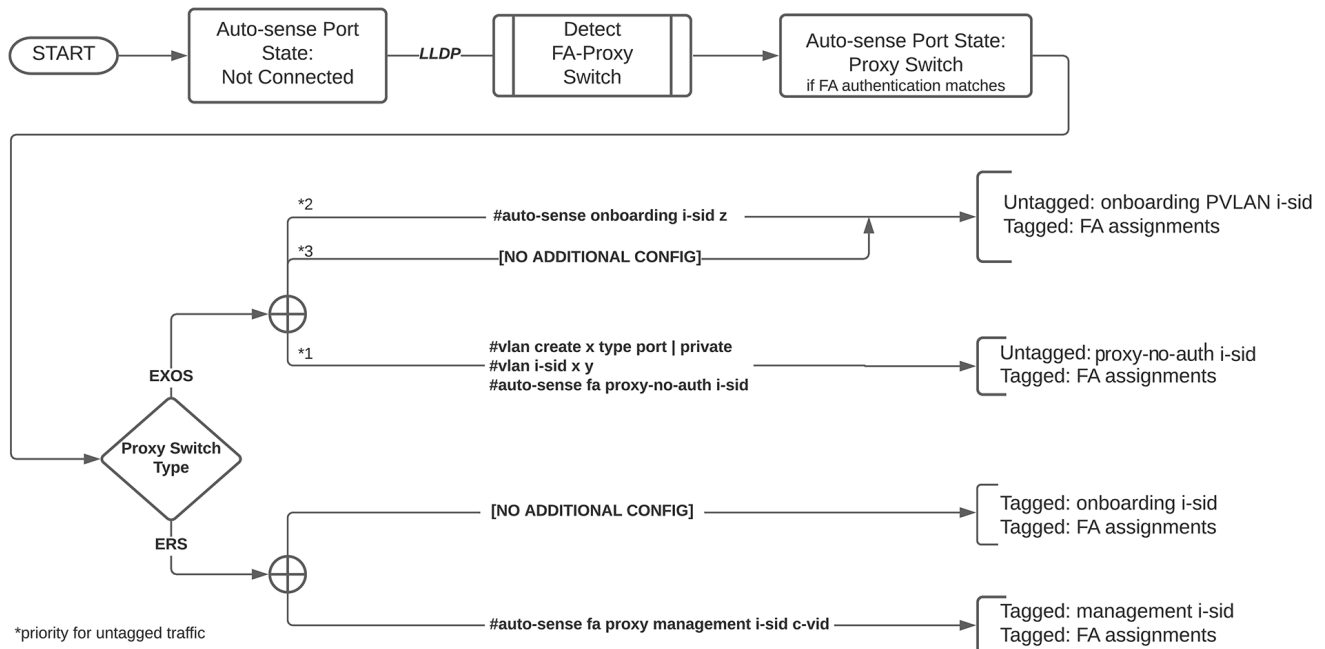


Figure 4: Auto-sense FA proxy switch

Auto-sense FA Ring Proxy Switch

Detection of Fabric Attach (FA) ring proxy switch results in untagged or tagged data based on system configuration.

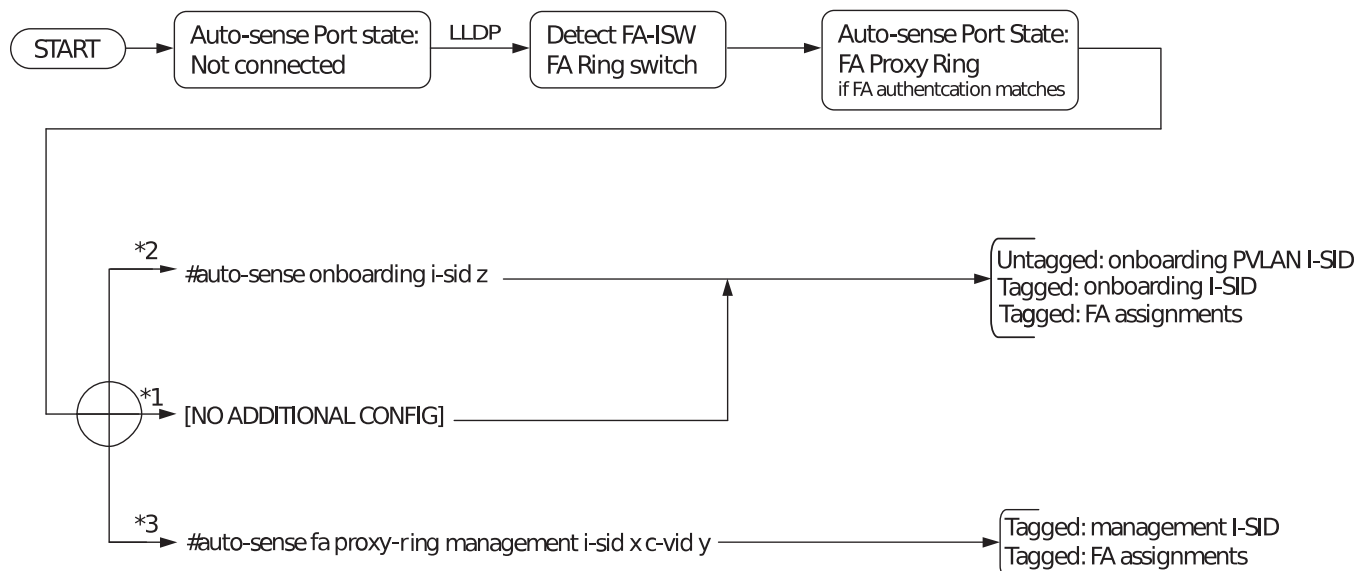


Figure 5: Auto-sense FA Proxy Ring Switch

Auto-sense FA WAP, Camera, or OVS without NAC

Detection of a Fabric Attach (FA) wireless access point (WAP), camera, or open virtual switch (OVS) without network access control (NAC) results in untagged or tagged data based on system configuration.




Figure 6: Auto-sense FA WAP / camera / OVS without NAC

Auto-sense FA WAP, Camera, or OVS with NAC

Detection of a Fabric Attach (FA) wireless access point (WAP), camera, or open virtual switch (OVS) with network access control (NAC) results in untagged or tagged data based on system configuration.

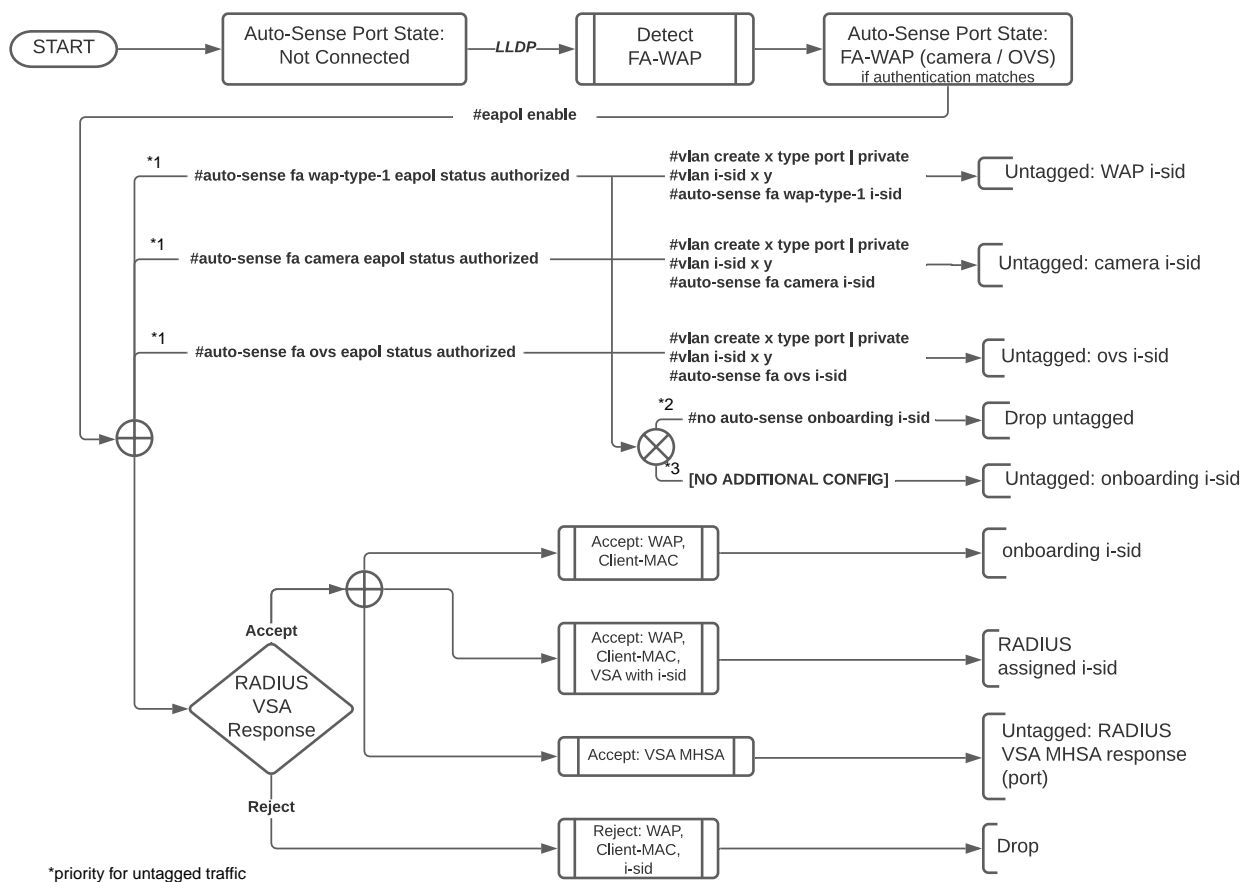


Figure 7: Auto-sense FA WAP / camera / OVS with NAC

Automatic Channelization with Auto-sense

The switch automatically channelizes a port if the port operates in Auto-sense mode and you insert a supported part. The following images show the impact of Auto-sense on a channelization-capable port.

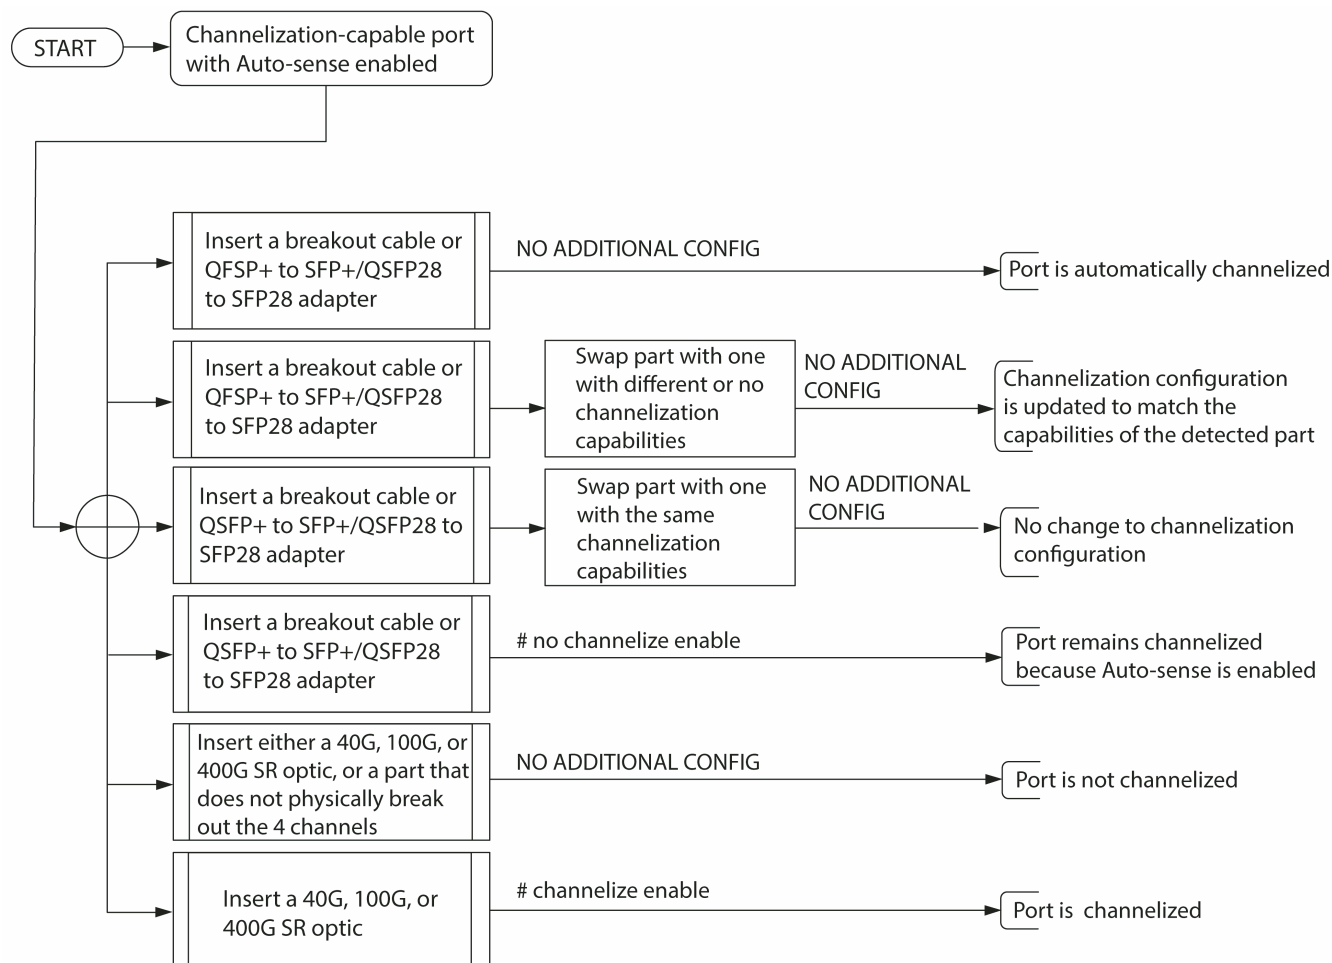


Figure 8: Automatic Channelization with Auto-sense enabled

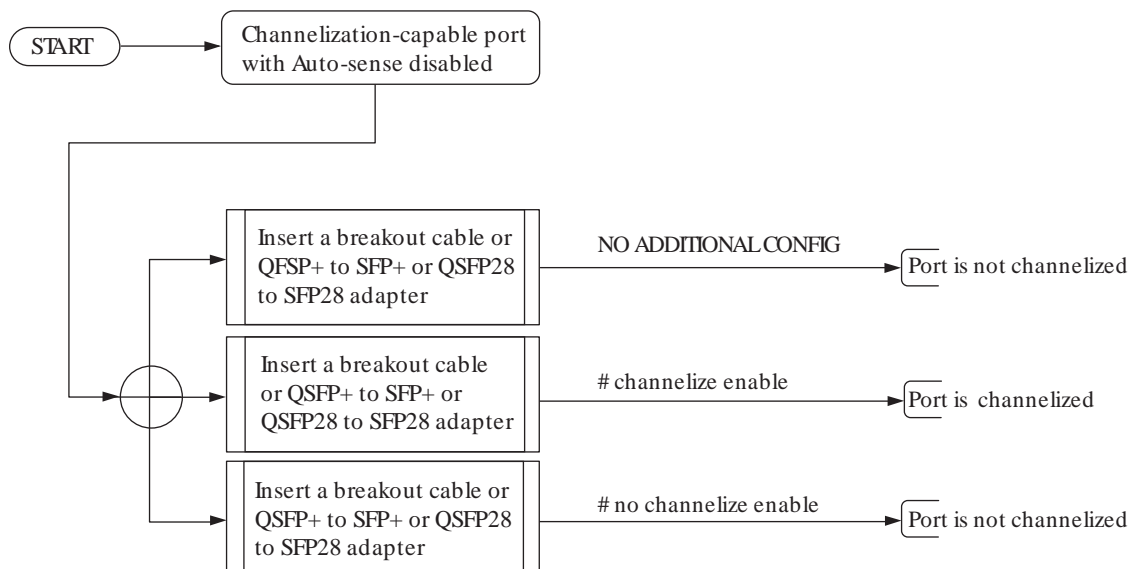


Figure 9: Automatic Channelization with Auto-sense disabled

For more information, see [Channelization](#) on page 704.

Auto-sense with ExtremeCloud SD-WAN

Detection of an SD-WAN Appliance results in untagged VLAN 4047 and the application of Fabric Extend tunnel assignments. If you use SD-WAN Orchestrator to enable a Local Breakout (LBO) branch, Auto-sense can also apply the LBO configuration to select VRFs.

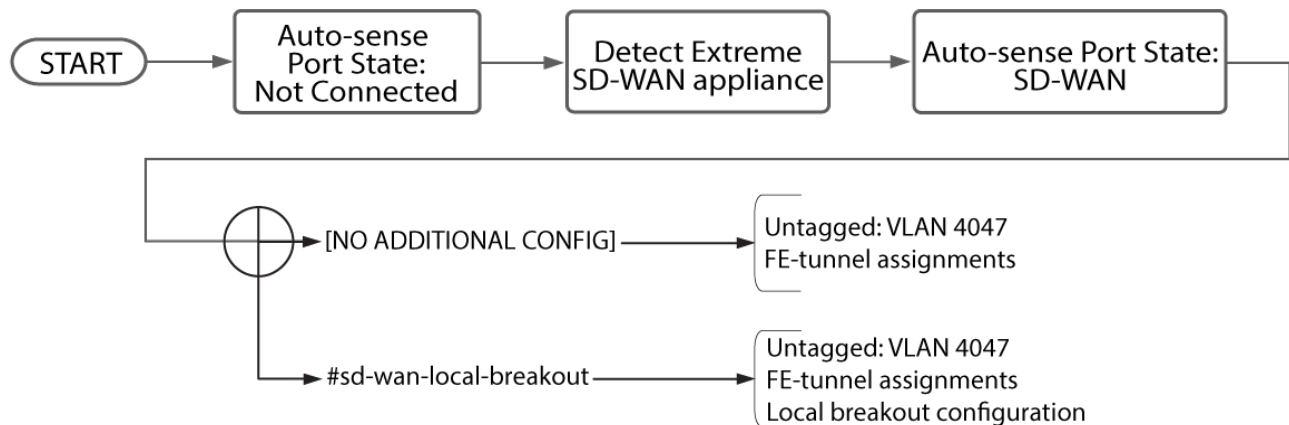


Figure 10: Auto-sense SD-WAN

IP Phone Support

Table 7: IP Phone product support

Feature	Product	Release introduced
IP Phone Support	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

The IP phone support feature focusses on the following key points:

- Works only on the Flex UNI-enabled and Auto-sense ports.
- For Avaya phones, you can choose to configure and send call server and file server Link Layer Discovery Protocol (LLDP) Type-Length-Value (TLV) options through the **lldp vendor-specific** CLI command.

- To reduce configuration overhead, this feature includes the Auto-sense voice mechanism to detect IP phones from LLDP signalling. After the switch detects the phone, this mechanism manages the following tasks:
 - Provides the voice VLAN to the phone, tagging, Differentiated Services Code Point (DSCP), and priority parameters through the LLDP Media Endpoint Discovery (MED) signaling options.
 - Configures a switched UNI for phone traffic and sends it to the Service Instance Identifier (I-SID) that is associated with the voice VLAN.
 - Handles the configuration, whether “trusted” or “untrusted” on the port and priority re-markings.
 - Integrates with the Auto-sense functionality. For more information on Auto-sense, see [Auto-sense](#) on page 21.

**Note**

This feature does not support auto-creation of voice VLAN and MultiLink Trunking (MLT) or Split Multi-Link Trunking (SMLT).

This feature has the following connectivity models:

- Standalone IP phone, which connects to a switch port.
- IP Phone with PC behind it, where the IP Phone has a small inbuilt bridge, and a PC connects to that bridge port.

**Note**

Phone traffic is tagged with the voice VLAN whereas the PC traffic is untagged. However, you can configure the phone to send the traffic as untagged.

The IP phone connectivity supports the following scenarios:

- Call server and file Server LLDP TLV options—These TLVs are Avaya proprietary. Use them only with the Avaya IP phones to detect the IP addresses of a Call server and File Server.
- Phone detection through LLDP messaging—Use the Capabilities and Enabled Capabilities field in the LLDP packet to detect a phone. A “T” capability identifies a phone.
- Auto-sense voice option, without Network Access Control (NAC)—Use this functionality to specify the voice VLAN and voice I-SID in a single CLI command.
- Auto-ISID-Offset—Use this functionality if the voice VLAN is received without an I-SID from a Radius response. The Auto-ISID-Offset functionality determines an I-SID automatically to send the data traffic.
- Auto-sense voice, LLDP authentication, and Non- EAP (NEAP) (MAC authentication) connectivity—If you have enabled NEAP, it authorizes all the MAC addresses received on the port and IP phone. With the LLDP authentication option, a device, such as phone, is trusted and does not require a Remote Authentication Dial-in User (RADIUS) authentication. For this authentication, the Extensible Authentication Protocol (EAP) is notified after a phone is detected and the port is in the Auto-sense voice state. Then, the MAC address of the phone is added to EAP or NEAP host table.

Auto-sense Voice

The Auto-sense voice feature is an addition to the Auto-sense module. Based on the events of the phone discovery in the network, you can use this feature to configure phone devices without manual intervention.

With the Auto-sense voice feature, you can configure the voice I-SID and the voice VLAN. If you configure the I-SID as untagged, the phone receives VLAN as zero. When you configure Auto-sense voice, switched UNI is configured in VOICE I-SID for each port that is in "voice" state. The switch adds the Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) policies for voice and voice-signaling Type-Length-Value (TLVs) to the LLDP packet and sends LLDP packets to the phone. It uses the configured voice VLAN and default values for Differentiated Services Code Point (DSCP) (46) and priority (6). After you run the **auto-sense voice** command, a filter is installed to prioritize the traffic that passes through the configured I-SID. The filter applies to traffic that reaches the VOICE I-SID, which implies the voice traffic.

After the switch discovers a Link Layer Discovery Protocol (LLDP) packet with phone capabilities, the port transitions to the "voice" state. The port receives a message on the voice event details.

After you configure auto-sense voice, following tasks take place:

- The I-SID <I-SID value> and C-VID <C-VID value> are saved in the control plane.
- The voice I-SID is created.
- The ports that are in the "voice" state process the voice configuration message and begin the dynamic configuration. This configuration includes the following tasks:
 - Creation of voice Switched UNI (S-UNI).
 - Deletion of onboarding and data I-SID, or S-UNI, if you configured auto-sense as "untagged".
 - Addition of LLDP-MED voice and voice-signaling TLVs to the LLDP packet and sending of LLDP packets to the phone.
- The voice filter is updated. After you run the **auto-sense voice** command, a filter is installed to prioritize the traffic that passes through the configured I-SID. The filter applies to the traffic in Voice I-SID. The traffic that passes through this I-SID is

internally prioritized with level 6 and forwarded with a dot1p value of 6, for tagged packets. For the IP packets, the DSCP value of 46 is forwarded.



Note

- To disable Auto-sense on the port but keep the dynamic configurations made by Auto-sense, use the command **no auto-sense enable convert-to-config**. The voice S-UNI loses its Auto-sense origin and has a config origin instead. The LLDP-MED policies installed by Auto-sense are preserved.
- If you use the **no auto-sense voice** command, the system removes the voice S-UNI and the LLDP-MED policies. The voice I-SID is removed if it was installed by using the **auto-sense voice** command. If the I-SID existed before you used the **auto-sense voice** command, the system does not remove the I-SID but the I-SID does lose its Auto-sense origin.
- A port exits the voice state in one of the following scenarios:
 - If the port is down
 - If the LLDP session fails between the switch and the phone
 - If Auto-sense is disabled on the port that connects to the IP phone

After a port exits the voice state, the Switched UNI (S-UNI), LLDP voice and voice-signaling are deleted.

IP Phone Configuration using CLI

This section provides procedures to configure IP phones using the command line interface (CLI).

Configure Auto-sense Voice Information for IP Phones

The switch applies the voice configuration on Auto-sense-enabled ports, after it discovers IP phones on the port through Link Layer Discovery Protocol (LLDP) packets.

Before You Begin

If you boot the switch with a configuration file, and not through Zero Touch Fabric Configuration, you must manually enable Auto-sense on specific ports.

About This Task

Perform this procedure to configure Auto-sense voice information for IP phones.

A global Auto-sense voice configuration does not require Auto-sense LLDP authentication based on the following cases.

- In a non NAC, a phone is classified based on the phones LLDP signaling.
- In a NAC, a phone is authenticated based on EAP/NEAP radius authenticated, or if configured, it is LLDP authenticated

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Choose one of the following options to configure Auto-sense voice information for IP phones:
 - To configure voice VLAN traffic for an IP phone with a PC behind it: `auto-sense voice i-sid <1-15999999> c-vid <1-4094>`
 - To configure voice VLAN traffic for a standalone IP phone:
 - a. Delete all Auto-sense voice configurations.
 - b. Configure a customer voice VLAN: `vlan create <2-4059> name WORD<0-64> type port-mstprstp <0-63>`
 - c. Associate a CVLAN I-SID to the customer voice VLAN: `vlan i-sid <1-4059> <1-16777215> [force]`
3. Configure the traffic as untagged:

```
auto-sense voice i-sid <1-15999999> untagged
```



Note

The phone receives VLAN ID as 0 and the tagging is configured as "untagged".

Example

Configure VLAN tagging as untagged:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense voice i-sid 1234 untagged
```

Variable Definitions

The following table defines parameters for the **auto-sense voice** command.

Variable	Value
<code>i-sid<1-15999999></code>	Specifies the service instance identifier (I-SID).
<code>c-vid<1-4094></code>	Specifies the customer VLAN ID.
<code>untagged</code>	Specifies the VLAN tagging type as untagged. Note: The phone receives VLAN ID as 0 and the tagging is configured as "untagged".

Enable Auto-sense LLDP Authentication of IP Phones

Before You Begin

- You must enable Extensible Authentication Protocol over LAN (EAPoL) globally.

About This Task

Perform this procedure to enable Link Layer Discovery Protocol (LLDP) authentication of IP phones. The switch authenticates the phone after it receives LLDP packets from the phone if EAP/NEAP is enabled.

Auto-sense LLDP authentication applies to Auto-sense ports in the VOICE state. Auto-sense LLDP authentication does not require a global Auto-sense voice configuration.

The **no auto-sense eapol voice lldp-auth** command removes all Auto-sense LLDP sessions and removes the Auto-sense LLDP authentication configuration.

The system removes the LLDP session for the following reasons:

- You disable EAPoL globally.
- You disable Auto-sense on the port.
- The LLDP neighbor is removed.

If the LLDP authentication configuration exists and one of the following situations occur, the LLDP session is recreated:

- You reenable EAPoL globally.
- You reenable Auto-sense on the port.
- The LLDP neighbor is recreated.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable LLDP authentication:

```
auto-sense eapol voice lldp-auth
```

Example

Enabling LLDP authentication on the switch:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#auto-sense eapol voice lldp-auth
```

Variable Definitions

The following table defines parameters for the **auto-sense eapol voice** command.

Variable	Value
<i>lldp-auth</i>	Enables Link Layer Discovery Protocol (LLDP) authentication of IP phones. By default, LLDP authentication of IP phones is disabled on the switch.

Configure LLDP Vendor Specific Information

About This Task

Use this procedure to configure the Link Layer Discovery Protocol (LLDP) vendor-specific information on a call server or a file server.



Note

After you configure LLDP vendor specific call server information, the SIP Proxy of the phone is configured as transport type Transport Layer Security (TLS) port 5061. This option is available depending on the operating system of the call server.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. To configure LLDP vendor-specific information for a call server server, enter:

```
lldp vendor-specific call-server <1-8> <A.B.C.D>
```
3. To configure LLDP vendor-specific information for a file server, enter

```
lldp vendor-specific file-server <1-4> <A.B.C.D>
```

Example

Configure the LLDP vendor-specific information on a call server:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#lldp vendor-specific call-server 1 192.0.2.0
```

Variable Definitions

The following table defines parameters for the **lldp vendor-specific** command.

Variable	Value
<i>call-server</i> <1-8> <A.B.C.D>	Specifies the Link Layer Discovery Protocol (LLDP) vendor specific information on the call server number and the IP address.
<i>file-server</i> <1-4> <A.B.C.D>	Specifies an LLDP vendor specific information on the file server number and the IP address.

View LLDP Vendor Specific Information

About This Task

Use this procedure to view the Link Layer Discovery Protocol (LLDP) vendor-specific information on a call server or a file server.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. To view LLDP vendor-specific information for a call server, enter:
`show lldp vendor-specific call-server`
3. To configure LLDP vendor-specific information for a file server, enter
`show lldp vendor-specific file-server`

Example

Display the LLDP vendor-specific information on a call server:

```
Switch:1>enable
Switch:1#show lldp vendor-specific call-server
=====
                        LLDP Call-Server
=====
NUM          IP
-----
1            192.0.2.0
2            198.51.100.0
-----
All 2 out of 2 Total Num of call-server entries displayed
```

View LLDP Neighbor Vendor Specific Information

About This Task

Use this procedure to view the remote Link Layer Discovery Protocol (LLDP) vendor-specific information on a call server or a file server.

Procedure

1. Enter Privileged EXEC mode:
`enable`

2. To view remote LLDP vendor-specific information for a call server, enter:
`show lldp neighbor vendor-specific call-server`
3. To view remote LLDP vendor-specific information for a file server, enter:
`show lldp neighbor vendor-specific file-server`

Example

Display remote LLDP vendor-specific information on a file server:

```
Switch:1>enable
Switch:1#show lldp neighbor vendor-specific file-server
=====
Remote LLDP File-Server IP Addresses
=====
PORT      IP
-----
203       192.0.2.0, 198.51.100.0, 203.0.113.0
-----
All 3 out of 3 Total Num of remote file-server entries displayed
```

Enable LLDP Voice Authentication on a Specific Port

About This Task

Perform this procedure to enable Link Layer Discovery Protocol (LLDP) voice authentication of IP phones on a port.

You cannot manually enable LLDP voice authentication on an Auto-sense-enabled port. If the system detects a phone on an Auto-sense port, then “eapol voice lldp-auth” configuration is automatically applied on the port that connects to the phone. This procedure applies to ports with Auto-sense disabled.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
`enable`
`configure terminal`
`interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]`
`[,...][slot/all][all]}`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable LLDP voice authentication on a specific port:
`eapol voice lldp-auth`

Example

Enabling LLDP voice authentication on a specific port:

```
Switch:1>enable
Switch:1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#eapol voice lldp-auth
```

Variable Definitions

The following table defines parameters for the **eapol voice** command.

Variable	Value
<i>lldp-auth</i>	Enables Link Layer Discovery Protocol (LLDP) voice authentication of IP phones on the selected port. By default, LLDP authentication of IP phones is disabled on the switch.

IP Phone Configuration using EDM

This section provides procedures to view IP phone configuration using EDM.

View Vendor Specific Call Server Information

About This Task

Perform this procedure to view inventory attributes for vendor-specific call server information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Vendor Specific**.
3. Select the **Call Server** tab.

Vendor Specific Call Server Field Descriptions

Use the data in the following table to use the **Call Server** tab.

Name	Description
CallServerNum	Specifies the call server ID.
CallServerAddressType	Specifies the IP address type of the call server.
CallServerAddress	Specifies the IP address of the call server.

View Vendor Specific File Server Information

About This Task

Perform this procedure to view inventory attributes for vendor-specific file server information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Vendor Specific**.
3. Select the **File Server** tab.

Vendor Specific File Server Field Descriptions

Use the data in the following table to use the **File Server** tab.

Name	Description
FileServerNum	Specifies the file server ID.
FileServerAddressType	Specifies the IP address type of the file server.
FileServerAddress	Specifies the IP address of the file server.

Zero Touch Deployment

Table 8: Zero Touch Deployment product support

Feature	Product	Release introduced
Zero Touch Deployment	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2 Not supported on VIM5-2Q
	VSP 7400 Series	VOSS 8.2

For the most current information on switches supported by ExtremeCloud™ IQ, see [ExtremeCloud IQ Learning What's New](#).

Zero Touch Deployment enables a switch to be deployed automatically with ExtremeCloud IQ but you still need to onboard the switch on the ExtremeCloud IQ side. When the switch powers on, the Dynamic Host Configuration Protocol (DHCP) Client obtains the IP address and gateway from a DHCP Server, and discovers the Domain Name Server, connecting the switch automatically to ExtremeCloud IQ Site Engine or to ExtremeCloud IQ cloud management application.

The switch integrates with ExtremeCloud IQ using IQAgent.

Zero Touch Provisioning Plus (ZTP+) provides ExtremeCloud IQ Site Engine connectivity to the switch.

For more information about ExtremeCloud IQ Agent, see [ExtremeCloud IQ Support](#) on page 1084. For more information about ZTP+, see [Zero Touch Provisioning Plus](#) on page 96.

To use zero touch functionality, your switch must be in a Zero Touch Deployment-ready configuration mode, which means the switch cannot have existing primary or secondary configuration files loaded. Factory shipped switches are Zero Touch Deployment ready because they deploy without configuration files. However, existing switches require manual preparation before Zero Touch Deployment can function.

To prepare an existing switch for Zero Touch Deployment, the switch must boot without a configuration file. Perform one of the following actions:

- Rename existing primary and secondary configuration files. Use the **mv** command to rename the existing configuration files. For example, **mv config.cfg config.cfg.backup**.

This is the preferred option as it ensures that the primary and secondary files are removed while making a backup of them at the same time. This option also ensures that the switch uses the default config.cfg file for the final configuration after it has successfully onboarded.

- Boot from non-existent configuration files. Use the **boot config choice** command to configure the primary and backup configuration files to reference files that do not exist on the switch:

```
boot config choice primary config-file nonexistent1.cfg
```

```
boot config choice primary backup-config-file nonexistent2.cfg
```

This option also works, however, after the switch has successfully onboarded, it does not use the default config.cfg file but uses the alternative configuration file name provided instead, which might not be desired.

- Delete the existing primary and secondary configuration files. Create a backup of these files before you delete them.

Configuration Considerations

The switch configuration depends on whether you use Zero Touch Deployment or factory default mode.

Zero Touch Deployment Configuration

With Zero Touch Deployment, the switch configuration consists of the following:

- The ssh and sshd boot configuration flags are enabled by default.
- The Auto-sense guest I-SID is disabled by default.
- All ports are Private VLAN isolated ports.

- VLAN 4048 is created as an *onboarding-vlan* for host-only connectivity for In Band management. All front panel ports are members of VLAN 4048.
- In Band management is enabled.
- Dynamic Host Configuration Protocol (DHCP) client requests are cycled between In Band and Out of Band ports, except on the 5320 Series. With the exception of 5320-24T-4X-XT and 5320-24T-24S-4XE-XT, 5320 Series supports In Band management only.
- If the switch resets after the IP address is obtained from the DHCP Server, the entire DHCP process does not need to be repeated. Instead, the switch can directly send the DHCP Request to the DHCP Server for the IP stored in the `/intflash/dhcp/dhclient.leases` file.
- Out of Band management is enabled, except on the 5320 Series. With the exception of 5320-24T-4X-XT and 5320-24T-24S-4XE-XT, 5320 Series supports In Band management only.
- All ports are administratively enabled.
- IQAgent is enabled by default.
- Zero Touch Provisioning Plus (ZTP+) for ExtremeCloud IQ Site Engine onboarding is enabled by default.
- Zero Touch Fabric Configuration is initiated.
- After the Zero Touch Fabric establishes successfully, the onboarding VLAN 4048 is automatically assigned to onboarding I-SID 15999999.

For more information about IQAgent, see [ExtremeCloud IQ Support](#) on page 1084.

Factory Default Mode

The switch continues to support the boot configuration flag **boot config flags factorydefaults** to return an existing switch to factory default configuration.



Note

Zero Touch Deployment does not run on a switch returned to factory default configuration in this manner.

For more information, see [Boot Sequence](#) on page 247.

The factory default boot configuration flag supports the following behaviors:

Table 9: Factory default boot configuration flag parameters

Parameter	Action
no parameter used	Removes the configuration files, primary and secondary config file names, all user account passwords, digital certificates, IKE/OSPF/IS-IS keys, and SNMP communities. All ports are disabled and assigned to VLAN 1. License files are not removed.
<i>config-only</i>	Boots the switch with a blank configuration. This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities. All ports are disabled and assigned to VLAN 1. Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.
<i>reset-all-files</i>	Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset. If required, you must revoke the license file. The 30-day factory license does not apply to 4220 Series. Note: You can also use the unconfigure switch command to achieve the same behavior.

Table 9: Factory default boot configuration flag parameters (continued)

Parameter	Action
<i>zero-touch</i>	Boots the switch with a default configuration that enables Auto-sense. This parameter resets secure files but keeps the security mode and performs a ZTP+ configuration after reboot. License files are not removed.
<i>zero-touch-config-only</i>	<p>Boots the switch with a blank configuration that enables Auto-sense and performs a ZTP+ configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities.</p> <p>Note: Feature support can vary across platforms. For more information, see <i>Fabric Engine and VOSS Feature Matrix</i>.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p>

For more information, see:

- [Configure Boot Flags](#) on page 775
- [Enable Factory Default Behavior](#) on page 363
- [Return an Existing Configuration to Factory Defaults](#) on page 366

Zero Touch Provisioning Plus

Table 10: Zero Touch Provisioning Plus product support

Feature	Product	Release introduced
Zero Touch Provisioning Plus	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2.5
	VSP 7400 Series	VOSS 8.2.5

With zero touch functionality, switches are automatically discovered on the network within minutes of when they are connected.

Zero Touch Provisioning Plus (ZTP+) enables you to deploy and configure switches in ExtremeCloud IQ Site Engine with minimal server configuration and intervention. ZTP+ enabled switches send information, such as the serial number, software version, MAC, management IP, and port information to ExtremeCloud IQ Site Engine automatically.

When the switch powers on, the DHCP Client obtains the IP address and gateway from the DHCP server, discovers the Domain Name Server, and connects the switch to ExtremeCloud IQ Site Engine.

If the switch cannot obtain an IP address from a DHCP server, it automatically assigns an IPv4 link-local address (169.254.x.x) to the management OOB port and the management VLAN by using the switch MAC address. The formula uses the second last byte of the MAC address as the third octet in the link-local IP address, and the last octet in the IP address is always .1 for the management OOB port and .4 for the management VLAN.

If you configure a manual, static IP address, the switch removes the IPv4 link-local address.



Note

With the exception of 5320-24T-4X-XT and 5320-24T-24S-4XE-XT, 5320 Series does not have dedicated management ports. You can use front panel ports for management connectivity for these switches.

The following list provides an example of the link-local addressing:

- BaseMacAddr: 00:51:00:f9:**20**:00 where 0x20=32
- IP address for the management OOB port: 169.254.32.1

- IP address for the management VLAN: 169.254.32.4

ZTP+ uses HTTPS for communication between the switch and the ExtremeCloud IQ Site Engine server. The switch discovers the ExtremeCloud IQ Site Engine server by resolving the DNS name *extremecontrol.<domain-name>*.



Important

This feature requires a Zero Touch Deployment-ready configuration. For more information, see [Zero Touch Deployment](#) on page 91.

ZTP+ Phases of Operation

Zero Touch Provisioning Plus (ZTP+) auto-provisioning occurs in phases after you connect the switch to the network, if the switch is in factory ship state with no valid configuration saved on the device.

Connect

The Connect phase is the first phase of ZTP+ during which the switch connects to the ExtremeCloud IQ Site Engine server on the network. The ExtremeCloud IQ Site Engine server is discovered by resolving the DNS name *extremecontrol.<domain-name>*.

If the attempt is successful, the ExtremeCloud IQ Site Engine server responds with an **Accept** message. When connectivity is established, the switch communicates with the ExtremeCloud IQ Site Engine server securely and transmits information, such as its serial number, model number. The switch then progresses to the next phase of ZTP+.

Configuration

The next phase after the image upgrade is ZTP+ Configuration phase. During this phase, the switch queries the ExtremeCloud IQ Site Engine server for configuration updates, and initiates auto-provisioning by transmitting information, such as the image version, model name, and serial number. The switch then attempts to apply the configuration that is pushed from the ExtremeCloud IQ Site Engine server.

If the switch can still communicate with the ExtremeCloud IQ Site Engine server after the configuration is applied, the new configuration is automatically saved on the switch. The switch can be managed through ExtremeCloud IQ Site Engine using Simple Network Management Protocol (SNMP). However, if the configuration that is pushed from the ExtremeCloud IQ Site Engine server breaks switch connectivity to the ExtremeCloud IQ Site Engine server, the switch reboots without saving the configuration. After the switch reboots, ZTP+ onboarding restarts.

Any configurations pushed from the ExtremeCloud IQ Site Engine server to devices using the initial ZTP+ configuration push are not displayed in the **show log file detail** command output. The logs associated with the Cloud connector are logged internally to `state_machine.txt` and `ztp_plus.txt` files located in `/intflash/cc/cc_logs/`.

ExtremeCloud IQ Site Engine uses ZTP+ to configure the following items:

- Application of a subscription-license to a switch
- Auto-sense parameters, including a port data I-SID
- CLIP configuration including assigning the associated GRT for switch management
- DVR role and domain

**Note**

On a factory-defaulted device, you can enable a DVR leaf node without rebooting the switch. You can restore the factory default state using the **boot config flags factorydefaults zero-touch** command. For more information, see [Enable Factory Default Behavior](#) on page 363.

- I-SID names configuration
- Link Layer Discovery Protocol (LLDP) neighbor discovery

**Note**

Based on the LLDP discovery, you can use port templates on the ExtremeCloud IQ Site Engine server. Enabling or disabling LLDP is not supported.

- Login
- Port configuration

You can use port templates on the ExtremeCloud IQ Site Engine server to configure the following:

- AutoNegotiation
- SLPP Guard
- Admin
- Port Alias
- Speed
- Duplex
- Management I-SID on-boarding
- Network Time Protocol (NTP)
- SNMP
- VLANs

You can use port templates on the ExtremeCloud IQ Site Engine server to manage the following:

- PVID
- Tagged
- Untagged

**Note**

Custom port templates overwrite Auto-sense default configuration.

- VRF creation including VLAN assignment

Upgrade

After a successful connect to the ExtremeCloud IQ Site Engine server, the next phase of ZTP+ is the Upgrade phase. This phase verifies that the switch is running the image file version that is currently selected as the reference version on the ExtremeCloud IQ Site Engine server.

Image file validation is initiated by the switch. After a successful connect, the switch sends an image file upgrade request to the ExtremeCloud IQ Site Engine server with details on the current image file version. If the image file versions on the switch and the ExtremeCloud IQ Site Engine server match, no upgrade is initiated, and the switch moves to the next phase of ZTP+. If the ExtremeCloud IQ Site Engine server detects a different image file version, ZTP+ initiates the .tgz image file download from a specified URL location.

After a successful image upgrade, the switch reboots and reconnects to the ExtremeCloud IQ Site Engine server. If there are errors in the image upgrade process, an event is added to the server log. The switch then retries the image upgrade.

ZTP+ Considerations

The following considerations apply to Zero Touch Provisioning Plus (ZTP+) :

- Fabric configurations are not supported with ZTP+. After ZTP+ is configured, ExtremeCloud IQ Site Engine server can use Simple Network Management Protocol (SNMP) to remotely configure Fabric-related configurations on the switch using SNMP MIBs.
- Only the Out-of-Band (OOB) port or the Management VLAN interface are used to connect the ExtremeCloud IQ Site Engine server.

Configure ZTP+ Using the CLI

This section provides procedures to configure and manage Zero Touch Provisioning Plus (ZTP+) using the Command Line Interface (CLI).

After your device is onboarded, you have access to ExtremeCloud IQ Site Engine.



Note

You must configure a Segmented Management Instance to use ZTP+. For more information, see [Segmented Management Instance Configuration using the CLI](#) on page 123.

For information about onboarding switches, see <https://www.extremenetworks.com/support>.

Configure the ZTP+ Server Address

About This Task

Configure the ExtremeCloud IQ Site Engine server address for ZTP+. After you configure a server, the software immediately attempts ZTP+ onboarding with the configured server.

This procedure is optional. By default, the ExtremeCloud IQ Site Engine server is discovered by resolving the DNS name *extremecontrol.<domain-name>*.

Procedure

1. Enter Application Configuration mode:

```
enable  
  
configure terminal  
  
application
```
2. Configure the IPv4, IPv6, or DNS address of the server:

```
auto-provision server WORD<0-256>
```

View ZTP+ Status

About This Task

Use this procedure to verify the status of Zero Touch Provisioning Plus (ZTP+) on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Verify that ZTP+ is enabled:

```
show application auto-provision
```

Example

The following is an example output of the **show application auto-provision** command:

```
Switch:1>show application auto-provision  
=====
```

Auto-provision Info	
=====	
Operational Status	: Complete
Configured Server Address:	extremecontrol
extremecontrol IP Address:	192.0.2.34
Last successful sync	: Mon Sep 16 19:21:29 2024 EEST

Zero Touch Fabric Configuration

Table 11: Zero Touch Fabric Configuration product support

Feature	Product	Release introduced
Zero Touch Fabric Configuration	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0.5
LLDP Fabric Connect TLV	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

You can use Zero Touch Fabric Configuration to deploy Fabric-capable switches in a plug and play manner with no initial configuration. The switches form a new Fabric automatically or they can connect to an existing Fabric that is Auto-sense-capable, obtain an IP address and Domain Name System (DNS) information from a Dynamic Host Configuration Protocol (DHCP) server using the onboarding I-SID/VLAN, which then permits the system to automatically onboard to the management servers, such as ExtremeCloud IQ or ExtremeCloud IQ Site Engine, to conduct actual provisioning deployment of the switch. For more information about Auto-sense, see [Auto-sense](#) on page 21.

Zero Touch Fabric Configuration automatically configures Shortest Path Bridging MAC (SPBM) and IS-IS without user intervention if you boot the switch in Zero Touch Deployment-ready configuration mode, meaning you boot without a configuration file. Zero Touch Fabric Configuration uses LLDP to signal Fabric capability and exchanges SPB backbone VLAN IDs information to ensure seamless joining to any existing fabric deployment. The switches use the chassis MAC addresses as their system ID. To ensure participation in the correct area, newly joining ZTF switches listen to IS-IS update

packets for area information and automatically self-assign a nickname. For more information, see [Zero Touch Deployment](#) on page 91.

Zero Touch Fabric Configuration uses the port-based Auto-sense features, that enables all ports on the switch, by default, and all ports operate in Auto-sense mode. With the support of Auto-sense, Zero Touch Fabric Configuration onboards all ports on the switch to an existing network, without having to manually enable each port. Auto-sense automatically detects neighbor capabilities and performs the configuration on the port to reach the desired connectivity with the neighbor without user invention.

With Auto-sense functionality, ports on a switch can detect whether they connect to a Shortest Path Bridging (SPB) device, a Fabric Attach (FA) client, FA Proxy, Voice IP devices, or an undefined host, and then make the necessary configuration. For more information about Auto-sense, see [Auto-sense](#) on page 21.

If you start two nodes in a network without an existing configuration file, then Zero Touch Fabric Configuration, through Auto-sense, dynamically establishes an IS-IS adjacency between them. For more information, see [Establishing IS-IS Adjacencies](#) on page 103.

Default IS-IS Parameters

Zero Touch Fabric Configuration automatically configures the Shortest Path Bridging (SPB) and Intermediate System-to-Intermediate System (IS-IS) infrastructure to enable Fabric architecture on a switch. The system initializes the following items after you start the switch in Zero Touch Fabric Configuration mode:

- Enables Shortest Path Bridging MAC (SPBM).
- Creates a private VLAN 4048.
- Creates the Auto-sense onboarding I-SID 159999999.
- Assigns the Auto-sense onboarding I-SID 159999999 to private VLAN 4048 and also includes the management VLAN.



Note

As a best practice, use the onboarding I-SID for onboarding purposes and, whenever possible, configure a management VLAN or management CLIP on a different I-SID after the onboarding procedures have been successfully completed.

- Enables Auto-sense on all ports.
- Configures Auto-sense access ports and layer 2 trusted Auto-sense ports.
- Creates an SPBM instance.
- Enables IS-IS globally.

Parameter	Default value
SPBM instance	1
Primary B-VLAN	4051
Secondary B-VLAN	4052

Parameter	Default value
Manual area	Initialize to 00.1515.fee1.900d.1515.fee1.900d Note: You can change the manual area dynamically, without disabling IS-IS, only when the area is the Zero Touch Fabric Configuration area. When IS-IS is enabled, you cannot delete the last manual area.
Auto-sense onboarding I-SID	15999999

Establishing IS-IS Adjacencies

Zero Touch Fabric Configuration automatically triggers when the switch boots without a configuration file, the platform enables Intermediate System-to-Intermediate System (IS-IS) without a nickname or manual area. The system creates default backbone VLANs (B-VLAN) (4051 and 4052) and IS-IS manual area values. As a result, if you start two nodes in a network without an existing configuration file, then Zero Touch Fabric Configuration dynamically establishes an IS-IS adjacency between them.

The switch uses the Auto-sense functionality with the Zero Touch Fabric Configuration feature to establish the adjacency between two nodes. For more information about how and when the system tries to establish the adjacency, see [Auto-sense Port States](#) on page 28 .

- If you manually configure an SPBM instance on a node, then the system removes the SPBM instance that is dynamically created by Zero Touch Fabric Configuration. The system uses the LLDP Fabric Connect TLV to send user-defined B-VLANs to other nodes in the network. Only the first pair of B-VLANs is learned. If the switch already learned the B-VLANs from neighbor_A, the switch ignores the B-VLANs received from neighbor_B, if those are different.
- If a switch operating in Zero Touch Fabric Configuration mode in the network receives B-VLANs from a neighboring switch, which do not match the default B-VLANs configured through Zero Touch Fabric Configuration, the switch performs the following actions:
 - Disables IS-IS.
 - Deletes its VLANs.
 - Unassigns the B-VLANs.
 - Assigns the values received through LLDP Fabric Connect TLV.
 - Creates the corresponding VLANs.
 - Re-enables IS-IS and log a message on the console.

LLDP Fabric Connect TLV

The system uses the Link Layer Discovery Protocol (LLDP) Fabric Connect Type-Length-Value (TLV) to communicate B-VLANs and system IDs between nodes in the SPB cloud.

For more information about LLDP and its interaction with Fabric Attach, see [Link Layer Discovery Protocol \(802.1AB\) Fundamentals](#) on page 2466.

Table 12: LLDP Fabric Connect TLV Format

TLV Type	TLV Length	OUI	Subtype	Fabric Connect Capability	B-VLANs Number	B-VLAN-1	B-VLAN-2	System ID Length	System ID
7 bits	9 bits	3 octets	1 byte	1 byte	1 byte	2 bytes	2 bytes	1 byte	6 bytes

Table 13: LLDP Fabric Connect TLV Field Descriptions

TLV Field	Description
OUI	Specifies the Extreme Networks OID value (0xd88466).
Fabric Connect capability	Fabric Connect capability is enabled on all nodes that support Zero Touch Fabric Configuration. The value is 0 if the LLDP Fabric Connect TLV is not carrying any information in it. By default, the value is set to 1 on all ports.
B-VLANs number	Specifies the number of B-VLANs that the TLV can carry. The LLDP Fabric Connect TLV supports an unlimited number of B-VLANs, but Zero Touch Fabric Configuration sends two B-VLANs through the TLV. The value is 0 if the node is sending default B-VLAN values.
B-VLANs	Specifies the B-VLAN that is user-configured or dynamically learned from a neighbor node in the network.
System ID Length	Specifies the length (in bytes) of the System ID.
System ID	Specifies the IS-IS system ID.

Example: Create an IS-IS Adjacency between Auto-sense and non-Auto-sense VOSS Switches

Link Layer Discovery Protocol (LLDP) stations that connect to a local area network (LAN) advertise the station capabilities to each other, allowing the discovery of physical topology information for network management. When the system enables a switch as a Fabric Attach (FA) server in the Shortest Path Bridging (SPB) network, it receives LLDP messages from the FA Client and the FA Proxy devices using the LLDP Fabric Connect Type-Length-Value (TLV). For more information, see [LLDP Fabric Connect TLV](#) on page 103.

Following are the steps to create an Intermediate-System-to-Intermediate-System (IS-IS) adjacency between Auto-sense and non-Auto-sense VOSS switches, such as VSP

8600 Series or switches running releases that do not include Auto-sense support. For feature support information, see *Fabric Engine and VOSS Feature Matrix*.

On the non-Auto-sense VOSS switch:

1. On VSP 8600 Series only, disable the High Availability-CPU mode.
2. Enable the Shortest path bridging MAC (SPBM) mode.
3. Create an SPBM instance.
4. Create an SPBM B-VLAN.
5. Configure the port that links to the Auto-sense switch.

On the Auto-sense switch:

1. Enable the Shortest path bridging MAC (SPBM) mode.
2. Enable Auto-sense on the port that links to the non-Auto-sense switch.

Example

Create an IS-IS adjacency:

On a VSP 8600 Series switch:

```
enable
configure terminal
no boot config flags ha-cpu
boot config flags spbm-config-mode
spbm
router isis
spbm 1
spbm 1 b-vid 100,101 primary 100
spbm 1 nick-name 1.44.66
manual-area c1
exit
vlan create 100 type spbm-bvlan
vlan create 101 type spbm-bvlan
vlan members remove 1 1/3
interface gigabitEthernet 1/3
isis
isis spbm 1
isis enable
router isis enable
```

On the Auto-sense switch:

```
enable
configure terminal
boot config flags spbm-config-mode
interface gigabitEthernet 1/3
auto-sense enable
1 YYYY-MM-DD HH:MM:SS.622Z Switch - 0x00374589 - 00000000 GlobalRouter FA INFO Fabric Attach Assignments will
be rejected since ISIS is disabled.
1 YYYY-MM-DD HH:MM:SS.779Z Switch - 0x001dc703 - 00000000 GlobalRouter ISIS INFO B-VLANs (100,101) dynamically
learnt through LLDP. ISIS Restarted
1 YYYY-MM-DD HH:MM:SS.779Z Switch - 0x002d0609 - 00000000 GlobalRouter LLDP INFO New LLDP Neighbor Discovered
on interface 1/3
1 YYYY-MM-DD HH:MM:SS.954Z Switch - 0x00004727 - 00000000 GlobalRouter SNMP INFO SPBM detected adj INIT on
Port1/3, neighbor f873.a201.03df
1 YYYY-MM-DD HH:MM:SS.984Z Switch - 0x00004727 - 00000000 GlobalRouter SNMP INFO SPBM detected adj UP on
Port1/3, neighbor f873.a201.03df
show isis adjacencies
=====
                        ISIS Adjacencies
=====
```

INTERFACE	L STATE	UPTIME	PRI	HOLDTIME	SYSID	HOST-NAME	STATUS	AREA	AREA-NAME
Port1/3	1 UP	00:05:18	127	21	f873.a201.03df	VSP-8608	ACTIVE	HOME	

```
-----
Home: 1 out of 1 interfaces have formed an adjacency
-----
```



Segmented Management

[Migration to Segmented Management Instance](#) on page 110

[Interface Types](#) on page 111

[Management Applications](#) on page 117

[DHCP Client for Segmented Management Instance](#) on page 119

[Dynamic Change Options for Segmented Management Instance](#)

[Attributes](#) on page 121

[Segmented Management Instance Configuration using the CLI](#) on page 123

[Segmented Management Instance Configuration using EDM](#) on page 156

Table 14: Segmented Management Instance product support

Feature	Product	Release introduced
Segmented Management Instance - Management Interface CLIP	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 14: Segmented Management Instance product support (continued)

Feature	Product	Release introduced
Segmented Management Instance - Management Interface OOB	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.2 Supported on 5320-24T-4X-XT and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2
	VSP 7400 Series	VOSS 8.2
Segmented Management Instance - Management Interface VLAN	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Segmented Management Instance — ability to migrate VLAN or loopback IP address	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2
	VSP 7400 Series	VOSS 8.2

Table 14: Segmented Management Instance product support (continued)

Feature	Product	Release introduced
Segmented Management Instance — DHCP Client for Management Interface OOB or Management Interface VLAN	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6 OOB supported on 5320-24T-4X-XT and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2
	VSP 7400 Series	VOSS 8.2
Segmented Management Instance — sflow, Application Telemetry, and IPFIX	4220 Series	Fabric Engine 9.2 No IPFIX support
	5320 Series	Fabric Engine 8.8 IPFIX supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4 Support for IPFIX only
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8

Table 14: Segmented Management Instance product support (continued)

Feature	Product	Release introduced
Segmented Management Instance — ability to create a VLAN in routing mode	4220 Series	Fabric Engine 9.3
	5320 Series	Fabric Engine 9.3
	5420 Series	Fabric Engine 9.3
	5520 Series	Fabric Engine 9.3
	5720 Series	Fabric Engine 9.3
	7520 Series	Fabric Engine 9.3
	7720 Series	Fabric Engine 9.3
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.3
	VSP 7400 Series	VOSS 9.3

A Management Instance is required to provide access to specific management applications.

With Segmented Management, the Management plane (management protocols) is separated from the Control Plane (routing plane) from a process and data-path perspective. Segmented Management is the only method to manage switches. One or a combination of the following management interface/management instance types can be used:

- Out-of-Band (OOB) management IP address (IPv4 and IPv6)
- In-band Loopback/circuitless IP (CLIP) management IP address (IPv4 and IPv6)
- In-band management VLAN IP address (IPv4 and IPv6)



Important

The Segmented Management Instance provides support for management interfaces that transmit and receive packets directly to and from the system native Linux IP stack. Unlike a traditional management interface, for example, a CLIP in the GRT that is part of the OS networking IP stack, Segmented Management Instance interfaces do not route packets through the OS networking IP stack.

Segmented Management provides better security because you cannot reach the management instance from outside the VRF (in case of CLIP) or outside VLAN/I-SID (in case of management VLAN), and because it has a built-in firewall for the management plane. There is also more predictability with symmetric traffic flows for management traffic originating from and terminating on the switch, for instance:

- Sessions originated from switch (client mode) - Source IP of packets is determined based on Management IP stack routing table weights (configurable).
- Sessions connecting to switch (server mode) - Source IP is derived from session connection and reply will go out on management interface packet.

Migration to Segmented Management Instance

Segmented Management Instance Migration

You have two command options to change the attributes of a Management Instance:

- **convert** command
- **migrate-to-mgmt** command

As a best practice, use the **convert** command.

With the **convert** command, you can dynamically change the attributes of a Management Instance while you actively manage the switch over that same Management Instance without requiring the switch to reboot. This option also has rollback functionality to recover from unwanted changes.

For more information, see the following sections:

- [Dynamic Change Options for Segmented Management Instance Attributes](#) on page 121
- [Change Attributes for the Management Instance](#) on page 138

You can use the **migrate-to-mgmt** command to move a management VLAN to a different VLAN ID, or a management CLIP to a different VRF. However, if you use this option, you must reboot your switch after you save the configuration changes.

The following is an outline of the steps required for management migration using the **migrate-to-mgmt** command:

1. Configure a new or existing VLAN or CLIP management interface using Interface Configuration mode in the CLI (**interface vlan <vlan_id>** or **interface loopback <clip_id>**) or EDM.



Important

The IP interface and all routing protocols attached to the original VLAN are deleted post migration.

2. Add required routes to reach management services and subnets from the new interface.
3. Test connectivity to the new interface using ping and traceroute, and from the switch to management stations and servers.
4. Use the **migrate-to-mgmt** command from the new interface CLI mode.
5. Save the configuration and reboot.



Note

During boot, the migrate-to-mgmt settings are parsed and override the existing management interface with the new interface.

6. Access and manage the switch from the new interface.

For more information see, [Migrate a VLAN or CLIP IP address to the Segmented Management Instance](#) on page 123 (using CLI) or [Migrate an IP Address to a Segmented Management Instance](#) on page 156 (using EDM).

Interface Types

The Management Instance supports the following interface types:

- [CLIP](#) on page 111
- [OOB](#) on page 113
- [VLAN](#) on page 113



Note

The switch also supports a management instance specific to ExtremeCloud SD-WAN. For more information, see [Fabric Extend \(FE\) States](#) on page 32.

You can configure a maximum of three Segmented Management Instance interfaces, one of each type.

You can configure the route priority for the Segmented Management Instance. The Source IP default route priority is management CLIP (weight 100), then management VLAN (weight 200), then management OOB interface (weight 300). You can route packets through a different management interface than the default configuration, but you must add a specific static route or change the default weight of the management interface.



Note

If you change the default route weight, the management interface with the lowest weight value becomes the default route for all segmented management interface traffic.

You can configure the default topology IP for LLDP and SONMP advertisements. Both LLDP and SONMP advertise the same topology IP. SONMP supports only IPv4 addresses. If multiple IPv4 addresses are configured on an OOB or VLAN management interface, the advertised IP priority is static IP address, then DHCP IP address, then link-local IP address.

IPsec is not supported on Segmented Management Instance management interfaces.

CLIP

You can use this interface type for CLIP management network routing in a Fabric network or Layer 3 routing network.



Important

A CLIP Management Instance is not a management CLIP created in the GRT. You must create the CLIP Management Instance using the Segmented Management Instance configuration.

The following list defines the abilities of this interface type:

- You can assign a circuitless management IP (CLIP) address bound to a VRF.
- You can associate only one VRF ID with a CLIP Management Instance IP address.

**Note**

For switches that support only a single active VRF, the VRF must be the active VRF. For more information about VRF support on these models, see [VRF Lite Configuration Rules](#) on page 4078.

- The IP address is not bound to a physical network; it does not transmit nor receive IPv4 Address Resolution Protocol (ARP) or IPv6 Neighbor Discovery (ND) messages.
- You do not need to configure a default or static route. This interface type uses all routing information learned by protocols attached to the associated VRF.
- Packets can ingress on any port or VLAN that belongs to the VRF associated with the CLIP Management Instance.
- You must configure accept policies or configure inter-VRF route redistribution to access the CLIP Management Instance from a different VRF. Inter-VRF access is not permitted with traditional IP routing using OSPF, BGP, or RIP. Packets ingressing the switch from a VLAN that belongs to a different VRF without a configured accept policy will not reach the CLIP Management Instance IP address. For more information, see [Redistribution of CLIP Segmented Management Instance Examples](#) on page 154.
- If you migrate the current IS-IS IP source address to the CLIP Management Instance, after the upgrade the IS-IS source IP address moves to the CLIP Management Instance. You must configure a new GRT CLIP using a different IP address and assign that as the new IS-IS source IP.
- Advertisement of the IPv4 or IPv6 address for the CLIP Management Instance to IS-IS in the GRT occurs automatically. Advertisement of the IPv4 or IPv6 address in the VRF Layer 3 VSN bound to the CLIP Management Instance occurs automatically. You must configure route redistribution to advertise the CLIP Management Instance to different protocols. For more information, see [Redistribution of CLIP Segmented Management Instance Examples](#) on page 154.
- If you configure both a Management VLAN in routing mode and a Management CLIP, the switch prefers the Management CLIP when it selects the source IP address and outgoing interface for traffic initiated by management client applications, for example RADIUS reachability. For management client traffic, ensure connectivity exists in the Management CLIP context.

OOB

You can use this interface type for OOB management network routing, as an alternative to in-band network routing management.

The following list defines the abilities of this interface type:

- You can assign a management IP address bound to the Out-of-Band (OOB) interface.
- You can associate only one OOB interface with an OOB Management Instance IP address.
- You can use a VLAN IP address including primary and secondary IP address on a local VLAN.
- The Dynamic Host Configuration Protocol (DHCP) Client can request an IPv4 address for the OOB Management Instance interface.
- You must configure a default or static route to reach the next-hop gateway; no routing protocol information is used to access off-link networks.
- You can configure only Layer 3 networking parameters in the Management Instance Configuration mode (**mgmt oob**) in CLI.
- You can configure only Layer 1 and Layer 2 networking parameters in the mgmtEthernet Interface Configuration mode (**interface mgmtEthernet mgmt**) in CLI.

VLAN

The Management VLAN interface type supports two modes:

- non-routing
- routing

You can use this interface type for management of Layer 2 and Layer 3 switches or for Zero-Touch onboarding of newly deployed devices. When the Management VLAN is in routing mode, packets incoming on other VLANs are routed from the underlying VLAN to the Management interface.

For more information on Zero-Touch onboarding, see [Zero Touch Capabilities](#) on page 20.

Abilities Common to Both Modes

The following list defines the abilities of this interface type common to both routing and non-routing modes:

- You can assign a Management Instance IP address to an inband VLAN.
- You can associate only one VLAN ID with a VLAN Management Instance IP address.

Abilities in Non-Routing Mode

The following list defines the abilities of this interface type specific to non-routing mode:

- For the VLAN Management Instance to take route priority when used in conjunction with the CLIP Management Instance, you must configure a default route for the VLAN Management Instance with a value lower than 100, or configure static routes for direct communication over the VLAN Management Instance and management networks.
- You can configure a Management Instance VLAN on a DvR Leaf node by specifying the I-SID. For more information, see [Management I-SID Assignment to DvR Leaf](#) on page 932.
- The DHCP Client can request an IPv4 address for the VLAN Management Instance interface.
- The interface resides on the physical VLAN segment, behaving as a host for sending and receiving IPv4 ARP and IPv6 ND messages.
- You must configure a default or static route to reach the next-hop gateway; no routing protocol information is used to access off-link (other subnets) networks.
- No internal routing occurs between the VLAN Management Instance and other non Management Instance VLANs. The VLAN Management Instance does not route to or from the GRT. Packets must ingress on one of the ports in the VLAN Management Instance.

Packets sent to the VLAN Management Instance IP address must ingress the switch from a VLAN or network-to-network interface (NNI) port (or contain the VLAN ID) associated with the VLAN Management Instance. The system does not route packets between the network operating system (NOS) routing VLAN and the VLAN Management Instance.

If you configure the same VLAN ID for NOS routing and for the VLAN Management Instance, the NOS routing stack transmits and receives all ARP, ND, and ICMP packets. In this scenario, the packets are only counted and shown in the NOS routing KHI port statistics. The management statistics and KHI management statistics do not count or show the packets.

- You can use secondary IP addresses if the VLAN Management Instance has its own VLAN ID that is not shared with a local VLAN.
- You can bind the VLAN Management Instance to an I-SID, which bridges all management traffic to a single I-SID in a Fabric network. Also, other normal VLAN related operations such as VLAN port member changes are valid.
- Bridged management traffic must ingress on the VLAN or I-SID.
- The VLAN Management Instance can be routed by upstream routers.
- You cannot configure a management VLAN in non-routing mode from the VLAN interface.
- The interface type displays as `VLAN-L2` in show command output.
- You can convert a management VLAN in non-routing mode to routing mode.

Abilities in Routing Mode

The following list defines the abilities of this interface type specific to routing mode:

- You can only create, or delete, a Management VLAN in routing mode from the VLAN interface. The Management VLAN inherits the VRF, MAC offset, and IPv4 and IPv6 address configuration from the VLAN interface.
- You cannot enable the management routing mode on a brouter port interface.
- You cannot enable the management routing mode on a VLAN with a One IP Anycast IP Gateway interface.
- You cannot configure a management VLAN in routing mode from the Management Instance configuration mode. No manual configuration is allowed in routing mode.
- You do not need to configure a default route. This interface type uses all routing information learned by protocols attached to the associated VRF.
- Packets can ingress on any port or VLAN that belongs to the VRF associated with the Management VLAN.
- You cannot create additional IPv4 addresses on the VLAN after you enable the Management VLAN in routing mode. You can create additional IPv6 addresses.
- The switch does not save the Management VLAN configuration in the configuration file or display it in the running configuration. The switch does display `mgmt enable` under the VLAN interface. The interface type displays as `VLAN-L3` in show command output.
- You can convert a Management VLAN in routing mode to non-routing mode.
- If you configure both a Management VLAN in routing mode and a Management CLIP, the switch prefers the Management CLIP when it selects the source IP address and outgoing interface for traffic initiated by management client applications, for example RADIUS reachability. For management client traffic, ensure connectivity exists in the Management CLIP context.

Coexistence Restrictions

IPv4 and IPv6 address coexistence for both a NOS routing VLAN and VLAN Management Instance is supported, however you must manually match both IP address configurations between the VLANs.

If you configure the VLAN Management Instance with a manual IPv4 address and a DHCP IPv4 address first, you cannot add a IPv4 address to a NOS routing VLAN.

If you configure the VLAN Management Instance with an IPv6 address first, you can only add one IPv6 global address to a NOS routing VLAN.

The following restrictions apply when a VLAN Management Instance coexists with a port-based VLAN or with a brouter port:

- If you want a dual stack IPv4 and IPv6 coexistence between a NOS VLAN and VLAN Management Instance, you must configure the same IPv4 and IPv6 addresses on the VLAN Management Instance and on the NOS VLAN.

You cannot configure the VLAN Management Instance with both IPv4 and IPv6 and configure the NOS VLAN with IPv4 or IPv6 only.

- If you disable NOS routing for IPv4, then you must disable routing for IPv6, and vice versa.
- If you configure the same VLAN ID for a local VLAN and for the VLAN Management Instance, you cannot use secondary IP addresses on the local VLAN. You cannot use the same VLAN ID if the local VLAN already uses secondary IP addresses.

For switches that support a single active VRF, the management VLAN must be part of the active VRF. For more information about VRF support on these models, see [VRF Lite Configuration Rules](#) on page 4078.

Configuration Example - Coexistence with Port-Based VLAN

The following example shows how a VLAN Management Instance in non-routing mode can be configured to share the same IP address as a routing port-based VLAN.

You can configure the NOS VLAN first, and then configure the VLAN Management Instance, or in reverse order. You can remove or add the coexistence at any time.



Note

With the coexistence between NOS routing stack and the VLAN Management Instance, packets sent to the VLAN Management Instance IP address must ingress the switch from a VLAN port (or contain the VLAN ID) associated with the VLAN Management Instance. The system does not route packets between the NOS routing VLAN and the VLAN Management Instance.

IPv4

```
vlan create 10 type port-mstprstp 0
vlan members add 10 1/1
interface vlan 10
ip address 192.0.2.0/24
exit
mgmt vlan 10
ip address 192.0.2.0/24
ip route 0.0.0.0/0 next-hop 192.0.2.1
enable
```

IPv6



Note

IPv6 does not apply to 7830 Series.

```
vlan create 10 type port-mstprstp 0
vlan members add 10 1/1
interface vlan 10
ipv6 interface address 2001:DB8::/32
ipv6 interface enable
exit
mgmt vlan 10
ipv6 address 2001:DB8::/32
ipv6 route 0::0/0 next-hop 2001::1
enable
```

Configuration Example - Coexistence with Brouter Port

The following example shows how the VLAN Management Instance can be configured to share the same IP address as a brouter interface.

You must configure the brouter interface before you enable the VLAN Management Instance. When the VLAN Management Instance is enabled, you must disable the VLAN Management Instance before you disable the brouter port.

IPv4

```
interface GigabitEthernet 1/1
no shutdown
brouter port 1/1 vlan 10 subnet 192.0.2.0/24
mgmt vlan 10
ip address 192.0.2.0/24
enable
```

IPv6

```
interface GigabitEthernet 1/1
no shutdown
ipv6 interface vlan 10
ipv6 interface address 2001:DB8::/32
ipv6 interface enable
mgmt vlan 10
ipv6 address 2001:DB8::/32
enable
```

Management Applications

The Segmented Management Instance provides support for management interfaces that transmit and receive packets directly to and from the system native Linux IP stack. Unlike a traditional management interface, for example, a CLIP in the GRT that is part of the networking IP stack, Segmented Management Instance interfaces do not route packets through the networking IP stack.

The following management applications use the Segmented Management Instance directly to transmit or receive packets with segmented management interfaces and addresses.

Segmented Management Instance Applications and Protocols	Client	Server	IPv4	IPv6
Digital Certificates	Yes		Yes	
DHCP Client	Yes		Yes	
DNS	Yes		Yes	Yes
FTP	Yes	Yes	Yes	Yes
HTTP/HTTPS		Yes	Yes	Yes
IQAgent	Yes		Yes	
NTPv4	Yes	Yes	Yes	Yes

Segmented Management Instance Applications and Protocols	Client	Server	IPv4	IPv6
OVSDB protocol support for VXLAN Gateway	Yes	Yes	Yes	
Ping	Yes	Yes	Yes	Yes
RADIUS	Yes		Yes	Yes
RADIUS Security (RADSec)	Yes		Yes	Yes
Representational State Transfer Configuration Protocol (RESTCONF)		Yes	Yes	
SSH/SCP/SFTP	Yes (SSH only)	Yes	Yes	Yes
Syslog	Yes		Yes	Yes
TACACS+	Yes		Yes	
Telnet	Yes	Yes	Yes	Yes
TFTP	Yes	Yes	Yes	Yes
Traceroute	Yes	Yes	Yes	Yes

The following management applications do not use the Segmented Management Instance to transmit or receive packets, but can integrate with segmented management interfaces and addresses.

Applications and Protocols	IPv4	IPv6
Application Telemetry	Yes	Yes
IPFIX	Yes	
Link Layer Discovery Protocol (LLDP)	Yes	Yes
SynOptics Network Management Protocol (SONMP)	Yes	
sFlow	Yes	
Remote Network Monitoring version 2 (RMON2)	Yes	

DHCP Client for Segmented Management Instance

To support Zero Touch Deployment, a DHCP Client is used for the Segmented Management Instance VLAN management interface (in non-routing mode only) or Out-of-Band (OOB) management interface.

**Note**

You must disable the DHCP Client if you plan to create a Management VLAN in routing mode.

The DHCP Client configuration supports a VLAN mode, OOB mode, and a cycle mode. DHCP Client cycle mode alternates IP address requests between the VLAN management interface and OOB management interface until an IP address is obtained on one of the interfaces. Priority is given to the OOB management interface.

You can also manually configure the DHCP Client to request an IPv4 address from a DHCP server for the In-band VLAN management interface, or the OOB management interface, or to cycle requests until an IP address is obtained on a VLAN or OOB management interface. The DHCP Client supports IPv4 addresses only, and cannot be enabled on multiple management interfaces simultaneously.

**Note**

If a default route is configured on an OOB or VLAN management interface, and then you configure DHCP so that it replaces the default route, the original default route is restored if you disable DHCP.

However, if the DHCP default route is updated or deleted after it is created by DHCP, the default route will not be replaced by the original route when DHCP is disabled.

DHCP Client Restrictions

DHCP Client for the Segmented Management Instance supports IPv4 addresses only, and cannot be enabled on multiple management interfaces simultaneously. The DHCP Client only supports the in-band VLAN management interface, or the OOB management interface, or to cycle requests on the the VLAN then OOB management interface until an IP address is obtained on one of the interfaces.

**Note**

The DHCP Client is disabled by default on previously configured or upgraded switches.

The DHCP Client is enabled by default in cycle mode when:

- The switch ships directly from manufacturing with Fabric Engine Release 8.6 or later.
- The primary and secondary configuration file is not on the switch.
- The primary and secondary configuration file fail to load on the switch.

The DHCP Client is not available if RMON2 is configured on a Management Instance, and RMON2 is not available if the DHCP Client is configured on a Management Instance.

When DHCP is enabled on a Management Instance interface, the DHCP Client initial broadcast discovery packet and initial response from the DHCP server are not counted or shown in KHI management statistics for the management interface. Only the packets after the DHCP IP address assignment completes are counted and shown. After an IP address is assigned, a UDP socket opens and packets are counted on the interface.

If you change the DHCP Client configuration between management VLAN, OOB, or cycle, the default route provided by the DHCP server might delete and add with a different nexthop or network. DHCP Client configuration changes can cause interruptions to existing management connections.

DHCP static routes are not saved in the configuration file or displayed in **show running-config**. You can view DHCP static routes with **show mgmt ip route static**. If the DHCP Client adds a default route to an interface, the previous default route is deleted. If you modify a default route created by the DHCP Client, the route type output of **show mgmt ip route static** changes from DHCP to STATIC. You can save the modified to static default route to the configuration file, but on reboot the DHCP Client deletes the modified default route and restores the default static route the DHCP server specifies.

DHCP Option 43

DHCP option 43 requests specific vendor options from the DHCP server. Only sub-option 226 (EXTREME.cloudiq-ip) is supported to change the value of the ExtremeCloud IQ server IP address on the switch.

With the support of DHCP option 43, DHCP can dynamically configure the IP address of a private/non-public ExtremeCloud IQ server for zero touch deployments when the default ExtremeCloud IQ server (hac.extremecloudiq.com) is not desired.

For information about configuring the switch to support ExtremeCloud IQ, see [ExtremeCloud IQ Support](#) on page 1084.

DHCP Option 43 Configuration Examples

This section provides examples to configure DHCP Option 43 on a Linux server and on Windows Server.

ISC DHCP Server configuration on Linux:

```
/etc/dhcp/dhcpd.conf
default-lease-time 60;
max-lease-time 7200;
option space EXTREME;
option EXTREME.cloudiq-ip code 226 = ip-address;
```

```

class "Edge-without-POE" {
    match if (option vendor-class-identifier = "EXTREME");
    vendor-option-space EXTREME;
    option EXTREME.cloudiq-ip 10.16.231.131;
}

subnet 30.30.30.0 netmask 255.255.255.0 {
    pool {
        range 30.30.30.10 30.30.30.20;
        allow members of "Edge-without-POE";
    }
    option domain-name-servers 10.1.10.1;
    option domain-name "labs.extremenetworks.com";
    option routers 30.30.30.250;
    default-lease-time 3600;
}

```

Windows Server configuration:

1. Go to scope options for defined DHCP pool.
2. Enter the following for Option 43: e2 04 0a 10 e7 83

Value	Description
e2 04	vendor ID prefix (e2 is the hexadecimal value of the code 226 used to identify sub-option EXTREME.cloudiq-ip and 04 the hexadecimal value of the length of an IP address in bytes)
0a 10 e7 83	IP address 10.16.231.131 converted to hexadecimal

Dynamic Change Options for Segmented Management Instance Attributes

You can dynamically change the attributes of a Management Instance while you actively manage the switch over that same Management Instance without requiring the switch to reboot. For example, if you onboard the switch using VLAN 4048, you can change that Management Instance VLAN to a new VLAN.

You can change the following attributes for the Management Instance:

- Management Instance VLAN:
 - change between Management VLAN modes (routing to non-routing, or non-routing to routing)
 - default gateways (IPv4 or IPv6)
 - IP address (IPv4 or IPv6)
 - I-SID (on a DvR Leaf)
 - tagged ports
 - untagged ports
 - VLAN ID
 - VRF
- Management Instance CLIP:
 - IP address (IPv4 or IPv6)
 - VRF

- Management Instance Out-of-Band (OOB)
 - IP address (IPv4 or IPv6)
 - default gateway (IPv4 or IPv6)

Operational Considerations

The following are operational considerations when you change Management Instance attributes using the **convert** or **convert-to-routing** commands:

- You cannot change parameters for more than one Management Instance operation at a time. You must issue the **mgmt convert-commit** command before you use the **convert** or **convert-to-routing** commands again for either the same or a different Management Instance.
- If you attempt to change attributes for an existing Management VLAN in non-routing mode, you cannot configure ports-tagged, ports-untagged, and I-SID parameters. Make configuration changes to the existing Management VLAN first before you use the **convert** command. If you change the switch to a DvR leaf node, you can change the I-SID parameter.
- If you attempt to change attributes for a Management VLAN in non-routing mode and the VLAN does not exist, a VLAN is automatically created in the background. You can specify ports-tagged, ports-untagged, and I-SID parameters to be associated with this new VLAN. The new VLAN is assigned to the default Spanning Tree Group, which is 0.
- If you attempt to change attributes for a Management VLAN in non-routing mode, the VLAN does not exist, and you do not specify ports-tagged, ports-untagged, or I-SID parameters, this is a special case. If any untagged ports in the old VLAN have dynamic Address Resolution Protocol (ARP) entries, then these ports automatically move from the old VLAN to the new VLAN. For example, VLAN 200 has port members 1/1, 1/2, and 1/3. ARP entries are configured on ports 1/1 and 1/2. VLAN 300 is created in the background and only ports 1/1 and 1/2 automatically move to this new VLAN.

If an MLT ID is associated with the old VLAN, the association is removed and re-added to the new Management VLAN ID.

- If you convert between Management VLAN modes, coexistence is maintained if it exists on the original Management VLAN.
- If you convert a Management VLAN in routing mode to non-routing mode, you must specify the gateway as part of the **convert** command syntax to install a default route in the new Management VLAN interface.
- The following list items apply if you convert a Management VLAN in non-routing mode to routing mode:
 - The MAC offset value is inherited from the VLAN.
 - An I-SID configuration is not converted.
 - Conversion is not supported on a DvR Leaf or on a VLAN with Secondary IP Interfaces.

- The IP address for the new Management VLAN is determined in the following order of preference:
 1. Specify in the **convert-to-routing** command syntax.
 2. Maintain the IP address from the new VLAN.
 3. Move the old VLAN IP address if the new VLAN has no address configured.



Note

When converting from a VLAN with a DHCP address, you must include an IP address in the **convert-to-routing** command syntax if one does not already exist on the VLAN. A Management VLAN in routing mode requires a static address.

- If the IP address remains the same, or is in the same subnet as the old IP address, static routes on the non-routing Management VLAN interface move to the NOS routing table if they do not already exist. The switch does not copy the static routes to the new Management VLAN routing interface because it relies on the NOS routing table. If you specify a gateway in the **convert-to-routing** command, the switch installs the static routes with that new default gateway. If you change the IP address to a different subnet, the switch only copies the static routes if you specify the gateway in the command syntax.
- The following list items apply when you change static routes attribute on a Management VLAN in non-routing mode or Management OOB interface:
 - If you provide the static route next-hop gateway, all next-hop gateway direct to the new gateway.
 - If you do not provide the static route next-hop gateway and the new subnet is the same as the old subnet, all routes are re-added as is.
 - If you do not provide a gateway and the new subnet is different, routes are discarded.
- If you attempt to change the VRF attribute for a Management CLIP, or a Management VLAN with the **convert-to-routing** command, but the VRF does not exist, a VRF is automatically created in the background. For the VRF to function properly, you must configure either SPBM Layer 3 VSN or IP interfaces and routing protocols.

The best practice is to configure and test VRF connectivity before you use the **convert** command.

- Dynamic routes added by DHCP convert to static routes.

Segmented Management Instance Configuration using the CLI

This section provides procedures to configure segmented management instance using the command line interface (CLI).

Migrate a VLAN or CLIP IP address to the Segmented Management Instance

Perform this procedure to migrate a new routing VLAN with a new IP address or a new loopback IP address under a different VRF to the Segmented Management

Instance. Alternatively, you can also use the **convert** command. For more information, see [Change Attributes for the Management Instance](#) on page 138.



Important

Choose a VLAN that does not have an IP interface on it. The upgrade process removes the IP configuration and network connectivity will be impacted.

About This Task



Note

Do not migrate interfaces used for routing purposes, for example, where you configure Layer 3 routing protocols.

This command does not apply to the OOB or mgmtEthernet interface. Releases that support this migration procedure automatically move the IP address on the mgmtEthernet interface from the routing stack to the Segmented Management Instance during the upgrade to this release.

Procedure

1. Enter Interface Configuration mode for either a VLAN or loopback interface:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059> or interface loopback <1-256>
```

2. Select the interface address for migration:

```
migrate-to-mgmt
```

3. View the designated interface addresses selected for migration:

```
show mgmt migration
```

4. Save the configuration selected for migration:

```
save config
```

Example

Identify an IP address currently assigned to an inband VLAN to migrate to the Management VLAN. The example assumes you already identified a CLIP address. The VRF column in **show mgmt migration** indicates where the interface is being moved from.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 20
Switch:1(config-if)#migrate-to-mgmt
Switch:1(config-if)#show mgmt migration
```

```
=====
                                Mgmt Migration Information
=====
```

IFINDEX	DESCR	VRF	IPV4	IPV6
1344	CLIP-1	GlobalRouter	192.0.2.102/32	10:0:0:0:0:0:1/128

```
2068      VLAN-20      GlobalRouter  198.51.100.6/24      20:0:0:0:0:0:1/64

2 out of 2 Total Num of mgmt migrate entries displayed
-----

Switch:1(config-if)#save config
```

Configure a Segmented Management Instance Using quick-config-mgmt Utility

Use the following procedure to run the quick-config-mgmt utility script to ease the transition to the Segmented Management Instance.

The quick-config-mgmt utility recognizes existing configuration. For the OOB Management Instance, you can overwrite the existing configuration only. However, for the VLAN Management Instance, you can overwrite the existing configuration, or you can migrate the existing configuration to a coexistence of IP on both the routing VLAN and the management VLAN.

The quick-config-mgmt utility supports the following:

- IPv4 only
- only one interface at one time
- Out-of-Band management and In-Band VLAN management

About This Task

You can use this procedure to help you transition to a new Segmented Management Instance. You can configure IPv4, static routes, and DHCP support for the Out-of-Band (OOB) Management Instance or for the In-Band VLAN Management Instance. If configuration exists for the interface type you selected, you are prompted to replace the configured interface or to quit the utility.



Important

If you configure DHCP, any other running DHCP instance is stopped and a new DHCP instance is created on the interface. This might cause loss of connectivity.

The default values are given in square brackets. You can input your values at the prompt or you can press `Enter` to accept the default values.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enter the following command to start the utility:

```
quick-config-mgmt
```



Important

If DHCP mode cycle is enabled, the following warning message displays to inform you that DHCP client will be disabled, if you continue.

Continuing will disable dhcp and may affect your connectivity to the DUT. Do you want to continue? y/n [n]:

Examples

The following examples show outputs from the **quick-config-mgmt** utility.

Configure the OOB Management Instance:

```
Switch:1(config)#quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the
configuration.

Once the basic parameters are configured, additional configuration can proceed using
other management interfaces.
Press q to abort at any time.
Management interface types:
    1 - Out of band management port
    3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]:
Please enter the Management Address IPv4 address, "d" for DHCP configuration or "q" to
quit [192.0.2.2]:
Please enter the Management Address Mask IPv4 address or "q" to quit [255.255.255.0]:
Please enter the Default Gateway Address IPv4 address, 0.0.0.0 for no default gateway,
or "q" to quit [192.0.2.5] :
Management interface created successfully
```

Replace an existing OOB Management Instance configuration:

```
Switch:1(config)#quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the
configuration.

Once the basic parameters are configured, additional configuration can proceed using
other management interfaces.
Press q to abort at any time.
Management interface types:
    1 - Out of band management port
    3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]:
MGMT OOB is already configured.
Continuing may remove parts or all of current config.
Do you want to continue? y/n [y]:
Please enter management interface type or "q" to quit. [1]:
Please enter the Management Address IPv4 address, "d" for DHCP configuration or "q" to
quit [192.0.2.2]:
Please enter the Management Address Mask IPv4 address or "q" to quit [255.255.255.0]:
Please enter the Default Gateway Address IPv4 address, 0.0.0.0 for no default gateway,
or "q" to quit [192.0.2.5] :
Management interface created successfully
```

Configure the In-band port-based VLAN Management Instance by removing parts of or all of the existing VLAN configuration:

```
Switch:1(config)#quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the
configuration.

Once the basic parameters are configured, additional configuration can proceed using
other management interfaces.
Press q to abort at any time.
Management interface types:
  1 - Out of band management port
  3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]: 3
MGMT VLAN is already configured.
Continuing may remove parts or all of current config.
Do you want to continue? y/n [n]: y
Please enter VLAN ID (2-4059) or "q" to quit [4059]: 2
VLAN 2 is already in use.
Do you want to re-use existing vlan configuration? y/n/q: [n]
This option will remove all current config on VLAN 2.
Please enter port to be added to the in-band management VLAN or "q" to quit [1/1]:
Please enter the Management Address Mask IPv4 address or "q" to quit [255.255.255.0]:
Please enter the Default Gateway Address IPv4 address, 0.0.0.0 for no default gateway,
or "q" to quit [192.0.2.5] :
Management interface created successfully
```

Configure the In-band port-based VLAN Management Instance by reusing the existing VLAN configuration:

```
Switch:1(config)#quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the
configuration.

Once the basic parameters are configured, additional configuration can proceed using
other management interfaces.
Press q to abort at any time.
Management interface types:
  1 - Out of band management port
  3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]: 3
MGMT VLAN is already configured.
Continuing may remove parts or all of current config.
Do you want to continue? y/n [n]: y
Please enter VLAN ID (2-4059) or "q" to quit [4059]: 2
VLAN 2 is already in use.
Do you want to re-use existing vlan configuration? y/n/q: [y]
Please enter port to be appended to the in-band management VLAN or leave empty
to keep currently configured ports or "q" to quit []: 1/1
Please enter the Management Address IPv4 address, "d" for DHCP configuration or "q" to
quit [192.0.2.2]:
Please enter the Management Address Mask IPv4 address or "q" to quit [255.255.255.0]:
Please enter the Default Gateway Address IPv4 address, 0.0.0.0 for no default gateway,
or "q" to quit [192.0.2.5] :
Management interface created successfully
```

Configure the In-band port-based VLAN Management Instance by reusing the existing VLAN configuration when IP address is configured and coexistence of mgmt and routing on same VLAN is desired:

```
Switch:1(config)#quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the
configuration.

Once the basic parameters are configured, additional configuration can proceed using
other management interfaces.
Press q to abort at any time.
Management interface types:
  1 - Out of band management port
  3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]: 3
MGMT VLAN is already configured.
Continuing may remove parts or all of current config.
Do you want to continue? y/n [n]: y
Please enter VLAN ID (2-4059) or "q" to quit [4059]: 2
VLAN 2 is already in use.
Do you want to re-use existing vlan configuration? y/n/q: [y]
Please enter port to be appended to the in-band management VLAN or leave empty
to keep currently configured ports or "q" to quit []: 1/1
IP address is already configured on VLAN 2.
Do you want to configure coexistence of mgmt and routing on the same vlan? y/n/q: [y]
Please enter the Management Address IPv4 address, "d" for DHCP configuration or "q" to
quit [192.0.2.2]:
Please enter the Management Address Mask IPv4 address or "q" to quit [255.255.255.0]:
Please enter the Default Gateway Address IPv4 address, 0.0.0.0 for no default gateway,
or "q" to quit [192.0.2.5] :
Management interface created successfully
```

Configure the In-band port-based VLAN Management Instance by reusing the existing VLAN configuration when IP address is configured but coexistence of mgmt and routing on same VLAN is not desired:

```
Switch:1(config)#quick-config-mgmt
Welcome to the management interface setup utility.
You will be requested for information to initially configure the switch.
When finished the information will be applied and stored as a part of the
configuration.

Once the basic parameters are configured, additional configuration can proceed using
other management interfaces.
Press q to abort at any time.
Management interface types:
  1 - Out of band management port
  3 - In-band port-based VLAN
Please enter management interface type or "q" to quit. [1]: 3
MGMT VLAN is already configured.
Continuing may remove parts or all of current config.
Do you want to continue? y/n [n]: y
Please enter VLAN ID (2-4059) or "q" to quit [4059]: 2
VLAN 2 is already in use.
Do you want to re-use existing vlan configuration? y/n/q: [y]
Please enter port to be appended to the in-band management VLAN or leave empty
to keep currently configured ports or "q" to quit []: 1/1
IP address is already configured on VLAN 2.
Do you want to configure coexistence of mgmt and routing on the same vlan? y/n/q: [n]
Management interface created successfully
```

Create a Segmented Management Instance

You must create a Management Instance to gain access to specific management applications. After you create the Management Instance, you can add an IP address to it and configure route redistribution to advertise reachability of the Management Instance to the rest of the network.

About This Task

The Management Instance supports different management interface types. When you create the Management Instance, you specify the interface type and the switch automatically creates the appropriate instance ID for that type.

Each Management Instance supports a IPv4 and IPv6 (global scope) management address for use by management applications.

This procedure creates a management CLIP, management OOB, or management VLAN in non-routing mode. To create a management VLAN in routing mode, see [Create a Management VLAN In Routing Mode](#) on page 130.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create the Management Instance required for your deployment:

- a. To create a management CLIP:

```
mgmt clip [vrf WORD<1-16>]
```

**Note**

If you do not specify a VRF, the management CLIP uses the GRT.

OR

- b. To create a management OOB:

```
mgmt oob
```

OR

- c. To create a management VLAN in non-routing mode and associate it with an existing port-based VLAN:

```
mgmt vlan <2-4059>
```

3. Enable the Management Instance:

```
enable
```

Example

Create and enable a Management CLIP:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config)#mgmt clip
Switch:1(mgmt:clip)#enable
```

Create and enable a Management OOB:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt oob
Switch:1(mgmt:oob)#enable
```

Create and enable a Management VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan 20
Switch:1(mgmt:vlan)#enable
```

Create a Management VLAN In Routing Mode

Before You Begin

Configure an IP address on the VLAN before you enable the management VLAN interface.

About This Task

You must create a Management Instance to gain access to specific management applications. Perform this procedure to create a Management VLAN in routing mode. To create a Management VLAN in non-routing mode, see [Create a Segmented Management Instance](#) on page 129 .

This Management VLAN inherits the VRF, MAC offset, and IPv4 and IPv6 address configuration from the VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```
2. Enable the management routing mode on the VLAN:

```
mgmt enable
```
3. Confirm the configuration:

```
show mgmt interface
```

Example

Create a Management VLAN in routing mode on VLAN 4.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 4
```



```
Switch:1(config-if)#mgmt enable
Switch:1(config-if)#show mgmt interface
```

Mgmt Interface Information								
INST	DESCR	TYPE	ADMIN	VLAN	ISID	PORT	VRF	PHYSICAL
4	Mgmt-vlan	VLAN-L3	enable	10	-	-	GlobalRouter	40:88:2f:e6:61:00

```
1 out of 1 Total Num of mgmt interfaces displayed
```

Delete a Segmented Management Instance

Use this task to delete a Management Instance. Deleting the Management Instance removes the IP address, and changes the associated VRF for a management CLIP.

This procedure deletes a management CLIP, management OOB, or management VLAN in non-routing mode. To delete a management VLAN in routing mode, see [Delete a Management VLAN In Routing Mode](#) on page 131.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Delete the Management Instance:

```
no mgmt {clip | oob | vlan}
```

Delete a Management VLAN In Routing Mode

About This Task

Perform this procedure to delete a Management VLAN in routing mode. To delete a Management VLAN in non-routing mode, see [Delete a Segmented Management Instance](#) on page 131.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```
2. Disable the management routing mode on the VLAN:

```
no mgmt enable
```

Configure the DHCP Client for a Segmented Management Instance

Use this task to configure the DHCP Client to obtain an IPv4 address for the Management Instance VLAN interface or Out-of-Band interface.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable and configure the DHCP Client for a management interface:

```
mgmt dhcp-client {cycle | oob | vlan}
```

Example

The following example configures the DHCP Client to cycle IPv4 requests for the management OOB interface, and then the In-Band management VLAN interface; priority is given to the OOB interface. The system cycles attempts until one management interface receives an IP address from the DHCP server:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt dhcp-client cycle
```

Variable Definitions

The following table defines parameters for the **mgmt dhcp-client** command.

Variable	Value
<i>cycle</i> Note: Exception: only supported on hardware with a physical management port	DHCP Client cycles IP requests for in-band VLAN and Out-of-Band management interfaces.
<i>oob</i> Note: Exception: only supported on hardware with a physical management port	DHCP Client requests an IP address for the Out-of-Band management interface.
<i>vlan</i>	DHCP Client requests an IP address for the VLAN management interface.

Configure an IP Address for a Segmented Management Instance

Use this task to add an IPv4 or IPv6 address to a Management Instance.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Before You Begin

- Ensure the IP address you plan to assign is not in use by an existing VLAN or CLIP IP subnet configured on the switch.
- If the DHCP client is configured for a Segmented Management Instance, you must manually disable the client. Configuring an IP address does not automatically disable the DHCP client.

Procedure

1. Enter Management Instance Configuration mode:

```
enable
```

```
configure terminal
```

```
mgmt {clip | oob | vlan}
```

2. Add an IPv4 address:

```
ip address {A.B.C.D [A.B.C.D] | A.B.C.D/X}
```

3. Add an IPv6 address:

```
ipv6 address WORD<0-255>
```

Example

Add an IPv4 address to the VLAN Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#ip address 192.0.2.12/24
```

Add an IPv4 address to the OOB Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt oob
Switch:1(mgmt:oob)#ip address 192.0.2.12 255.255.255.0
```

Add an IPv6 address to the CLIP Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt clip
Switch:1(mgmt:clip)#ipv6 address 2001:DB8::1/128
```

Configure a Segmented Management Instance Interface as Default Topology IP

Use this task to configure a Management Instance with a default topology IP. You can only configure one Management Instance interface as the default topology IP.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Procedure

1. Enter Management Instance Configuration mode:

```
enable

configure terminal

mgmt {clip | oob | vlan}
```

2. Configure the Management Instance as the default topology IP.

```
force-topology-ip
```

Example

The following example configures the Segmented Management Instance VLAN as the default topology IP:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#force-topology-ip
```

Configure Static Routes for a Management Instance

Use this task to configure static routes for Management Instances.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

About This Task

For the Management Instance CLIP, you do not need to configure a default or static route. This interface type uses all routing information learned by protocols attached to the VRF. For more information about how to associate a VRF with the CLIP interface, see [Create a Segmented Management Instance](#) on page 129.

For the Management Instance OOB and VLAN, you must configure a default or static route to reach the next-hop gateway; no routing protocol information is used to access off-link networks.

You can configure up to 100 IPv4 and IPv6 static routes.

Procedure

1. Enter Management Instance Configuration mode:

```
enable
```

```
configure terminal
```

```
mgmt {clip | oob | vlan}
```

2. Configure a static route:

```
ip route <A.B.C.D A.B.C.D | A.B.C.D/X> next-hop <A.B.C.D> [weight <1-65535>]
```

OR

```
ipv6 route WORD<0-255> [next-hop WORD<0-255>] [weight <1-65535>]
```

Example

Add a static route to configure routing for a Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#ip route 192.0.2.2/24 next-hop 198.51.100.1
```

Variable definitions

The following table defines parameters for the **ip route** and **ipv6 route** commands.

Variable	Value
<A.B.C.D A.B.C.D A.B.C.D/X>	Specifies the IP address and mask in one of the following formats: <ul style="list-style-type: none"> • A.B.C.D A.B.C.D • A.B.C.D/X
next-hop <A.B.C.D> OR next-hop WORD<0-255>	Specifies the next hop address for the static route. Use an IP in the same subnet as the management VLAN IP address.
weight <1-65535>	Specifies the static route cost. The default is 100 for CLIP, 200 for VLAN, and 300 for OOB. The management CLIP uses an internal static route with a weight of 100. If you use both CLIP and VLAN and need to force all default traffic out the management VLAN interface, configure a default static route with a weight lower than 100.
WORD<0-255>	Specifies the IPv6 address.

Configure Fragmented ICMP Packet Filtering on a Segmented Management Instance

About This Task

Use this task to enable fragmented ICMP packet filtering on a Segmented Management Instance.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Procedure

1. Enter Management Instance Configuration mode:

```
enable
```

```
configure terminal
```

```
mgmt {clip | oob | vlan}
```

2. Enable fragmented ICMP packet filtering:

- For IPv4:

```
ip icmp drop-fragments
```

- For IPv6:

```
ipv6 icmp drop-fragments
```

View Fragmented ICMP packet filtering Statistics on a Segmented Management Instance

About This Task

Use this task to view Fragmented ICMP packet filtering details on a Management Instance.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View Fragmented ICMP packet filtering:

- For IPv4:

```
show mgmt ip icmp [clip | oob | vlan]
```

- For IPv6:

```
show mgmt ipv6 icmp [clip | oob | vlan]
```

Variable Definitions

The following table defines parameters for the **show mgmt ip icmp** command.

Variable	Value
<i>clip</i>	Displays the IPv4 or IPv6 ICMP information specific to the management CLIP.
<i>oob</i>	Displays the IPv4 or IPv6 ICMP information specific to the management OOB.
<i>vlan</i>	Displays the IPv4 or IPv6 ICMP information specific to the management VLAN.

Configure MAC-offset for a Management VLAN Instance

Use this task to configure MAC-offset for a Management VLAN instance.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Procedure

1. Enter VLAN Management Instance Configuration mode:

```
enable
```

```
configure terminal
```

```
mgmt vlan
```

2. Enter the MAC-offset value:

```
mac-offset <0-511>
```

Variable Definitions

The following table defines parameters for the **mac-offset** command.

Variable	Value
<i><0-511></i>	Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated.

Change Attributes for the Management Instance

Use this task to change the IP address, VLAN, VRF, or default gateway for a Management Instance while you actively manage the switch over the same instance.



Important

Change the parameters in the following order:

1. VLAN or VRF
2. ports-tagged, ports-untagged-, or I-SID
3. IP address or default gateway

About This Task

Use this procedure if you onboard a new switch using default values but need the switch to be part of an existing Fabric deployment that uses non-default management VLAN attributes.

You cannot change parameters for more than one Management Instance operation at a time.

Procedure

1. Enter Management Instance Configuration mode:

```
enable
```

```
configure terminal
```

```
mgmt {clip | oob | vlan}
```

2. Use one of the following interfaces to configure the new values:

- Change parameters for an existing non-routed Management VLAN, or convert a routed Management VLAN to non-routing mode:
 - `convert vlan <1-4059> [rollback <0-3600>]`
 - `convert ip <A.B.C.D A.B.C.D | A.B.C.D/X> [rollback <0 | 30-3600>]`
 - `convert ipv6 WORD<0-255> [rollback <0 | 30-3600>]`
 - `convert gateway <A.B.C.D> [rollback <0 | 30-3600>]`
 - `convert gateway-v6 WORD<0-255> [rollback <0 | 30-3600>]`
 - `convert i-sid <1-16777215> [rollback <0 | 30-3600>]`
 - `convert vlan <1-4059> [ports-tagged {slot/port[/sub-port] [-slot/port[/sub-port]][, ...][slot/all][all]]} [ports-`


```

untagged {slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/
all][all]]} [rollback <0 | 30-3600>]
° convert vrf WORD <1-16> [rollback <0 | 30-3600>]

```



Important

After you configure the new values, the existing Management Instance VLAN is deleted and connectivity to the switch can be lost. You must reconnect to the switch before you can issue the `mgmt convert-commit` command.

- Change parameters for an existing routed Management VLAN, or convert a no-routing Management VLAN to routing mode:
 - ° `convert-to-routing [ip <A.B.C.D A.B.C.D | A.B.C.D/X>] [gateway {A.B.C.D}] [ipv6 <WORD<0-255>] [gateway-v6 WORD<0-255>] [rollback <0-3600>]`
 - ° `convert-to-routing [vlan <1-4059>] [vrf WORD<1-16>] [ports-untagged {slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]]} [ports-tagged {slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]]} [ip <A.B.C.D A.B.C.D | A.B.C.D/X>] [gateway {A.B.C.D}] [ipv6 WORD<0-255>] [gateway-v6 WORD<0-255>] [rollback <0 | 30-3600>]`
- Configure new Management OOB parameters:


```

convert [ip {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>}] [gateway <A.B.C.D>]
[rollback <0 | 30-3600>]

```
- Configure new Management CLIP parameters:


```

convert [vrf WORD <1-16>] [ip {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>}]
[rollback <0 | 30-3600>]

```



Note

If the VRF does not exist before you issue the **convert** command, the VRF is automatically created in the background. For this VRF to function properly, you must configure either SPBM Layer 3 VSN or IP interfaces and routing protocols.

3. Log on to Global Configuration mode to commit the parameter changes:

```
mgmt convert-commit
```



Note

Commit the change within 120 seconds (default) of issuing the **convert** or **convert-to-routing** command. Otherwise, the configuration changes automatically roll back to the previous configuration.

Examples

The following examples show change options for a Management Instance VLAN.

Convert a management VLAN to a new VLAN ID:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300
Switch:1(mgmt:vlan)#1 2021-09-24T21:41:54.627Z 4902 CP1 - 0x003c8677 -
00000000 GlobalRouter NLS_BASE INFO Mgmt convert: Dynamically moved the following ports
from vlan 10 to
new mgmt vlan 300 that had ARP entries: 1/1,
1 2021-09-24T21:41:54.627Z 4902 CP1 - 0x003c8671 -00000000 GlobalRouter NLS_BASE INFO
Mgmt convert: new
vlan 300 created successfully
1 2021-09-24T21:41:54.651Z 4902 CP1 - 0x003c8672 -00000000 GlobalRouter NLS_BASE INFO
Mgmt convert: existing
mgmt vlan instance deleted successfully
1 2021-09-24T21:41:54.719Z 4902 CP1 - 0x003c8673 -00000000 GlobalRouter NLS_BASE INFO
Mgmt convert: new mgmt
vlan instance created successfully with IP address 100.1.1.66/24
1 2021-09-24T21:41:54.719Z 4902 CP1 - 0x003c867e - 00000000 GlobalRouter NLS_BASE INFO
Convert on mgmt vlan
instance: Mgmt convert executed successfully

.
<reconnect to switch...>
Login: rwa
Password: ***

Mgmt convert: Please issue 'mgmt convert commit' before 120 seconds rollback timer
expires otherwise mgmt vlan
config change will be reverted
```

Convert a management VLAN to new IP address in the same subnet and in the same VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert ip 10.10.10.30/24
```

Convert a management VLAN to new IP address in a different subnet.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert ip 11.11.11.30/24 gateway 11.11.11.1
```

Convert a management VLAN to new VLAN ID with specified port or ports:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300 ports-untagged 1/2 ports-tagged 1/4
```

Convert a management VLAN to new VLAN ID with a specified I-SID:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300 i-sid 4300
```

Convert a management VLAN to new VLAN ID with a new IP address and default Gateway:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300 ip 11.11.11.30/24 gateway 11.11.11.1
```

Convert a management VLAN with all options:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300 ports-untagged 1/2 ports-tagged 1/4 i-sid 43000
ip 11.11.11.30/24 gateway 11.11.11.1
```

Convert a management VLAN to new I-SID (DvR Leaf only):

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert i-sid 4000
```

Convert a management VLAN with a faster rollback option (default is 120 seconds):

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300 rollback 60
```

Convert a management VLAN with no rollback option:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert vlan 300 rollback 0
```

Convert a management VLAN in non-routing mode to routing mode:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt vlan
Switch:1(mgmt:vlan)#convert-to-routing vlan 10 ip 192.168.10.1/24
```

The following examples show change options for Management Instance CLIP.

Convert a management CLIP from one VRF to another VRF. The IP address is the same:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt clip
Switch:1(mgmt:clip)#convert vrf blue
```

```
WARNING: the specified vrf does not exist - connectivity to the mgmt clip will be lost
until l3vsrn or
ip interfaces for the given vrf are provisioned.
Continue with this operation (y/n) ? n
```



Note

If the VRF does not exist before you issue the **convert** command, the VRF is automatically created in the background. For this VRF to function properly, you must configure either SPBM Layer 3 VSN or IP interfaces and routing protocols.

Convert a management CLIP to a new IP address. The VRF is the same:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt clip
Switch:1(mgmt:vlan)#convert ip 30.30.30.100/32
```

Convert a management CLIP to new IP address and VRF:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt clip
Switch:1(mgmt:vlan)#convert vrf blue ip 30.30.30.100/32
```

The following examples show the change options attributes for Management Instance Out-Of-Band (OOB).

Convert a management OOB to a new IP address in the same subnet:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt oob
Switch:1(mgmt:vlan)#convert ip 20.20.20.100/24
```

Convert a management OOB IP address to a different subnet:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#mgmt oob
Switch:1(mgmt:vlan)#convert ip 21.21.21.100/24 gateway 21.21.21.1
```

Variable Definitions

The following table defines parameters for the **convert** command.

Variable	Value
<code>gateway <A.B.C.D></code>	<p>Specifies the gateway IPv4 address.</p> <p>This parameter applies to the following Management Instance interface types:</p> <ul style="list-style-type: none"> • OOB • VLAN <p>If you convert a Management VLAN in routing mode to non-routing mode, you must specify the gateway as part of the convert command syntax to install a default route in the new Management VLAN interface.</p>
<code>gateway-v6 WORD<0-255></code>	<p>Specifies the gateway IPv6 address for the VLAN. This parameter applies to the VLAN interface type.</p> <p>If you convert a Management VLAN in routing mode to non-routing mode, you must specify the gateway as part of the convert command syntax to install a default route in the new Management VLAN interface.</p>
<code>ip <A.B.C.D A.B.C.D A.B.C.D/X></code>	Specifies the IPv4 address and subnet mask.
<code>ipv6 WORD<0-255></code>	Specifies the IPv6 address for the VLAN interface type.
<code>i-sid <1-16777215></code>	<p>Specifies the service instance identifier (I-SID).</p> <p>This parameter applies only to the VLAN interface type, in non-routing mode.</p> <p>Note: You can specify the I-SID for a switch configured as a DvR Leaf node.</p>
<code>rollback <0 30-3600></code>	<p>Specifies the time in seconds between when the command is issued and when the command changes automatically roll back to the previous configuration.</p> <p>The default is 120 seconds. To disable the rollback, enter 0.</p>

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>vlan <2-4059></code>	<p>Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.</p>
<code>vrf WORD<0-16></code>	<p>Specifies the vrf name.</p> <p>This parameter applies to the Management Instance CLIP and VLAN interface types only.</p>

The following table defines parameters for the **convert-to-routing** command.

Variable	Value
<code>gateway <A.B.C.D></code>	<p>Specifies the gateway IPv4 address.</p> <p>If you change the IP address for the Management VLAN to a different subnet, the switch only copies the static routes if you specify the gateway in the command syntax.</p>
<code>gateway-v6 WORD<0-255></code>	<p>Specifies the gateway IPv6 address for the VLAN.</p> <p>If you change the IP address for the Management VLAN to a different subnet, the switch only copies the static routes if you specify the gateway in the command syntax.</p>

Variable	Value
<code>ip <A.B.C.D A.B.C.D A.B.C.D/X></code>	Specifies the IP address and mask for the Management VLAN in one of the following formats: <ul style="list-style-type: none"> A.B.C.D A.B.C.D A.B.C.D/X If you do not specify an IP address, the existing VLAN IP address is maintained.
<code>ipv6 WORD<0-255></code>	Specifies the IPv6 address for the Management VLAN.
<code>ports-tagged {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code>	Specifies the tagged port list.
<code>ports-untagged {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code>	Specifies the untagged port list.
<code>rollback <0 30-3600></code>	Specifies the time in seconds between when the command is issued and when the changes will be undone if you do not commit the change. The default is 120 seconds. To disable the rollback, enter 0.
<code>vlan <1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD<1-16></code>	Specifies the VRF name to associate with the VLAN.

View Segmented Management Instance Information

Use this task to view Management Instance information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View general configuration information:

```
show mgmt interface [clip | oob | vlan]
```

3. View operational routes for the Management Instance:

```
show mgmt ip route [clip | oob | vlan]
```

OR

```
show mgmt ipv6 route [clip | oob | vlan]
```

**Note**

Routes with a type of LOCAL have a metric equal to 1.

4. View configured static routes for the Management Instance:

```
show mgmt ip route static [vlan | oob | clip]
```

OR

```
show mgmt ipv6 route static [vlan | oob | clip]
```

**Note**

Routes with a type of LOCAL have a metric equal to 256.

5. View the ARP or Neighbor Discovery cache information for the Management Instance:

```
show mgmt ip arp [clip | oob | vlan]
```

OR

```
show mgmt ipv6 neighbor [clip | oob | vlan]
```

Examples**Note**

For information on the Mgmt-sdwan interface, see [Fabric Extend \(FE\) States](#) on page 32.

Switch:1>show mgmt interface

Mgmt Interface Information								
INST	DESCR	TYPE	ADMIN	VLAN	ISID	PORT	VRF	PHYSICAL
1	Mgmt-oob1	Oob	enable	-	-	oob	-	40:e3:17:cc:d4:81
3	Mgmt-clip	CLIP	disable	-	-	-	GlobalRouter	00:00:00:00:00:00
4	Mgmt-vlan	VLAN-L2	enable	10	-	-	-	40:e3:17:cc:d5:00
5	Mgmt-sdwan	VLAN	enable	4047	-	-	-	40:e3:17:cc:d5:01

4 out of 4 Total Num of mgmt interfaces displayed

Switch:1>show mgmt interface vlan

Mgmt Interface Information								
INST	DESCR	TYPE	ADMIN	VLAN	ISID	PORT	VRF	
PHYSICAL								
4	Mgmt-vlan	VLAN-L2	enable	2	-	-	-	192.0.2.188

1 out of 1 Total Num of mgmt interfaces displayed

Switch:1>show mgmt ip route

Mgmt IPv4 Route Information - Table main				
DEST/MASK	NEXTHOP	METRIC	INTERFACE	TYPE
0.0.0.0/0	0.0.0.0	100	Mgmt-clip	INTERNAL
0.0.0.0/0	0.0.0.0	300	Mgmt-oob1	DHCP
192.0.2.189/24	0.0.0.0	256	Mgmt-vlan	LOCAL
172.16.0.0/29	0.0.0.0	1	Mgmt-sdwan	LOCAL

4 out of 4 Total Num of mgmt ip route displayed

Switch:1>show mgmt ip route static

Mgmt IPv4 Static Route Information - Table main					
INTERFACE	DEST/MASK	NEXTHOP	METRIC	STATE	TYPE
Mgmt-vlan	192.0.2.1/24	10.0.0.30	200	ACTIVE	STATIC
Mgmt-vlan	198.51.100.5/24	10.0.0.40	200	ACTIVE	STATIC
Mgmt-oob1	0.0.0.0/0	192.0.2.5	300	ACTIVE	DHCP

Switch:1>show mgmt ipv6 route static

Mgmt IPv6 Static Route Information - Table main				
INTERFACE	DEST/MASK	NEXTHOP	METRIC	STATE
Mgmt-vlan	40:0:0:0:0:0:0:0/64	10:0:0:0:0:0:0:40	200	ACTIVE
Mgmt-vlan	50:0:0:0:0:0:0:0/64	10:0:0:0:0:0:0:50	200	ACTIVE

Switch:1>show mgmt ip arp

Mgmt IP ARP Information			
IP_ADDRESS	INTERFACE	MAC_ADDRESS	STATE
10.10.10.1	Mgmt-vlan	00:1d:af:64:a2:14	REACHABLE
10.10.10.22	Mgmt-vlan	00:18:b0:5a:92:14	STALE
10.10.10.33	Mgmt-vlan	00:50:56:8c:43:55	FAILED
172.16.0.1	Mgmt-sdwan	00:07:32:a6:8e:d4	STALE

4 out of 4 Total Num of mgmt ip arp displayed

Switch:1>show mgmt ipv6 neighbor

Mgmt IPv6 Neighbor Information			
IPV6_ADDRESS	INTERFACE	MAC_ADDRESS	STATE
10::1	Mgmt-vlan	00:1d:af:64:a2:14	REACHABLE
10::22	Mgmt-vlan	00:18:b0:5a:92:14	STALE
10::33	Mgmt-vlan	00:50:56:8c:43:53	FAILED

View IP Address Information for a Segmented Management Instance

Use this task to view the IP address information for a Management Instance.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View Segmented Management Instance IP address information:

```
show mgmt ip [<clip | oob | vlan>]
```

OR

```
show mgmt ipv6 [<clip | oob | vlan>]
```

3. View Segmented Management Instance topology IP address information:

```
show mgmt topology-ip
```

Example

```
Switch:1>show mgmt ip vlan
```

```
=====
Mgmt IP Information
=====
INST      DESCR      IPV4      IPV6 GLOBAL/PREFIX LENGTH IPV6 LINKLOCAL
-----
4          Mgmt-vlan  192.0.2.12/24      0:0:0:0:0:0:0:0/0      0:0:0:0:0:0:0:0
1 out of 1 Total Num of mgmt ip displayed
=====
```

```
Switch:1>show mgmt topology-ip
```

```
=====
Mgmt Topology IP Information
=====
IPv4:
    Address: 192.0.2.10
    Instance: 1
    Description: oob1

IPv6:
    No address to display

Force-topology-ip setting: none
```

View Segmented Management Instance Statistics

Use this task to view Management Instance statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View general Segmented Management Instance statistics:

```
show mgmt statistics [clip | oob | vlan]
```

3. View ICMP statistics for the Management Instance:

```
show mgmt ip icmp-statistics
```

OR

```
show mgmt ipv6 icmp-statistics
```

4. View IP statistics for the Management Instance:

```
show mgmt ip ip-statistics
```

OR

```
show mgmt ipv6 ip-statistics
```

5. View the TCP connections for the Management Instance:

```
show mgmt ip tcp-connections
```

OR

```
show mgmt ipv6 tcp-connections
```

6. View the TCP statistics for the Management Instance:

```
show mgmt ip tcp-statistics
```

OR

```
show mgmt ipv6 tcp-statistics
```

7. View the UDP endpoints for the Management Instance:

```
show mgmt ip udp-endpoints
```

OR

```
show mgmt ipv6 udp-endpoints
```

8. View the UDP statistics for the Management Instance:

```
show mgmt ip udp-statistics
```

OR

```
show mgmt ipv6 udp-statistics
```

9. Enter Privileged EXEC mode:

```
enable
```

10. (Optional) Clear all of the statistics for the Management Instance:

```
clear mgmt statistics
```

Examples**Note**

For information on the `Mgmt-sdwan` interface, see [Fabric Extend \(FE\) States](#) on page 32.

```
Switch:1>show mgmt statistics
```

Mgmt Interface Stats Information							
INST	DESCR	RX-PKTS	RX-ERROR	RX-DROP	TX-PKTS	TX-ERROR	TX-DROP
1	Mgmt-oob1	111667	0	0	21412	0	0
3	Mgmt-clip	0	0	0	0	0	0
4	Mgmt-vlan	28	0	0	7	0	0
5	Mgmt-sdwan	43337	0	0	99	0	0

4 out of 4 Total Num of mgmt interfaces displayed

```
Switch:1>show mgmt ip icmp-statistics
```

```
=====
Mgmt ICMP Statistics Information
=====
```

```
InMsgs      : 44
InErrors    : 0
InCsumErrors : 0
InDestUnreaches : 44
InTimeExcds : 0
InParmProbs : 0
InSrcQuenchs : 0
InRedirects : 0
InEchos     : 0
InEchoReps  : 0
InTimestamps : 0
InTimestampReps : 0
InAddrMask  : 0
InAddrMaskReps : 0
OutMsgs     : 53
OutErrors   : 0
```

```
Switch:1>show mgmt ipv6 icmp-statistics
```

```
=====
Mgmt ICMPv6 Statistics Information
=====
```

```
InMsgs      : 58
InErrors    : 0
InCsumErrors : 0
InDestUnreaches : 0
InTimeExcds : 0
InParmProbs : 0
InPktTooBigs : 0
InRedirects : 0
InEchos     : 0
InEchoReps  : 0
InGroupMembQueries : 0
InGroupMembReductions : 0
InRouterSolicits : 0
InRouterAdvertisements : 55
InNeighborSolicits : 0
InNeighborAdvertisements : 3
InMLDv2Reports : 0
InType134    : 55
InType136    : 3
OutMsgs     : 69
OutErrors   : 0
OutDestUnreaches : 0
OutTimeExcds : 0
OutParmProbs : 0
OutPktTooBigs : 0
OutRedirects : 0
OutEchos     : 0
OutEchoReps  : 0
OutGroupMembQueries : 0
OutGroupMembResponses : 0
OutGroupMembReductions : 0
OutRouterSolicits : 0
OutRouterAdvertisements : 0
OutNeighborSolicits : 13
OutNeighborAdvertisements : 0
```

```

OutMLDv2Reports      : 56
OutType133           : 56
OutType135           : 13
OutType143           : 0

```

```

-----
Switch:1>show mgmt ip ip-statistics

```

```

=====
                        Mgmt IP Statistics Information
=====
InReceives           : 1231729
InHdrErrors           : 0
InAddrErrors          : 489
InUnknownProtos       : 0
InDiscards            : 0
InDelivers            : 1221886
OutRequests           : 1212585
OutDiscards           : 20
OutNoRoutes            : 0
ForwDatagrams          : 0
ReasmTimeout          : 0
ReasmReqds            : 0
ReasmOKs              : 0
ReasmFails            : 0
FragOKs               : 0
FragFails             : 0
FragCreates           : 0

```

```

-----
Switch:1>show mgmt ipv6 ip-statistics

```

```

=====
                        Mgmt IPv6 Statistics Information
=====
InReceives           : 226
InHdrErrors           : 0
InAddrErrors          : 0
InUnknownProtos       : 0
InDiscards            : 0
InDelivers            : 62
InTooBigErrors        : 0
InNoRoutes            : 0
InTruncatedPkts       : 0
InMcastPkts           : 224
InOctets              : 20556
InMcastOctets         : 20416
InBcastOctets         : 0
InNoECTPkts           : 226
InECT1Pkts            : 0
InECT0Pkts            : 0
InCEPkts              : 0
OutRequests           : 71
OutDiscards           : 0
OutNoRoutes            : 0
OutForwDatagrams       : 0
OutMcastPkts          : 69
OutOctets              : 5412
OutMcastOctets        : 5272
OutBcastOctets        : 0
ReasmTimeout          : 0
ReasmReqds            : 0
ReasmOKs              : 0
ReasmFails            : 0

```

```

FragOKs      : 0
FragFails    : 0
FragCreates  : 0

```

```

Switch:1>show mgmt ip tcp-connections

```

```

=====
Mgmt IP TCP connections
=====
STATE          RECV-Q    SEND-Q    Local Address:Port    Peer Address:Port
-----
LISTEN         0         5         0.0.0.0:ftp           0.0.0.0:*
LISTEN         0         5         0.0.0.0:telnet        0.0.0.0:*
LISTEN         0         40        0.0.0.0:https         0.0.0.0:*
LISTEN         0         1         0.0.0.0:login         0.0.0.0:*
ESTAB          0         0         192.0.2.10:https      198.51.100.1:50694
ESTAB          0         3         192.0.2.10:telnet     198.51.100.1:58862
ESTAB          0         0         192.0.2.10:https      198.51.100.1:59774
=====

```

```

Switch:1>show mgmt ipv6 tcp-connections

```

```

=====
Mgmt IPv6 TCP connections
=====
STATE          RECV-Q    SEND-Q    Local Address:Port    Peer Address:Port
-----
LISTEN         0         5         *:ftp                 *: *
LISTEN         0         5         *:telnet              *: *
LISTEN         0         40        *:https               *: *
LISTEN         0         1         *:login               *: *
=====

```

```

Switch:1>show mgmt ip tcp-statistics

```

```

=====
Mgmt Combined IPv4/v6 TCP Statistics Information
=====
TcpActiveOpens   : 9571
TcpPassiveOpens  : 9658
TcpAttemptFails  : 17
TcpEstabResets   : 86
TcpInSegs        : 1207867
TcpOutSegs       : 1199088
TcpRetransSegs   : 42
TcpInErrs        : 0
TcpOutRsts       : 89
TcpInCsumErrors  : 0
=====

```

```

Switch:1>show mgmt ipv6 tcp-statistics

```

```

=====
Mgmt Combined IPv4/v6 TCP Statistics Information
=====
TcpActiveOpens   : 9626
TcpPassiveOpens  : 9713
TcpAttemptFails  : 17
TcpEstabResets   : 86
TcpInSegs        : 1212159
TcpOutSegs       : 1203293
TcpRetransSegs   : 42
TcpInErrs        : 0
=====

```

```
TcpOutRsts      : 89
```

```
Switch:1>show mgmt ip udp-endpoints
```

```
=====
```

Mgmt IP UDP endpoints				
STATE	RECV-Q	SEND-Q	Local Address:Port	Peer Address:Port
UNCONN	0	0	0.0.0.0:bootpc	0.0.0.0:*
UNCONN	0	0	0.0.0.0:tftp	0.0.0.0:*
UNCONN	0	0	192.0.2.10:ntp	0.0.0.0:*
UNCONN	0	0	0.0.0.0:ntp	0.0.0.0:*
UNCONN	0	0	0.0.0.0:snmp	0.0.0.0:*

```
=====
```

```
Switch:1>show mgmt ipv6 udp-endpoints
```

```
=====
```

Mgmt IPv6 UDP endpoints				
STATE	RECV-Q	SEND-Q	Local Address:Port	Peer Address:Port
UNCONN	0	0	[0:0:0:0:0:0:1]:domain	*:*
UNCONN	0	0	*:tftp	*:*
UNCONN	0	0	[fe80:0:0:0:f66e:95ff:fe9f:81]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:f66e:95ff:fe9f:a5]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:a8bb:ccff:fedd:ee01]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:fc7:79ff:fe04:999c]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:609a:4aff:fe4e:cf04]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:c482:2aff:fe75:2e66]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:c80a:73ff:fe00:364e]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:80cc:e9ff:fec7:b9e2]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:64da:41ff:fec5:489e]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:f8d7:d6ff:feda:62bc]:ntp	*:*
UNCONN	0	0	[fe80:0:0:0:f66e:95ff:fe9f:0]:ntp	*:*
UNCONN	0	0	[0:0:0:0:0:0:1]:ntp	*:*
UNCONN	0	0	*:ntp	*:*
UNCONN	0	0	*:snmp	*:*

```
=====
```

```
Switch:1>show mgmt ip udp-statistics
```

```
=====
```

Mgmt UDP Statistics Information	
UdpInDatagrams	: 63622
UdpNoPorts	: 44
UdpInErrors	: 0
UdpOutDatagrams	: 63666
UdpIgnoredMulti	: 0
UdpRcvbufErrors	: 0
UdpSndbufErrors	: 0
UdpInCsumErrors	: 0

```
=====
```

```
Switch:1>show mgmt ipv6 udp-statistics
```

```
=====
```

Mgmt UDP6 Statistics Information	
Udp6InDatagrams	: 0
Udp6NoPorts	: 0
Udp6InErrors	: 0
Udp6OutDatagrams	: 0
Udp6IgnoredMulti	: 0
Udp6RcvbufErrors	: 0
Udp6SndbufErrors	: 0
Udp6InCsumErrors	: 0

```
=====
```

Redistribution of CLIP Segmented Management Instance Examples

The CLIP Management Instance is added as a LOCAL route in the Control Processor Route Table Manager table and change list infrastructure. Existing route redistribution mechanisms redistribute local routes into the desired routing protocols within the associated VRF or across VRF instances.

Redistribute IPv4 CLIP Management Instance to OSPF in GRT

```
router ospf
 redistribute direct
 redistribute direct enable
 exit
 ip ospf apply redistribute direct
```

Redistribute IPv6 CLIP Management Instance to OSPF in GRT

```
router ospf
 ipv6 redistribute direct
 ipv6 redistribute direct enable
 ipv6 ospf apply redistribute direct
```

Redistribute IPv4 CLIP Management Instance to OSPF in VRF

```
router vrf red
 ip ospf redistribute direct
 ip ospf redistribute direct enable
 exit
 ip ospf apply redistribute direct vrf red
```

Redistribute IPv6 CLIP Management Instance to OSPF in VRF

```
router vrf red
 ipv6 ospf redistribute direct
 ipv6 ospf redistribute direct enable
 exit
 ipv6 ospf apply redistribute direct vrf red
```

Redistribute IPv4 CLIP Management Instance to BGP in GRT

```
router bgp
 redistribute direct
 redistribute direct enable
 ip bgp apply redistribute direct
```

Redistribute IPv6 CLIP Management Instance to BGP in GRT

```
router vrf
 redistribute ipv6-direct
 redistribute ipv6-direct enable
 ipv6 bgp apply redistribute direct
```

Redistribute IPv4 CLIP Management Instance to BGP in VRF

```
router vrf red
 ip bgp redistribute direct
 ip bgp redistribute direct enable
 exit
 ip bgp apply redistribute direct vrf red
```


Redistribute IPv6 CLIP Management Instance to BGP in VRF

```
router vrf red
ip bgp redistribute ipv6-direct
ip bgp redistribute ipv6-direct enable
exit
ipv6 bgp apply redistribute direct vrf red
```

Accept Policy for IPv4 CLIP Management Instance in GRT to VRF Red (I-SID 200)

```
#grt-->vrf
router vrf red
isis accept i-sid 0
isis accept i-sid 0 enable
isis redistribute direct
isis redistribute direct enable
exit
isis apply accept vrf red
isis apply redistribute direct vrf red

#vrf-->grt
router isis
accept i-sid 200
accept i-sid 200 enable
exit
isis apply accept
```

Accept Policy for IPv6 CLIP Management Instance in GRT to VRF Red (I-SID 200)

```
#grt-->vrf
router vrf red
ipv6 isis accept i-sid 0
ipv6 isis accept i-sid 0 enable
ipv6 isis redistribute direct
ipv6 isis redistribute direct enable
exit
ipv6 isis apply accept vrf red
ipv6 isis apply redistribute direct vrf red

#vrf-->grt
router isis
ipv6 accept i-sid 200
ipv6 accept i-sid 200 enable
exit
ipv6 isis apply accept
```

Accept Policy for IPv4 CLIP Management Instance in VRF Blue (I-SID 300) to GRT

```
#vrf --> grt
router isis
accept i-sid 300
accept i-sid 300 enable
redistribute direct
redistribute direct enable
exit
isis apply accept
isis apply redistribute direct

#grt --> vrf
router vrf blue
isis accept i-sid 0
isis accept i-sid 0 enable
```

```
exit
isis apply accept vrf blue
```

Accept Policy for IPv6 CLIP Management Instance in VRF Blue (I-SID 300) to GRT

```
#vrf --> grt
router isis
ipv6 accept i-sid 300
ipv6 accept i-sid 300 enable
ipv6 redistribute direct
ipv6 redistribute direct enable
exit
ipv6 isis apply accept
ipv6 isis apply redistribute direct

#grt --> vrf
router vrf blue
ipv6 isis accept i-sid 0
ipv6 isis accept i-sid 0 enable
exit
ipv6 isis apply accept vrf blue
```

Segmented Management Instance Configuration using EDM

This section provides procedures to configure Segmented Management Instance using EDM.

Migrate an IP Address to a Segmented Management Instance

Perform this procedure to migrate a new routing VLAN with a new IP address or a new loopback IP address under a different VRF to the Segmented Management Instance. Alternatively, you can also use the **convert** command.



Important

Choose a VLAN that does not have an IP interface on it. The upgrade process removes the IP configuration and network connectivity will be impacted.

About This Task

You cannot migrate interfaces used for routing purposes, for example, where you configure Layer 3 routing protocols.

This command does not apply to the OOB or mgmtEthernet interface. Releases that support this migration procedure automatically move the IP address on the mgmtEthernet interface from the routing stack to the Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Migrate** tab.
4. Select **Insert**.
5. Select the instance type, either **clip** or **vlan**.

6. Specify the existing VLAN or loopback ID.
7. Select **Insert**.

Migrate Field Descriptions

Use the data in the following table to use the **Migrate** tab.

Name	Description
InstanceId	Specifies the interface instance to migrate.
InterfaceIndex	Shows the interface index of the identified interface.
InterfaceType	Shows the interface type.
Description	Shows the interface description.
VlanId	Specifies the VLAN ID for a port-based VLAN.
LoopbackId	Specifies the loopback ID.
VrfName	Shows the VRF associated with the loopback interface.
IpAddress	Shows the IPv4 address to migrate.
IpMask	Shows the subnet mask for the IPv4 address.
Ipv6Address	Shows the IPv6 address to migrate.
Ipv6PrefixLength	Shows the prefix length for the IPv6 address.

Configure a Segmented Management Instance

You must create a Management Instance to gain access to specific management applications.

About This Task

The Management Instance supports different management interface types. When you create the Management Instance, you specify the interface type and the switch automatically creates the appropriate instance ID for that type.

Each Management Instance supports a IPv4 and IPv6 (global scope) management address for use by management applications.

This procedure creates a management CLIP, management OOB, or a management VLAN in non-routing mode. To create a management VLAN in routing mode, see [Create a Management VLAN In Routing Mode](#) on page 159.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.

3. Select the **Interface** tab.
4. Select **Insert**.
5. In **Instanceld**, select the type of Management Instance to create.
6. (Optional) For a CLIP Management Instance, in **VrfName**, type the VRF name to associate with the CLIP instance.

**Note**

To associate the GRT with the CLIP instance, type **GlobalRouter** in **VrfName**.

7. For a VLAN Management Instance, in **VlanId**, type the VLAN ID to associate the management VLAN with an existing port-based VLAN.
8. For an OOB Management Instance, in **OOBIndex**, select the interface port number to associate for Out-of-Band management.
9. Select **State** to enable the instance.
10. To specify the interface as the default topology IP for LLDP advertisements, select **InterfaceTopologyIpFlag**.
11. To administratively enable RMON for the interface, select **RmonAdminEnable**.
12. For a DvR Leaf node, in **Isid**, type the I-SID value to associate with the Management Instance VLAN.
13. Select **Insert**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Instanceld	Indicates the Management Instance type associated with this entry.
InterfaceType	Indicates the interface type.
VlanId	Specifies the VLAN ID to associate with the management VLAN. This action creates a Management VLAN in non-routing mode.
OOBIndex	Specifies the interface ID to associate with the management OOB.
VrfName	Specifies the VRF name to associate with the management CLIP.
State	Indicates if the interface is enabled for this instance. The default is disabled.
InterfaceMacAddr	Indicates the MAC address for the interface.
InterfaceName	Indicates the interface name.
InterfaceTopologyIpFlag	Specifies if the interface is the default topology source IP.
ZtpOn	Identifies the Zero Touch Provisioning status for the interface.
RmonAdminEnable	Specifies if RMON is administrative enabled for the interface.
RmonOperEnable	Indicates the RMON operational status for the interface.

Name	Description
RmonIpAddress	Indicates the RMON IP address for the interface.
MacOffset	Translates the IP address into a MAC address.
DropIcmpFragEnable	Enables IPv4 Fragmented ICMP packet filtering on the Management Instance. The default is disabled.
DropIcmpv6FragEnable	Enables IPv6 Fragmented ICMP packet filtering on the Management Instance. The default is disabled.
Isid	Specifies the I-SID number that associated with the Management Instance VLAN for the DvR Leaf node. For non DvR node, the default I-SID value is 0 and this value cannot be edited.

Create a Management VLAN In Routing Mode

Before You Begin

Configure an IP address on the VLAN before you enable the management VLAN interface.

About This Task

You must create a Management Instance to gain access to specific management applications. Perform this procedure to create a Management VLAN in routing mode. To create a Management VLAN in non-routing mode, see [Configure a Segmented Management Instance](#) on page 157 .

This Management VLAN inherits the VRF, MAC offset, and IPv4 and IPv6 address configuration from the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Vlan Router** tab.
4. Select **Insert**.
5. Type the VLAN ID or select **Vlan** to choose from a list of available VLAN IDs, and then select **Ok**.
6. Select **Insert**.
7. Select **Apply**.

Vlan Router *Field Descriptions*

Use the data in the following table to use the **Vlan Router** tab.

Name	Description
VlanId	Specifies the ID of the VLAN interface to use as the Management VLAN.
VrfName	Displays the name of the VRF associated with the VLAN.

Configure DHCP Client for a Segmented Management Instance

Use this task to configure a DHCP Client on a management interface.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Dhcp** tab.
4. In **Client**, select an option to configure the DHCP Client.
5. Select **Apply**.

Dhcp Field Descriptions

Use the data in the following table to use the **Dhcp** tab.

Name	Description
Client Note: Exception: oob and cycle are only supported on hardware with a physical management port	Specifies the DHCP client configuration: <ul style="list-style-type: none">• oob - DHCP Client on the Out-of-Band management interface.• vlan - DHCP Client on the VLAN management interface.• cycle - DHCP Client cycles between in-band and Out-of-Band management interfaces until an IP address is obtained on one management interface.• disable - DHCP Client is disabled.
ClientPreferredInterface	Shows the DHCP Client preferred management interface when in cycle mode. On reboot, the system first attempts to acquire a DHCP IP address on the preferred interface.

View IPv4 ARP Information for a Segmented Management Instance

Use this task to view IPv4 Address Resolution Protocol (ARP) information.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **IpArp** tab.

IpArp Field Descriptions

Use the data in the following table to use the **IpArp** tab.

Name	Description
Address	Shows the IPv4 address of the ARP entry.
Instance	Shows the Management Instance ID.
IntfName	Shows the Management Instance interface name for the ARP entry.
MacAddr	Shows the MAC address for the ARP entry.
State	Shows the state of the ARP entry. The state can be one of the following: <ul style="list-style-type: none">• reachable• stale• permanent• failed• delay

View IPv6 ND Information for a Segmented Management Instance

Use this task to view IPv6 Neighbor Discovery (ND) information.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Ipv6Neighbor** tab.

Ipv6Neighbor Field Descriptions

Use the data in the following table to use the **Ipv6Neighbor** tab.

Name	Description
Addr	Shows the IPv6 address of the neighbor entry.
Instance	Shows the Management Instance ID.
IntfName	Shows the Management Instance interface name for the neighbor entry.
MacAddr	Shows the MAC address for the neighbor entry.
State	Shows the state of the neighbor entry. The state can be one of the following: <ul style="list-style-type: none">• reachable• stale• permanent• failed• delay

Configure IPv4 Static Routes for a Management Instance

Use this task to configure IPv4 static routes for Management Instances.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

About This Task

For the Management Instance CLIP, you do not need to configure a default or static route. This interface type uses all routing information learned by protocols attached to the VRF. For more information about how to associate a VRF with the CLIP interface, see [Configure a Segmented Management Instance](#) on page 157.

For the Management Instance OOB and VLAN, you must configure a default or static route to reach the next-hop gateway; no routing protocol information is used to access off-link networks.

You can configure up to 100 IPv4 and IPv6 static routes.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **IpStaticRoute** tab.
4. Select **Insert**.
5. For the **Instance**, select the type of Management Instance interface.
6. Type the destination IP address and mask.
7. Type the next hop IP address.
8. (Optional) Type a metric value.
9. Select **Insert**.

IpStaticRoute Field Descriptions

Use the data in the following table to use the **IpStaticRoute** tab.

Name	Description
Instance	Specifies the Management Instance.
DestAddr	Specifies the destination IP address.
DestMask	Specifies the destination mask.
NextHop	Specifies the next hop address for the static route. Use an IP address in the same subnet as the management VLAN IP address.
IntfName	Specifies the Management Instance interface name for the route entry.

Name	Description
Metric	Specifies the static route cost. The default is 200. The management CLIP uses an internal static route with a weight of 100. If you use both CLIP and VLAN and need to force all default traffic out the management VLAN interface, configure a default static route with a weight lower than 100.
State	Shows if the route is active or inactive.

Configure IPv6 Static Routes for a Management Instance

Use this task to configure IPv6 static routes for Management Instances.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

About This Task

For the Management Instance CLIP, you do not need to configure a default or static route. This interface type uses all routing information learned by protocols attached to the VRF. For more information about how to associate a VRF with the CLIP interface, see [Configure a Segmented Management Instance](#) on page 157.

For the Management Instance OOB and VLAN, you must configure a default or static route to reach the next-hop gateway; no routing protocol information is used to access off-link networks.

You can configure up to 100 IPv4 and IPv6 static routes.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Ipv6StaticRoute** tab.
4. Select **Insert**.
5. For the **Instance**, select the type of Management Instance interface.
6. Type the destination IPv6 address and prefix length.
7. Type the next hop IPv6 address.
8. (Optional) Type a metric value.
9. Select **Insert**.

Ipv6StaticRoute Field Descriptions

Use the data in the following table to use the **Ipv6StaticRoute** tab.

Name	Description
Instance	Specifies the Management Instance.
DestAddr	Specifies the destination IP address.
DestPrefixLen	Specifies the destination prefix length.
NextHop	Specifies the next hop address for the static route. Use an IP address in the same subnet as the management VLAN IP address.
IntfName	Specifies the Management Instance interface name for the route entry.
Metric	Specifies the static route cost. The default is 200. The management CLIP uses an internal static route with a weight of 100. If you use both CLIP and VLAN and need to force all default traffic out the management VLAN interface, configure a default static route with a weight lower than 100.
State	Shows if the route is active or inactive.

View IPv4 Operational Routes for a Segmented Management Instance

Use this task to view IPv4 operational routes.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **IpRoute** tab.

IpRoute Field Descriptions

Use the data in the following table to use the **IpRoute** tab.

Name	Description
DestAddr	Shows the destination address of the route entry.
DestMask	Shows the destination mask of the route entry.
Metric	Shows the metric, or cost, assigned to the route entry. If multiple entries exist to the same destination, the metric determines which route is used. Routes with a type of LOCAL have a metric equal to 1.
Instance	Shows the Management Instance ID.
NextHop	Shows the next hop for the route entry.

Name	Description
IntfName	Shows the Management Instance interface name for the route entry.
Type	Shows the type of route entry.

View IPv6 Operational Routes for a Segmented Management Instance

Use this task to view IPv6 operational routes.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Ipv6Route** tab.

Ipv6Route Field Descriptions

Use the data in the following table to use the **Ipv6Route** tab.

Name	Description
DestAddr	Shows the destination address of the route entry.
PrefixLen	Shows the destination prefix length of the route entry.
Metric	Shows the metric, or cost, assigned to the route entry. If multiple entries exist to the same destination, the metric determines which route is used. Routes with a type of LOCAL have a metric equal to 256.
Instance	Shows the Management Instance ID.
NextHop	Shows the next hop for the route entry.
IntfName	Shows the Management Instance interface name for the route entry.
Type	Shows the type of route entry.

View Topology IP for a Segmented Management Instance

Use this task to view the default topology IP address for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **TopologyIp** tab.

TopologyIp Field Descriptions

Use the data in the following table to use the **TopologyIp** tab.

Name	Description
AddrType	Shows the IP address type for the topology IP.
Addr	Shows the IP address for the topology IP.
InterfaceName	Shows the interface name of the identified interface for the topology IP.
InstanceId	Specifies the interface instance for the topology IP.

Configure an IP Address for a Segmented Management Instance

Use this task to configure and view IPv4 address information for a Segmented Management Instance.

**Note**

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Before You Begin

- Ensure the IP address you plan to assign is not in use by an existing VLAN or CLIP IP subnet configured on the switch.
- If the DHCP client is configured for a Segmented Management Instance, you must manually disable the client. Configuring an IP address does not automatically disable the DHCP client.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **IpAddress** tab.
4. Select **Insert**.
5. Select the Segmented Management Instance interface type.
6. Type the address information.
7. Select **Insert**.

IpAddress Field Descriptions

Use the data in the following table to use the **IpAddress** tab.

Name	Description
Instanceld	Specifies the interface instance.
Address	Specifies IPv4 address for the interface instance. Ensure that the management CLIP IP address does not fall into the range of a configured VLAN IP address range as this is not allowed.
Mask	Specifies the subnet mask of the IP address.
AddrOrigin	Shows the IP address origin.
IntfName	Shows the interface name.

Configure an IPv6 Address for a Segmented Management Instance

Use this task to configure or view IPv6 address information for a Segmented Management Instance.



Note

You can use this task to configure a Segmented Management Instance VLAN only if the VLAN operates in non-routing mode.

Before You Begin

- Ensure the IP address you plan to assign is not in use by an existing VLAN or CLIP IP subnet configured on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Ipv6Address** tab.
4. Select **Insert**.
5. Select the Segmented Management Instance interface type.
6. Type the address information.
7. Select **Insert**.

Ipv6Address Field Descriptions

Use the data in the following table to use the **Ipv6Address** tab.

Name	Description
InstanceId	Specifies the interface instance.
Address	Specifies the IPv6 address for the interface instance. Ensure that the management CLIP IP address does not fall into the range of a configured VLAN IP address range as this is not allowed.
PrefixLength	Specifies the prefix length for the IPv6 address.
AddrOrigin	Shows the IPv6 address origin.
IntfName	Shows the interface name.
DadStatus	Shows the IPv6 DAD status of the address.

View Segmented Management Instance Statistics

View operational statistics for the Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Mgmt**.
3. Select the **Interface** tab.
4. Select a Management Instance by placing the cursor in a cell within the applicable row.
5. Select **Graph**.

Interface Counters Field Descriptions

Use the data in the following table to use the **Interface Counters** tab.

Name	Description
RxPkts	Counts the packets received on the Segmented Management Instance.
RxError	Counts the packets received with errors on the Segmented Management Instance.
RxDrop	Counts the packets received and dropped on the Segmented Management Instance.
TxPkts	Counts the packets transmitted on the Segmented Management Instance.
TxError	Counts the packets transmitted with errors on the Segmented Management Instance.
TxDrop	Counts the packets dropped before transmission on the Segmented Management Instance.

View IP Address Statistics for a Segmented Management Instance

Use this task to view IP address statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IP** tab.
4. To clear IP statistics, select **Clear Stats**.
5. To clear IP counters, select **Clear Counters**.

IP Field Descriptions

Use the data in the following table to use the **IP** tab.

Name	Description
InReceives	Shows the inbound packet statistics.
InHdrErrors	Shows the inbound packets with header errors statistics.
InAddrErrors	Shows the inbound packets with address errors statistics.
InUnknownProtos	Shows the inbound packets with unknown protocols statistics.
InDiscards	Shows the inbound packets discarded statistics.
InDelivers	Shows the inbound packets delivered statistics.
OutRequests	Shows the outbound packet requests statistics.
OutDiscards	Shows the outbound packets discarded statistics.
OutNotRoutes	Shows the outbound packets with no routes statistics.
ForwDatagrams	Shows the forwarded datagram packets statistics.
ReasmTimeout	Shows the packet reassembly timeouts statistics.
ReasmReqds	Shows the packet reassembly requests statistics.
ReasmOKs	Shows the successfully reassembled packets statistics.
ReasmFails	Shows the failed reassembled packets statistics.
FragOKs	Shows the successfully fragmented packets statistics.

Name	Description
FragFails	Shows the failed fragmented packets statistics.
FragCreates	Shows the fragments created statistics.

View IPv6 Address Statistics for a Segmented Management Instance

Use this task to view IPv6 address statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IPv6** tab.
4. To clear IP statistics, select **Clear Stats**.
5. To clear IP counters, select **Clear Counters**.

IPv6 Field Descriptions

Use the data in the following table to use the **IPv6** tab.

Name	Description
InReceives	Shows the inbound packet statistics.
InHdrErrors	Shows the inbound packets with header errors statistics.
InAddrErrors	Shows the inbound packets with address errors statistics.
InUnknownProtos	Shows the inbound packets with unknown protocols statistics.
InDiscards	Shows the inbound packets discarded statistics.
InDelivers	Shows the inbound packets delivered statistics.
InTooBigErrors	Shows the inbound packets too big errors statistics.
InNoRoutes	Shows the inbound packets with no routes statistics.
InTruncatedPkts	Shows the inbound packets truncated statistics.
InMcastPkts	Shows the inbound multicast packets statistics.
InOctets	Shows the inbound octets statistics.
InMcastOctets	Shows the inbound multicast octets statistics.

Name	Description
InBcastOctets	Shows the inbound broadcast octets statistics
InNoECTPkts	Shows the inbound packets with no Explicit Congestion Notification (ECN) statistics.
InECTIPkts	Shows the inbound packets with ECT(1) statistics.
InECT0Pkts	Shows the inbound packets with ECT(0) statistics.
InCEPkts	Shows the inbound packets with Congestion Encountered (CE) statistics.
OutRequests	Shows the outbound packet requests statistics.
OutDiscards	Shows the outbound packets discarded statistics.
OutNoRoutes	Shows the outbound packets with no routes statistics.
OutForwDatagrams	Shows the forwarded datagram packets statistics.
OutMcastPkts	Shows the outbound multicast packets statistics.
OutOctets	Shows the outbound octets statistics.
OutMcastOctets	Shows the outbound multicast octets statistics.
OutBcastOctets	Shows the outbound broadcast octets statistics.
ReasmTimeout	Shows the packet reassembly timeouts statistics.
ReasmReqds	Shows the packet reassembly requests statistics.
ReasmOKs	Shows the successfully reassembled packets statistics.
ReasmFails	Shows the failed reassembled packets statistics.
FragOKs	Shows the successfully fragmented packets statistics.
FragFails	Shows the failed fragmented packets statistics.
FragCreates	Shows the fragments created statistics.

View IP ICMP Statistics for a Segmented Management Instance

Use this task to view IP ICMP statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IP-ICMP** tab.
4. To clear IP statistics, select **Clear Stats**.
5. To clear IP counters, select **Clear Counters**.

IP-ICMP Field Descriptions

Use the data in the following table to use the **IP-ICMP** tab.

Name	Description
InMsgs	Shows the inbound messages statistics.
InErrors	Shows the inbound errors statistics.
InCsumErrors	Shows the inbound checksum errors statistics.
InDestUnreachs	Shows the inbound destination unreachable statistics.
InTimeExcds	Shows the inbound time exceeded statistics.
InParmProbs	Shows the inbound parameter problems statistics.
InSrcQuenchs	Shows the inbound source quenchs statistics.
InRedirects	Shows the inbound redirects statistics.
InEchos	Shows the inbound echos statistics.
InEchoReps	Shows the inbound echo replies statistics.
InTimestamps	Shows the inbound timestamps statistics.
InTimestampsReps	Shows the inbound timestamp replies statistics.
InAddrMasks	Shows the inbound address masks statistics.
InAddrMaskReps	Shows the inbound address mask replies statistics.
OutMsgs	Shows the outbound messages statistics.
OutErrors	Shows the outbound errors statistics.
OutDestUnreachs	Shows the outbound destination unreachable statistics.
OutTimeExcds	Shows the outbound time exceeded statistics.
OutParmProbs	Shows the outbound parameter problems statistics.

Name	Description
OutSrcQuenchs	Shows the outbound source quenchs statistics.
OutRedirects	Shows the outbound redirects statistics.
OutEchos	Shows the outbound echos statistics.
OutEchoReps	Shows the outbound echo replies statistics.
OutTimestamps	Shows the outbound timestamps statistics.
OutTimestampReps	Shows the outbound timestamps replies statistics.
OutAddrMasks	Shows the outbound address masks statistics.
MsgInType0	Shows the inbound Type0 messages statistics.
MsgOutType8	Shows the outbound Type8 messages statistics.

View IPv6 ICMP Statistics for a Segmented Management Instance

Use this task to view IPv6 ICMP statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IPv6-ICMP** tab.
4. To clear IP statistics, select **Clear Stats**.
5. To clear IP counters, select **Clear Counters**.

IPv6-ICMP Field Descriptions

Use the data in the following table to use the **IPv6-ICMP** tab.

Name	Description
InMsgs	Shows the inbound messages statistics.
InErrors	Shows the inbound errors statistics.
InCsumErrors	Shows the inbound checksum errors statistics.
InDestUnreachs	Shows the inbound destination unreachable statistics.
InTimeExcds	Shows the inbound time exceeded statistics.
InParmProbs	Shows the inbound parameter problems statistics.

Name	Description
InPktTooBigs	Shows the inbound packets too big statistics.
InRedirects	Shows the inbound redirects statistics.
InEchos	Shows the inbound echos statistics.
InEchoReplies	Shows the inbound echo replies statistics.
InGroupMembQueries	Shows the inbound group member queries statistics.
InGroupMembResponses	Shows the inbound group member responses statistics.
InGroupMembReductions	Shows the inbound group member reductions statistics.
InRouterSolicits	Shows the inbound router solicits statistics.
InRouterAdvertisements	Shows the inbound router advertisements statistics.
InNeighborSolicits	Shows the inbound neighbor solicits statistics.
InNeighborAdvertisements	Shows the inbound neighbor advertisements statistics.
InMLDv2Reports	Shows the inbound MLDv2 reports statistics.
InType134	Shows the inbound type134 statistics.
InType136	Shows the inbound type136 statistics.
OutMsgs	Shows the outbound messages statistics.
OutErrors	Shows the outbound errors statistics.
OutDestUnreachs	Shows the outbound destination unreachable statistics.
OutTimeExcds	Shows the outbound time exceeded statistics.
OutParmProbs	Shows the outbound parameter problems statistics.
OutPktTooBigs	Shows the outbound packets too big statistics.
OutRedirects	Shows the outbound redirects statistics.
OutEchos	Shows the outbound echos statistics.
OutEchoReps	Shows the outbound echo replies statistics.
OutGroupMembQueries	Shows the outbound group member queries statistics.
OutGroupMembResponses	Shows the outbound group member responses statistics.

Name	Description
OutGroupMembReductions	Shows the outbound group member reductions statistics.
OutRouterStatistics	Shows the outbound router statistics
OutRouterAdvertisements	Shows the outbound router advertisements statistics.
OutNeighborSolicits	Shows the outbound neighbor solicits statistics.
OutNeighborAdvertisements	Shows the outbound neighbor advertisements statistics.
OutMLDv2Reports	Shows the outbound MLDv2 reports statistics.
OutType133	Shows the outbound Type133 statistics.
OutType135	Shows the outbound Type135 statistics.
OutType143	Shows the outbound Type143 statistics.

View UDP Statistics for a Segmented Management Instance

Use this task to view UDP statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IP/IPv6 UDP** tab.

IP/IPv6 UDP Field Descriptions

Use the data in the following table to use the **IP/IPv6 UDP** tab.

Name	Description
IPVersion	Shows the IP address version as ipv4 or ipv6.
InDatagrams	Shows the input datagram statistics.
NoPorts	Shows the number of ports statistics.
InErrors	Shows the input errors statistics.
OutDatagrams	Shows the output datagram statistics.
IgnoredMulti	Show the ignored multiport statistics.
RcvbufErrors	Shows the received buffer errors statistics.
SndbufErrors	Shows the send buffer errors statistics.
InCsumErrors	Shows the input checksum errors statistics.
Clear	Specifies to clear the statistics. Default is false.

View TCP Statistics for a Segmented Management Instance

Use this task to view TCP statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IP/IPv6 TCP** tab.

IP/IPv6 TCP Field Descriptions

Use the data in the following table to use the **IP/IPv6 TCP** tab.

Name	Description
IPVersion	Shows the IP address version as ipv4 or ipv6.
ActiveOpens	Shows the active open TCP connections statistics.
PassiveOpens	Shows the passive open TCP connections statistics.
AttemptFails	Shows the TCP connection attempt failures statistics.
EstabResets	Shows the TCP connection established resets statistics.
InSegs	Shows the input segments statistics.
OutSegs	Shows the output segments statistics.
RetransSegs	Shows the retransmit segments statistics.
InErrs	Shows the input checksum errors statistics.
OutRsts	Shows the output resets statistics.
InCsumErrors	Shows the input checksum errors statistics.
Clear	Specifies to clear the statistics. Default is false.

View TCP Connections and UDP Endpoints Statistics for a Segmented Management Instance

Use this task to view TCP connections and UDP endpoints statistics for a Segmented Management Instance.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **IP/IPv6 Socket (TCP/UDP)** tab.

IP/IPv6 Socket (TCP/UDP) *Field Descriptions*

Use the data in the following table to use the **IP/IPv6 Socket (TCP/UDP)** tab.

Name	Description
IPVersion	Shows the IP address version as ipv4 or ipv6.
Type	Shows the connection type as tcp or udp.
Index	Shows the index ID for the connection.
State	Shows the link state for the connection.
RecvQ	Shows the connection received quantity.
SendQ	Shows the connection sent quantity.
LocalAddressAndPort	Shows the local IP address and port.
PeerAddressAndPort	Shows the peer IP address and port.



Command Line Interface

[Command Line Interface Fundamentals](#) on page 178

[CLI Procedures](#) on page 193

Table 15: Command Line Interface product support

Feature	Product	Release introduced
Command Line Interface (CLI)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

For information about specific commands, see *Fabric Engine Command References*.

Command Line Interface Fundamentals

This section describes the Command Line Interface (CLI).

CLI is an industry standard command line interface that you can use for single-device management.

CLI Command Modes

CLI command modes provide specific sets of CLI commands. When you log onto the switch, you are in User EXEC mode with limited commands. While in a higher mode, you can access most commands from lower modes, except if they conflict with commands of your current mode.

There are two categories of CLI commands: show commands and configuration commands. You can use show commands from multiple command modes with the same results; they show the same configuration information regardless of the

command mode. Configuration command results, however, might be dependent on the command mode from which a configuration command is used. For example, an **enable** command used in Global Configuration mode will enable a feature globally, and the same command used from one of the interface command modes will enable a feature for a specific interface only.

The following figure illustrates the navigation paths for the various command modes:

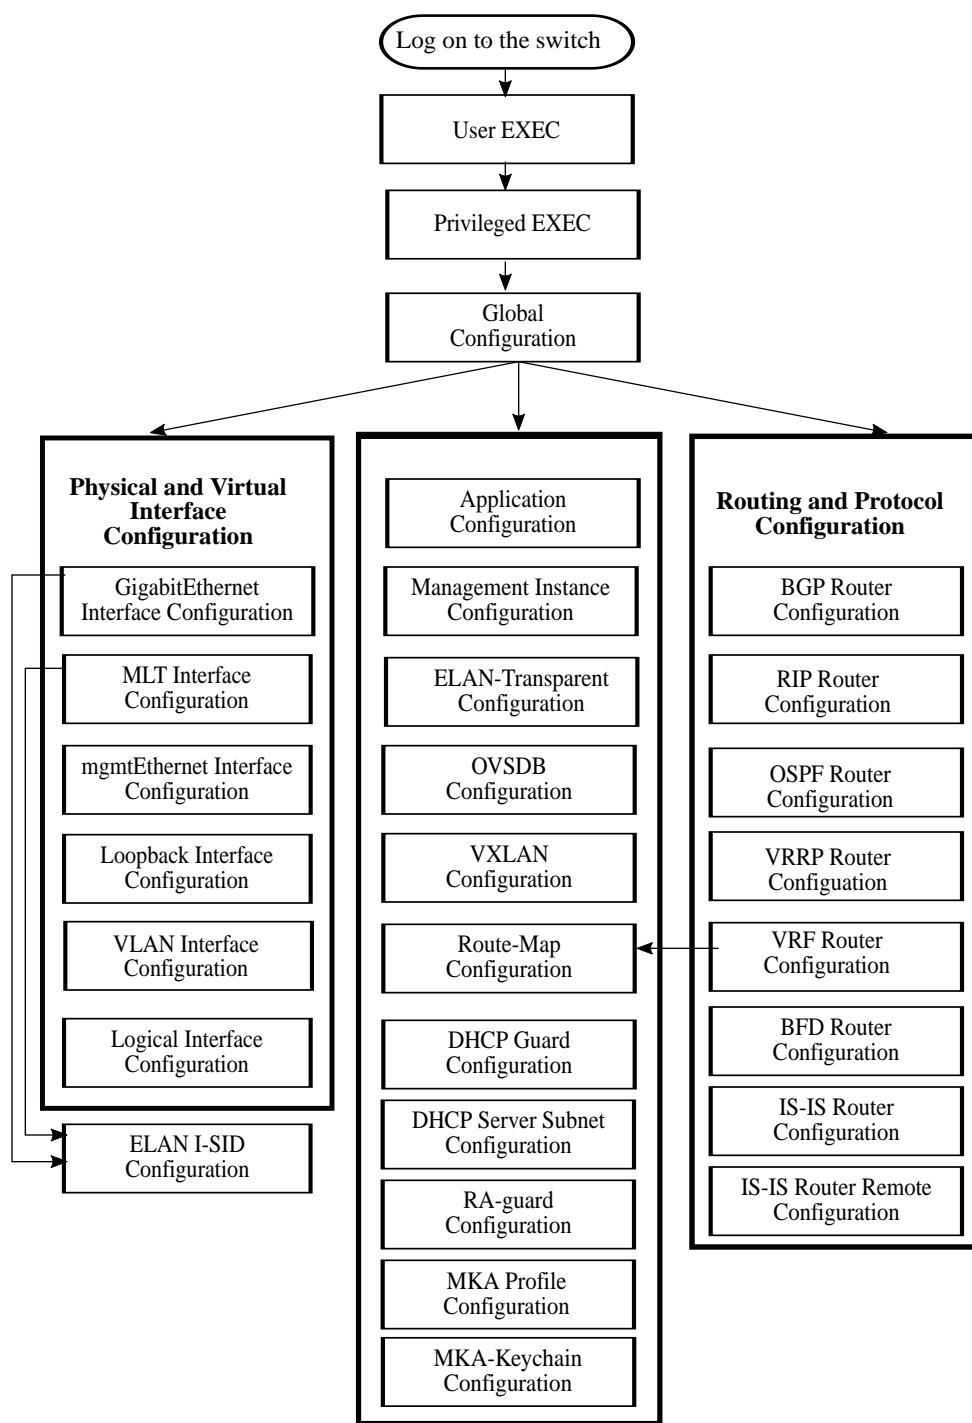


Figure 11: CLI Command Mode Navigation

Your user authorization credentials determine what commands are available to you in Privileged EXEC mode and all higher-level modes. See [Managing the Switch](#) on page 246 for more information.

To navigate from higher-level modes to lower-level modes, use the following commands:

- **exit** to navigate from a higher-level mode to a lower-level mode, down to Privileged EXEC mode
- **end** to navigate from any command mode directly to Privileged EXEC mode
- **disable** to navigate from Privileged EXEC mode to User EXEC mode
- **logout** to terminate the CLI session from any command mode

The following table describes the various command modes, including the CLI command to access each mode, the command prompt that displays in each mode, and a description of the purpose of the mode.



Note

Some command modes are hardware dependent. If any of the following commands modes do not display on your hardware, they are not supported or applicable.

Table 16: CLI Command Mode Summary

Command mode	Command to access mode	Prompt displayed in mode	Description
User EXEC	None required; default mode	>	View configuration settings and connection status.
Privileged EXEC	enable	#	Configure limited device-wide settings.
	Note: Depending on feature configuration, you can be prompted to enter a user name and password to access Privileged EXEC mode. For more information, see Authentication for Privileged EXEC Command Mode on page 193.		
Global Configuration	configure {terminal network}	(config) #	From a terminal or TFTP server, configure device-wide global parameters on a running configuration, or specify the file name of a configuration file.

Table 16: CLI Command Mode Summary (continued)

Command mode	Command to access mode	Prompt displayed in mode	Description
GigabitEthernet Interface Configuration	interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}	(config-if) #	Configure chassis operations and features on a physical port.
MLT Interface Configuration	interface mlt <1-512>	(config-mlt) #	Configure an MLT interface.
mgmtEthernet Interface Configuration	interface mgmtEthernet mgmt	(config-if) #	Configure a dedicated physical management port (if supported on the hardware).
Loopback Interface Configuration	interface loopback <1-256>	(config-if) #	Configure a loopback CLIP interface.
VLAN Interface Configuration	interface vlan <1-4059>	(config-if) #	Configure port-based, policy-based, private, or SPBM B-VLANs
Logical Interface Configuration	logical-intf isis <1-255>	Layer 2: (config-isis-<1-255>) # Layer 3: (config-isis-<1-255>-<A.B.C.D>) #	Configure a logical Layer 2 or Layer 3 IS-IS interface.
BGP Router Configuration	router bgp	(router-bgp) #	Configure device-wide BGP routing protocol settings.
RIP Router Configuration	router rip	(config-rip) #	Configure device-wide RIP routing protocol settings.
OSPF Router Configuration	router ospf	(config-ospf) #	Configure device-wide OSPF routing protocol settings.
IS-IS Router Configuration	router isis	(config-isis) #	Configure device-wide IS-IS routing protocol settings.

Table 16: CLI Command Mode Summary (continued)

Command mode	Command to access mode	Prompt displayed in mode	Description
IS-IS Router Remote Configuration	router isis remote	(config-isis-remote) #	Configure the Multi-area SPB parameters like area virtual node, Shortest Path Bridging MAC (SPBM), manual area and so on.
VRF Router Configuration	router vrf WORD<1-16>	(router-vrf) #	Configure a VRF instance.
VRRP Router Configuration	router vrrp	(config-vrrp) #	Configure device-wide VRRP protocol settings.
Application Configuration	application	(config-app) #	Configure custom applications, such as RESTCONF.
Management Instance Configuration	mgmt <clip oob vlan>	(mgmt:clip) # or (mgmt:oob) # or (mgmt:vlan) #	Configure a segmented management CLIP, Out-of-Band (OOB), or VLAN instance.
ELAN I-SID Configuration	i-sid <1-16777215> [elan]	(elan:<1-16777215>) #	Add ports and traffic to a Switched UNI I-SID on a GigabitEthernet or MLT interface.
ELAN-Transparent Configuration	i-sid <1-16777215> elan-transparent	(elan-tp:<1-16777215>) #	Add ports and MLT interfaces to an ELAN-Transparent based service.
OVSDB Configuration	ovsdb	(config-ovsdb) #	Configure OVSDB protocol support for VXLAN Gateway.
Route-Map Configuration	route-map WORD<1-64> <1-65535>	(route-map) #	Configure device-wide or VRF instance-specific route map policy settings.
DHCP-guard Configuration	ipv6 fhs dhcp-guard policy WORD<1-64>	(config-dhcpguard) #	Configure DHCPv6 for advertised address-based, prefix-based, and preference-based filtering.

Table 16: CLI Command Mode Summary (continued)

Command mode	Command to access mode	Prompt displayed in mode	Description
DHCP Server Subnet Configuration	ip dhcp-server subnet <A.B.C.D/X>	(config-dhcp) #	Configure a subnet for clients that request addresses from the hosted DHCP Server.
RA-guard Configuration	ipv6 fhs ra-guard policy WORD<1-64>	(config-raguard) #	Configure RA Guard for advertised IPv6 and MAC address-based, IPv6 prefix-based, preference-based, hop count limit-based, and default router preference-based filtering.
VXLAN Configuration	vnid <1-16777215> i-sid <1-16777215>	(vxlan:<1-16777215>) #	Associate port or MLT interface VLANs, configure VXLAN endpoints and untagged traffic.
MKA Profile Configuration	macsec mka profile WORD<1-16>	(mka-profile) #	Configure replay protection and confidentiality offset for an MKA profile.
MKA Keychain Configuration	macsec mka keychain WORD<5-32>	(mka-keychain)	Configure an MKA keychain.
BFD Router Configuration	router bfd	(router-bfd) #	Configure device-wide BFD settings.

Special CLI Command Modes

A special CLI command mode provides a set of specific commands that are different in operation and syntax from the standard switch CLI command modes and the commands available in them. For example, a set of CLI commands that configure services on a Virtual Machine (VM).



Note

Special CLI command modes can be hardware dependent. For information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

The following table describes the special command mode.

Table 17: Special CLI Command Mode Summary

Special command mode	Command mode navigation	Command to access mode	Prompt displayed in mode	Description
Fabric IPsec Gateway	Accessible from Privileged EXEC	virtual-service WORD<1-128> console Note: Type CTRL+Y to exit the console.	FIGW>	Configure services like IPsec, fragmentation and reassembly, and to manage the Fabric IPsec Gateway VM.

Default User Names and Passwords for CLI

The following table contains the default user names and passwords that you can use to log on to the switch using the command line interface (CLI). For more information about how to change passwords, see [Security](#) on page 3367.

Table 18: CLI default user names and passwords

User name	Password	Description
rwa	rwa	read-write-all
rw	rw	read-write
ro	ro	read-only
l1	l1	layer 1
l2	l2	layer 2
l3	l3	layer 3

You can create up to a maximum of 10 CLI users for each role. For more information, see [Multiple CLI Users for Each Role](#) on page 263.

If you enable enhanced secure mode, the user names and passwords are different than the default values documented in the preceding table. For more information on enhanced secure mode, see [Enhanced Secure Mode](#) on page 266.



Important

The default passwords and community strings are documented and well known. As a best practice, change the default passwords and community strings immediately after you first log on. For more information about how to change user names and passwords, see [Security](#) on page 3367.

Documentation convention for the port variable

Commands that require you to enter one or more port numbers on the switch use the parameter `{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}` in the syntax. The following table specifies the rules for using `{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}`.

Syntax	How to use
<code>{slot/ port[/sub- port]}</code>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. For example, 1/1 indicates the first port on slot 1. 1/41/1 indicates the first channel on slot 1, port 41.
<code>{slot/ port[/sub- port] [- slot/port[/ sub-port]] [, ...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. For example, 1/1-1/3 indicates ports 1 to 3 on slot 1, or 1/41/1,1/41/3 indicates the first and third channels of slot 1, port 41.</p>

Command Completion

The CLI provides potential command completions to the command string. Completions are provided by using a question mark (?) or by using the CLI autocompletion feature.

Question Mark Command Completion

The ? command completion is available for any valid command. By typing a command and using a ? as the last argument in the command, the system returns a list of possible command completions from the point of the ?. A short description is provided with each possible completion.

If you enter the following command:

```
Switch:1(config-isis)#redistribute ?
```

CLI provides a list of completions for the **redistribute** ? command.

```
Switch:1(config-isis)#redistribute ?
  direct      isis redistribute direct command
  ospf        isis redistribute ospf command
  rip         isis redistribute rip command
  static      isis redistribute static command
```

All the parameters listed under redistribute indicate sub-context commands.

You must use one of the available completions, and if necessary, use the command completion help again to find the next completion.

```
Switch:1(config-isis)#redistribute direct ?
enable      Enable isis redistribute direct command
metric      Isis route redistribute metric
metric-type Set isis redistribute metric type
route-map   Set isis redistribute direct route-policy
subnets    Set isis redistribute subnets
<cr>
```

When you see `<cr>` (Carriage Return/Enter Key) in the list with the additional choices, this means that no additional parameters are required to run the CLI command. However, the additional choices listed could be peer commands or sub-context commands.

For example, the parameters listed under **redistribute direct ?** are peer commands. You can enter these peer commands on the same line as the root command, for example **redistribute direct enable**. However, the `<cr>` indicates that you can also enter the **redistribute direct** command only and this command does not require any additional parameters at this level.

CLI Autocompletion

CLI autocompletion is a feature that you can use to automatically fill in the unique parts of a command string rather than typing the entire command. Autocompletion makes the CLI experience easier and prevents mistakes in spelling that force you to re-enter the command.

Autocompletion completes the token in the command as soon as it becomes unique.

The **Tab** key autocompletes the command without running the command, and places the cursor immediately after the last character. The **Enter** key autocompletes the command and runs it.

To enable redistribution of IS-IS direct routes, when you use `redistribute direct ?`, you see multiple possible sub-context commands.

```
Switch:1(config-isis)#redistribute direct ?
enable      Enable isis redistribute direct command
metric      Isis route redistribute metric
metric-type Set isis redistribute metric type
route-map   Set isis redistribute direct route-policy
subnets    Set isis redistribute subnets
<cr>
```

If you type `redistribute direct m`, and then press the **Tab** key, the system completes the command to the following point:

```
Switch:1(config-isis)#redistribute direct m<TAB>
Switch:1(config-isis)#redistribute direct metric
```

To configure the metric type, the letter **m** is not unique to that parameter. You must type **-t**, and then press **Tab** to autocomplete the syntax:

```
Switch:1(config-isis)#redistribute direct metric-t<TAB>
Switch:1(config-isis)#redistribute direct metric-type
```



Note

Keep in mind that new releases can modify or add commands, which means the uniqueness for command completion can change.

Command Completion for VRF Names

Command completion using the question mark and **Tab** key also applies to VRF names in show and configuration commands. Ensure the VRF name is unique before you press **Enter** to run the command. For example, a switch has the following user-created VRF names:

- accounting
- test1
- test2



Note

This function translates a VRF name to lower case letters and uses that format even if you created the VRF with a mix of lower case and uppercase letters.

The following table illustrates the results of various actions.

Table 19: Command completion results for VRF name

Command action	Result
<code>show ip route vrf te<Tab></code>	Displays <code>show ip route vrf test</code> at the prompt. Because <code>test</code> is not unique, you cannot press Enter to run the command without further modification of the VRF name.
<code>show ip route vrf te<Enter></code>	Because <code>test</code> is not unique, the switch displays <code>% Ambiguous command. See '^' marker.</code> , and the command does not run.
<code>show ip route vrf a<Tab></code>	Displays <code>show ip route vrf accounting</code> at the prompt. Because <code>accounting</code> is unique, you can then press Enter to run the command.
<code>show ip route vrf a<Enter></code>	Because <code>accounting</code> is unique, the command runs and the switch displays the command output.



Note

Autocompletion of VRF names does not apply to unrestricted and untrusted Segmented VRFs.

Using the question mark (?) to see the full list of VRF names for command completion includes the GlobalRouter and, on switches that support multiple VRFs, the MgmtRouter.

default command operator

You can reset the modified configuration of a command to the default configuration by using the default operator. For more information about the default value for each command, see *Fabric Engine Command References*.

Use the ? command completion along with the default keyword in each configuration mode, to view the list of commands that support the default operator. For more information, see [Command Completion](#) on page 186.

Examples

Configure **csnp-interval** to its default value. The default value of **csnp-interval** is 10 seconds.

```
Switch:1>show isis
```

```
=====
                        ISIS General Info
=====
AdminState : disabled
RouterType : Level 1
System ID : e45d.523c.6484
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 200
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name :
ip source-address :
ipv6 source-address :
ip tunnel source-address :
Tunnel vrf :
ip tunnel mtu :
Num of Interfaces : 1
Num of Area Addresses : 0
Inband Mgmt Clip Ip :
backbone : disabled
Dynamically Learned Area : 00.0000.0000
Hello Padding : enabled
FAN Member : Yes
Multi-Area OperState : disabled
Multi-Area Flags :
```

```
Switch:1>enable
```

```
Switch:1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config)#router isis
```

```
Switch:1(config-isis)#default csnp-interval
```

```
Switch:1(config-isis)#show isis
```

```
=====
                        ISIS General Info
=====
```

```

=====
AdminState : disabled
RouterType : Level 1
System ID : e45d.523c.6484
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name :
ip source-address :
ipv6 source-address :
ip tunnel source-address :
Tunnel vrf :
ip tunnel mtu :
Num of Interfaces : 1
Num of Area Addresses : 0
Inband Mgmt Clip Ip :
backbone : disabled
Dynamically Learned Area : 00.0000.0000
Hello Padding : enabled
FAN Member : Yes
Multi-Area OperState : disabled
Multi-Area Flags :

```

View the IP configuration commands for an MLT interface that support the default operator.

```

Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#default ?
Default settings
fa                Set Fabric Attach configuration to default on mlt
flex-uni          Set flex-uni to default on mlt interface
ip               Default IP configurations on MTL interface
isis             Set interface level isis parameters to default value
lacp             Set lacp for specific mlt to default
smlt             Create default smlt on a specific mlt
svlan-prototype  Set vlan port type to default
virtual-ist      Create virtual-ist on MLT with default value
Switch:1(config-mlt)#default ip ?
Default IP configurations on MLT interface
arp-inspection   Default arp inspection configuration
dhcp-snooping    Default dhcp snooping configuration
Switch:1(config-mlt)#default ip arp-inspection ?
<cr>

```

default all command

You can use the **default all** command to return all parameters configured on a port to their default values instead of using the individual default operator with each command on the port. Additionally, you can use the *auto-sense* parameter with this command to return all values to the existing configuration and enable the Auto-sense

feature on the port. You can append the **-y** parameter to bypass the confirmation question.

```
Switch:1(config)#interface gigabitethernet 1/5
Switch:1(config-if)#default all
Warning: This command will default all port settings, including removal of VLAN, MLT,
IS-IS, ACL or any other configuration relating to the port. Are you sure you want to
continue (y/n) ? y
```

no command operator

You can use the no operator in a command to negate a configuration. Based on the functionality of the command, you can perform negations, such as disable, delete, remove, or reset to the default configuration. For more information about the no operator for each command, see *Fabric Engine Command References*.

Use the ? command completion along with the no keyword to view the list of commands that support the no operator in each configuration mode. For more information, see [Command Completion](#) on page 186.

Negate the automatic virtual link that provides automatic dynamic backup link for OSPF traffic.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router ospf
Switch:1(config-ospf)#no auto-vlink
```

Remove an IP address configuration from VLAN.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 3
Switch:1(config-if)#no ip address 192.0.2.4
```

View the commands that can negate a configuration in RIP router configuration mode.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router rip
Switch:1(config-rip)#no ?
Negate a command or set its defaults
ipv6          Disable ipv6 configurations
network       Disable rip on an ip network
redistribute  To disable/delete redistribute globally
Switch:1(config-rip)#no network ?
{A.B.C.D} Network ip address
Switch:1(config-rip)#no network 192.0.2.4 ?
<cr>
```

GREP with CLI show command

You can use Global Regular Expression Print (GREP) with **show** commands to filter the output based on match criteria.

Enter the **show** command followed by the pipe (|) character, followed by the GREP filter command. The **show** command output contains only the lines that match the GREP filter pattern.



Note

The **show fulltech** command does not support GREP filters.

The following GREP filter commands are supported.

GREP filter function	Description
<i>begin</i>	Displays the output of a command starting from the first line, which matches the given pattern.
<i>count</i>	Counts the number of lines in the output of a command.
<i>exclude</i>	Displays only the output lines which do not match the given pattern. The lines matching the pattern are discarded.
<i>head</i>	Limits the output of a command to the first few lines. If a number is not specified then only the first 10 lines display.
<i>include</i>	Displays only the output lines which match the given pattern.
<i>no-more</i>	Temporarily disables pagination for the output of an CLI command. When the lines of output exceed the terminal length, you are not prompted to continue or quit but the entire output of the command continues to be displayed. The effect is similar to setting terminal length 0 but only for the current command.
<i>tail</i>	Limits the output of a command to the last few lines. If a number is not specified then only the last 10 lines display.

Timestamp in show command outputs

The output for all CLI show commands includes a timestamp header to indicate when the command output was generated. This information can be helpful when communicating with Support.

The following command output shows a timestamp example.

```
Switch:1#show alarm statistics
*****
                Command Execution Time: Wed Nov 07 19:55:15 2018 UTC
*****
=====
                        ALARM STATISTICS
=====
```

PERSISTENT ALARM 0	PERSISTENT ACTIVE 0	PERSISTENT CLEARED 0	PERSISTENT WRPRD 0	DYNAMIC ALARM 11	DYNAMIC ACTIVE 8	DYNAMIC CLEARED 3	DYNAMIC WRPRD 0
--------------------------	---------------------------	----------------------------	--------------------------	------------------------	------------------------	-------------------------	-----------------------

Authentication for Privileged EXEC Command Mode

For enhanced security, you can configure the switch to authenticate the user before providing access to Privileged EXEC command mode. When you configure password authentication, the switch prompts the user to enter a user name and password to access Privileged EXEC command mode from User EXEC command mode. Use the same user name and password used to establish the Telnet or SSH connection to the switch.

For more information about how to configure Privileged EXEC authentication, see [Authentication for Privileged EXEC Command Mode](#) on page 3375.

CLI Procedures

This section contains information about common CLI tasks. You can access CLI during runtime to manage the switch.

Log on to the Software

Before You Begin

- The first time you connect to the switch, you must log on to CLI using the direct console port.

About This Task

After you first connect to CLI you can log on to the software using the default user name and password. For more information about the default user names and passwords, see [Default User Names and Passwords for CLI](#) on page 185.

Procedure

1. At the login prompt, enter the user name.
2. At the password prompt, enter the password.

View the Configuration

You can view the running configuration using the show command.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View the running configuration:
`show running-config`

Example

```

5320-48T-8XE-FabricEngine:1>enable
5320-48T-8XE-FabricEngine:1#show running-config
Preparing to Display Configuration...
#
# Mon Dec 13 23:44:36 2021 EET
# box type           : 5320-48T-8XE-FabricEngine
# software version    : 8.6.0.0
# cli mode            : ECLI
#
#Card Info :
# Slot 1 :
#   CardType          : 5320-48T-8XE-FabricEngine
#   CardDescription    : 5320-48T-8XE-FabricEngine
#   CardSerial#        : TB042128K-H0030
#   CardPart#          : 801107-00-02
#   CardAssemblyDate   : 20210712
#   CardHWRevision     : 02
#   CardHWConfig       :
#   AdminStatus        : up
#   OperStatus         : up
#
#!end
#
config terminal
#
# BOOT CONFIGURATION
#
boot config flags ftpd
boot config flags sshd
#boot config sio console baud 115200 1
# end boot flags
#
# SPBM CONFIGURATION
#
spbm
spbm ethertype 0x8100
#
# CLI CONFIGURATION
--More-- (q = quit)

```

Save the Configuration

After you change the configuration, you must save the changes to the module. Save the configuration to a file to retain the configuration settings.

About This Task

File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP) support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Save the running configuration:
`save config [backup WORD<1-99>] [file WORD<1-99>] [verbose]`

Example

Save the configuration to the default location:

```
Switch:1#save config
```

Identify the file as a backup file and designate a location to save the file:

```
Switch:1#save config backup 198.51.100.1/configs/backup.cfg
```

Variable Definitions

The following table defines parameters for the **save config** command.

Variable	Value
<i>backup WORD<1-99></i>	Saves the specified file name and identifies the file as a backup file. <i>WORD<1-99></i> uses one of the following formats: <ul style="list-style-type: none">• a.b.c.d:<file>• /intflash/<file> The file name, including the directory structure, up to 1 to 99 characters.
<i>file WORD<1-99></i>	Specifies the file name in one of the following formats: <ul style="list-style-type: none">• /intflash/<file>• a.b.c.d:<file> The file name, including the directory structure, up to 1 to 99 characters.
<i>verbose</i>	Saves the default and current configuration. If you omit this parameter, the command saves only parameters you change.

Configure the Web Server

Perform this procedure to enable and manage the web server using the Command Line Interface (CLI). After you enable the web server, you can connect to EDM.

About This Task

This procedure assumes that you use the default port assignments. You can change the port number used for HTTP and HTTPS.



Important

To enable HTTP access to the device, you must disable the secure-only option. To enable HTTPS access to the device, the secure-only option is enabled by default.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable the web server:

```
web-server enable
```



Important

When you enable the web management interface using the **web-server enable** command, the system prompts you to change the admin and read-only user default passwords. If you do not change the default passwords, you are prompted to change them every time you use the **web-server enable** command.

3. (Optional) Disable the secure-only option (for HTTP access) :

```
no web-server secure-only
```

4. Enable read-only user:

```
web-server read-only-user enable
```

5. Display the web server status:

```
show web-server
```

Example

```
Switch:1>show web-server
Web Server Info :
```

```

Status                : on
Secure-only            : enabled
TLS-minimum-version    : tlsv12
RO Username Status     : enabled
RO Username            : user
RO Password            : *****
RWA Username           : admin
RWA Password           : *****
Def-display-rows       : 30
Inactivity timeout     : 900 sec
Html help tftp source-dir :
HttpPort               : 80
HttpsPort              : 443
NumHits                : 0
NumAccessChecks        : 0
NumAccessBlocks        : 0
```

```

NumRxErroRs          : 0

NumTxErrors           : 0
NumSetRequest         : 1
Minimum password length : 8
Last Host Access Blocked : 0.0.0.0
In use certificate     : Self signed
Certificate Truspoint CA Name :
Certificate with Subject Name :

Ciphers-Tls           : TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
                        TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
                        TLS_RSA_WITH_AES_256_CBC_SHA256
TLS_RSA_WITH_AES_256_CBC_SHA
                        TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA
SSL renegotiation      : enabled

```

Variable Definitions

Use the data in the following table to use the **web-server** command.

Variable	Value
<i>def-display-rows</i> <10-100>	Configures the number of rows each page displays. The default is 30.
<i>enable</i>	Enables the web interface. The default is disabled.
<i>help-tftp</i> <WORD/0-256>	Configures the TFTP or FTP directory for Help files, in one of the following formats: a.b.c.d:/ peer:/ [<dir>]. The path can use 0-256 characters. The following example paths illustrate the correct format: <ul style="list-style-type: none"> • 192.0.2.1:/help • 192.0.2.1:/
<i>http-port</i> <80-49151>	Configures the web server HTTP port. The default port is 80.
<i>https-port</i> <443-49151>	Configure the web server HTTPS port. The default port is 443.
<i>inactivity-timeout</i> <30-65535>	Configures the web-server session inactivity timeout. The default is 900 seconds (15 minutes).
<i>password</i> {ro rwa} WORD<1-80>	Configures the user names and passwords for the web interface. The default user name for the RO account is user. The default user name for the RWA account is admin.
<i>password min-passwd-len</i> <1-80>	Configures the minimum password length. By default, the minimum password length is 8 characters.

Variable	Value
<i>read-only-user</i>	Enables read-only user for the web server. The default is disabled.
<i>secure-only</i>	Enables secure-only access for the web server.
<i>ssl-renegotiation</i>	Enables SSL renegotiation in the web server. The default is enabled.
<i>tls-min-ver</i> <tlsv10 tlsv11 tlsv12>	<p>Configures the minimum version of the TLS protocol supported by the web-server. You can select from the following:</p> <ul style="list-style-type: none"> tlsv10 – Configures the version to TLS 1.0. <p>Note: tlsv10 is not supported in enhanced secure mode.</p> <ul style="list-style-type: none"> tlsv11 – Configures the version to TLS 1.1. tlsv12 – Configures the version to TLS 1.2 <p>The default is tlsv12.</p>

Use GREP CLI Show Command Filters

Use the following GREP filters to output only the command lines specified by the filter.

Procedure

- Count the number of lines in the output:

```
<CLI command> | count
```
- Display the output of a command starting from the first line that matches the given pattern:

```
<CLI command> | begin WORD<0-255> [field <number>] [ignore-case]
[header <number>]
```
- Display only the output lines that match the given pattern:

```
<CLI command> | include <pattern> [field <number>] [ignore-case]
[header <number>]
```
- Display only the output lines that do not match the given pattern:

```
<CLI command> | exclude <pattern> [field <number>] [ignore-case]
[header <number>]
```
- Temporarily disable pagination for the output of a CLI command:

```
<CLI command> | no-more
```

There is no prompt to continue or to quit when the lines of output exceed the terminal length.

6. Limit the output of a command to the first few lines:

```
<CLI command> | head [<number>]
```

If a number is not specified, the first 10 lines display.

7. Limit the output of a command to the last few lines:

```
<CLI command> | tail [<number>] [from-line <number>] [header <number>]
```

If a number is not specified, the last 10 lines display.

Example

```
Switch:1>enable
Switch:1#configure terminal
```

Count the number of lines in the output:

```
Switch1:#show vlan basic | count
Count: 17 lines
```

Display only the output lines that match the given pattern:

```
Switch:1(config)#show vlan basic | include byPort field 3 header 6
```

```
=====
Vlan Basic
=====
VLAN
ID  NAME          TYPE      MSTP
INST_ID PROTOCOLID  SUBNETADDR  SUBNETMASK  VRFID  ORIGIN
-----
1   Default      byPort    0           none       N/A       N/A         0        CONFIG
3   VLAN3        byPort    3           none       N/A       N/A         0        CONFIG
4   VLAN4        byPort    4           none       N/A       N/A         0        CONFIG
5   VLAN5        byPort    5           none       N/A       N/A         0        CONFIG
8   VLAN-8       byPort    8           none       N/A       N/A         0        CONFIG
9   VLAN-9       byPort    9           none       N/A       N/A         0        CONFIG
11  VLAN-11      byPort    11          none       N/A       N/A         0        CONFIG
12  VLAN-12      byPort    12          none       N/A       N/A         0        CONFIG
20  VLAN-20      byPort    0           none       N/A       N/A         0        CONFIG
```

```
Switch:1(config)#show vlan basic | include private field 3 header 6
```

```
=====
Vlan Basic
=====
VLAN
ID  NAME          TYPE      MSTP
INST_ID PROTOCOLID  SUBNETADDR  SUBNETMASK  VRFID  ORIGIN
-----
6   VLAN6        private   40          none       N/A       N/A         0        CONFIG
7   VLAN7        private   41          none       N/A       N/A         0        CONFIG
```

Display only the output lines that do not match the given pattern:

```
Switch:1(config)#show vlan basic | exclude private field 3 header 6
```

```
=====
Vlan Basic
=====
VLAN
ID  NAME          TYPE      MSTP
INST_ID PROTOCOLID  SUBNETADDR  SUBNETMASK  VRFID  ORIGIN
-----
---
1   Default      byPort    0           none       N/A       N/A         0        CONFIG
3   VLAN3        byPort    3           none       N/A       N/A         0        CONFIG
4   VLAN4        byPort    4           none       N/A       N/A         0        CONFIG
```

```

5      VLAN5      byPort      5      none      N/A      N/A      0      CONFIG
8      VLAN-8     byPort      8      none      N/A      N/A      0      CONFIG
9      VLAN-9     byPort      9      none      N/A      N/A      0      CONFIG
11     VLAN-11    byPort      11     none      N/A      N/A      0      CONFIG
12     VLAN-12    byPort      12     none      N/A      N/A      0      CONFIG
20     VLAN-20    byPort      0      none      N/A      N/A      0      CONFIG

```

```
Switch:1(config)#show vlan basic | exclude byPort field 3 header 6
```

```

=====
Vlan Basic
=====
VLAN
ID   NAME           TYPE           MSTP
INST_ID PROTOCOLID    SUBNETADDR    SUBNETMASK    VRFID  ORIGIN
-----
6     VLAN6         private        40            none        N/A      N/A          0      CONFIG
7     VLAN7         private        41            none        N/A      N/A          0      CONFIG

```

Display the output of a command starting from the first line that matches the given pattern:

```
Switch:1(config)#show vlan basic | begin 8 header 6
```

```

=====
Vlan Basic
=====
VLAN
ID   NAME           TYPE           MSTP
INST_ID PROTOCOLID    SUBNETADDR    SUBNETMASK    VRFID  ORIGIN
-----
8     VLAN-8         byPort         8             none        N/A      N/A          0      CONFIG
9     VLAN-9         byPort         9             none        N/A      N/A          0      CONFIG
11    VLAN-11        byPort         11            none        N/A      N/A          0      CONFIG
12    VLAN-12        byPort         12            none        N/A      N/A          0      CONFIG
20    VLAN-20        byPort         0             none        N/A      N/A          0      CONFIG

```

Display the entire output of the command:

```
Switch:1(config)#show vlan basic | no-more
```

```

=====
Vlan Basic
=====
VLAN
ID   NAME           TYPE           MSTP
INST_ID PROTOCOLID    SUBNETADDR    SUBNETMASK    VRFID  ORIGIN
-----
1     Default        byPort         0             none        N/A      N/A          0      CONFIG
3     VLAN3          byPort         3             none        N/A      N/A          0      CONFIG
4     VLAN4          byPort         4             none        N/A      N/A          0      CONFIG
5     VLAN5          byPort         5             none        N/A      N/A          0      CONFIG
6     VLAN6          private        40            none        N/A      N/A          0      CONFIG
7     VLAN7          private        41            none        N/A      N/A          0      CONFIG
8     VLAN-8         byPort         8             none        N/A      N/A          0      CONFIG
9     VLAN-9         byPort         9             none        N/A      N/A          0      CONFIG
11    VLAN-11        byPort         11            none        N/A      N/A          0      CONFIG
12    VLAN-12        byPort         12            none        N/A      N/A          0      CONFIG
20    VLAN-20        byPort         0             none        N/A      N/A          0      CONFIG

```

Display only the first few lines of output:

```
Switch:1(config)#show vlan basic | head 9
```

```
=====
```

Vlan Basic								
VLAN ID	NAME	TYPE	MSTP INST_ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRfid	ORIGIN
1	Default	byPort	0	none	N/A	N/A	0	CONFIG
3	VLAN3	byPort	3	none	N/A	N/A	0	CONFIG

Display only the last few lines of output:

```
Switch:1(config)#show vlan basic | tail 8 header 6
```

Vlan Basic								
VLAN ID	NAME	TYPE	MSTP INST_ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRfid	ORIGIN
8	VLAN-8	byPort	8	none	N/A	N/A	0	CONFIG
9	VLAN-9	byPort	9	none	N/A	N/A	0	CONFIG
11	VLAN-11	byPort	11	none	N/A	N/A	0	CONFIG
12	VLAN-12	byPort	12	none	N/A	N/A	0	CONFIG
20	VLAN-20	byPort	0	none	N/A	N/A	0	CONFIG

```
Switch:1(config)#show vlan basic | tail from-line 15 header 6
```

Vlan Basic								
VLAN ID	NAME	TYPE	MSTP INST_ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRfid	ORIGIN
9	VLAN-9	byPort	9	none	N/A	N/A	0	CONFIG
11	VLAN-11	byPort	11	none	N/A	N/A	0	CONFIG
12	VLAN-12	byPort	12	none	N/A	N/A	0	CONFIG
20	VLAN-20	byPort	0	none	N/A	N/A	0	CONFIG

Variable definitions

The GREP filters use the following parameters:

Parameter	Description
<i>field<number></i>	Specifies the field in each line to match against the pattern. Fields are separated by white spaces and are counted starting with 1 for the left-most field. If the output is formatted as a table, whitespaces are not counted as fields.
<i>from-line <number></i>	Specifies the remaining output starting with a given line.
<i>head<number></i>	Specifies the number of lines to keep from the beginning of the output.
<i>header<number></i>	Specifies a number of lines from the start of the output to display unchanged before trying to match the pattern. This parameter is useful to keep the header of a table intact. This filter skips the header lines.
<i>ignore-case</i>	Specifies letters to match in the pattern regardless of case.

Parameter	Description
<i><number></i>	Specifies the number of lines of output to keep, either from the beginning of the output or from the end of the output.
<i><pattern></i>	Specifies the regular expression to match against each line of output. Use quotations if the parameter contains spaces.

Return a port to Default Values

About This Task

You can use the **default all** command to return all parameters configured on a port to their default values instead of using the individual default operator with each command on the port and optionally enable Auto-sense.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. (Optional) Return all parameters configured on a port to their default values:

```
default all [-y]
```

3. (Optional) Return all parameters configured on a port to their default values and enable Auto-sense:

```
default all auto-sense [-y]
```

4. Save the configuration:

```
save config
```

Variable Definitions

The following table defines parameters for the **default all** command.

Variable	Value
<i>auto-sense</i>	Clears the existing configuration and enables Auto-sense on the port.
<i>-y</i>	Suppresses the confirmation message. If you omit this parameter, you must confirm the action.



Enterprise Device Manager

[Enterprise Device Manager Fundamentals](#) on page 203

[EDM Interface Procedures](#) on page 214

[File Management in EDM](#) on page 221

Table 20: Enterprise Device Manager product support

Feature	Product	Release introduced
Enterprise Device Manager (EDM)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Read-Only user for EDM	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Enterprise Device Manager Fundamentals

This section details Enterprise Device Manager (EDM).

EDM is a web-based graphical user interface (GUI) you can use to configure a single switch. EDM runs from the switch and you can access it from a web browser.

Supported Browsers

Use the following browser versions to access Enterprise Device Manager (EDM):

- Microsoft Edge 146
- Google Chrome 146
- Mozilla Firefox 149
- Safari 26.2



Important

For optimal performance, use Mozilla Firefox or Google Chrome.

Default User Names and Passwords for EDM

The following table provides the default user names and passwords for EDM:

Table 21: Default EDM access

Access level	Default user name	Default password
RWA	admin	password
RO	user	12345678



Important

The default passwords and community strings are documented and well known. If the admin and user accounts for web access use default passwords, the system prompts you to change them when you enable the web management interface. For more information, see [Configure the Web Server](#) on page 195 and [Enable the Web Management Interface](#) on page 303.

For information about creating CLI accounts for each user role on the switch, see [Multiple CLI Users for Each Role](#) on page 263.

Device Physical View

After you access EDM, the system displays a real-time physical view of the front panel of the device on the **Device Physical View** tab in the content pane. From the front panel view, you can view fault, configuration, and performance information for the device or a single port.

You can use the device view to determine the operating status of the various ports in your hardware configuration. You can also use the device view to perform management tasks on specific objects. In the device view, you can select a port or the entire chassis. EDM outlines the selected object in yellow.

The conventions on the device view are similar to the actual device appearance. The port LEDs and the ports are color-coded to provide status. Green indicates the module or port is up and running, red indicates the module or port is disabled, dark pink indicates a protocol is down, and amber indicates an enabled port that is not connected to anything. For information about LED behavior, see your hardware documentation.

**Note**

While the EDM view typically mirrors the physical appearance of most devices, the 7830 Series is an exception by design. Although the physical hardware includes only one LED per port, EDM intentionally displays four LEDs per port to represent the different channels of a channelized port. This design choice enhances visibility and management of individual port channels within the EDM interface on the 7830 Series.

EDM Window

The following list identifies the different sections of the EDM window:

- **Navigation pane**—Located on the left side of the window, the navigation pane displays all the available command tabs in a tree format. A row of buttons at the top of the navigation pane provides a quick method to perform common functions.
- **Content pane**—Located on the right side of the window, the content pane displays the tabs and dialog boxes where you can view or configure parameters on the switch.
- **Menu bar**—Located at the top of the content pane, the menu bar shows the most recently accessed primary tabs and their respective secondary tabs.
- **Toolbar**—Located just below the menu bar, the toolbar provides quick access to the most common operational commands such as Apply, Refresh, and Help.

The following figure shows an example of two tabs open in the content pane of the EDM window.

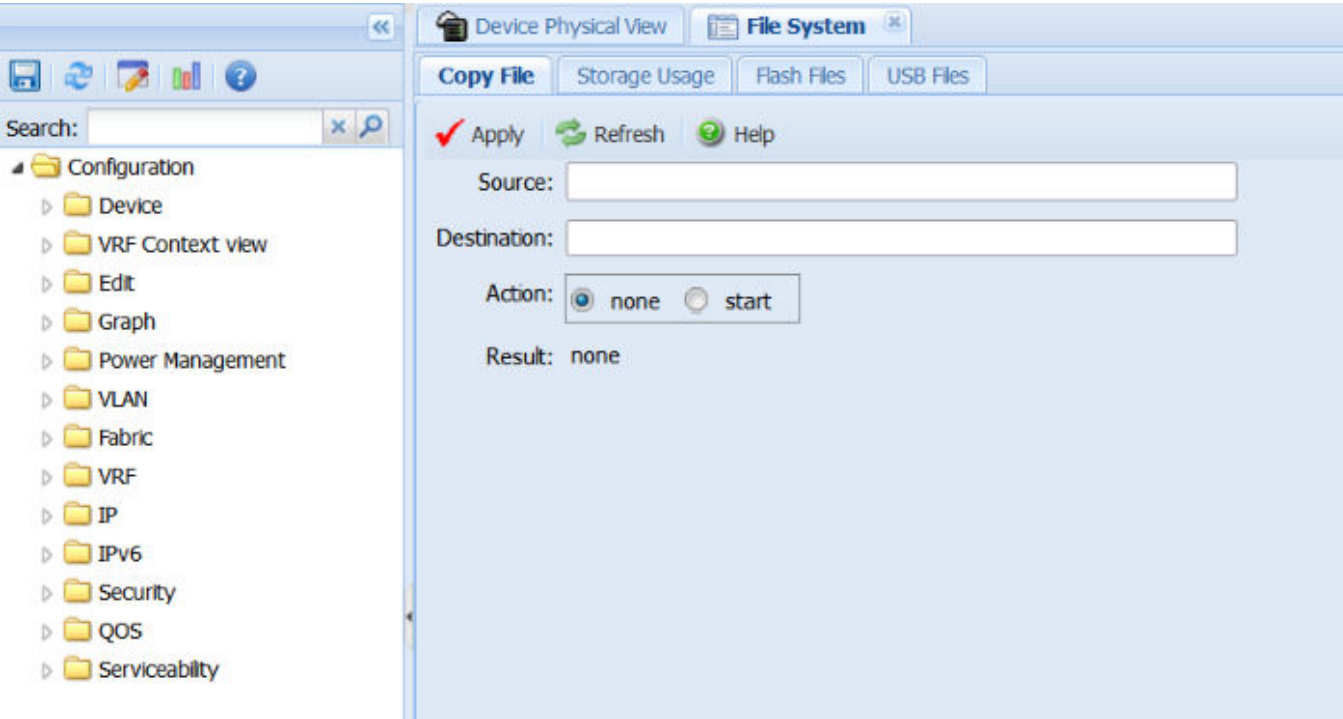


Figure 12: EDM window

Navigation Pane

You can use the navigation pane to see what commands are available and to quickly browse through the command hierarchy. A row of buttons at the top of the navigation pane provides a quick method to perform common functions.



Note



For module-based chassis, menu options related to a specific module are activated only after you install and select the required module.

The following table describes the buttons that display at the top of the navigation pane.

Table 22: Navigation pane buttons

Button	Name	Description
	Save Config	Saves the running configuration.
	Refresh Status	Refreshes the Device Physical View.
	Edit	Edits the selected item in the Device Physical View.
	Graph	Opens the graph options for the selected item in the Device Physical View.
	Help Setup Guide	Opens instructions about how to install the Help files and configure EDM to use the Help files.

Table 22: Navigation pane buttons (continued)

Button	Name	Description
	Help Documentation	Displays links to the following information: <ul style="list-style-type: none">• Software Release Notes• Documentation collections• Hardware and Software Compatibility Matrices• Documentation for Extreme Networks optics• RESTCONF Reference Documentation• The support portal for Software, MIB, Vulnerability/CVE and Field Notices• GitHub information for GNS3 images
	Close all tabs	Closes all open tabs in the EDM session.

Expand a folder by selecting the directional arrow next to the folder name. Some folders have sub-folders such as the Edit folder, which has the Port, NTP, and other sub-folders.

Within each folder and sub-folder, there are numerous options, which provide access to tabs. To open an option, select it. The selected tab displays in the menu bar and

opens in the content pane. The following table describes the top-level folders in the navigation pane.



Note

Feature support differences affect folder availability. For information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Table 23: Navigation Pane Folders

Menu	Description
Device	<p>Use the Device menu to refresh and update device information or enable polling.</p> <ul style="list-style-type: none"> • Preference Setting — Enable polling or hot swap detection. Configure the frequency to poll the device. • Refresh Status — Use this option to refresh the device view. • Rediscover Device — Use this to trigger a rediscovery to update all of the device information.
VRF Context view	<p>Use the VRF Context view to switch to another VRF context when you use the embedded EDM. GlobalRouter is the default view at log in. You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view. You can open only five tabs for each EDM session.</p>
Edit	<p>Use the Edit menu to view and configure parameters for the chassis hardware or for the currently selected hardware component, including one or more ports. You can also use the Edit menu to perform the following tasks:</p> <ul style="list-style-type: none"> • check and configure ports, including but not limited to, the management ports and the internal Extreme Integrated Application Hosting ports, on the device • change the configuration of many features, including but not limited to, the file system, NTP, OVSDDB, BFD, SMTP, Link-state tracking, Endpoint Tracking, VTEP, Mgmt Instance, and SNMPv3
Graph	<p>Use the Graph menu to view and configure EDM statistics and to produce graphs of the chassis or port statistics.</p>

Table 23: Navigation Pane Folders (continued)

Menu	Description
Power Management	Use the Power Management menu to view and configure Energy Saver.
VLAN	Use the VLAN menu to view and configure the following: <ul style="list-style-type: none">• VLANs• Spanning Tree• MLT/LACP• SMLT• SLPP
Fabric	Use the Fabric menu to view and configure the following: <ul style="list-style-type: none">• IS-IS• Shortest Path Bridging MAC (SPBM)• ISID• Fabric Attach• DvR• Multi-area SPB• Stats
VRF	Use the VRF menu to view and create VRFs.

Table 23: Navigation Pane Folders (continued)

Menu	Description
IP	<p>Use the IP menu to view and configure IP routing functions for the system, including but not limited to the following:</p> <ul style="list-style-type: none">• IP-VPN• IP-MVPN• IP• TCP/UDP• OSPF• RIP• VRRP• RSMPT• BGP• Multicast• MSDP• IGMP• IPFIX• PIM• SPB-PIM-GW• DHCP Relay• DHCP Server• DHCP Snooping• ARP Inspection• Source Guard• UDP Forwarding• IS-IS• Policy• BFD

Table 23: Navigation Pane Folders (continued)

Menu	Description
IPv6	<p>Use the IPv6 menu to view and configure IPv6 routing functions, including but not limited to the following:</p> <ul style="list-style-type: none"> • IPv6 • IPv6 - VPN • TCP/UDP • Tunnel • OSPFv3 • VRRP • BGP+ • RSMLT • DHCP Relay • Policy • FHS • IS-IS • IPv6 RIPng • IPv6 PIM • IPv6 MLD • IPv6 Mroute • IPv6 BFD
Security	<p>Use the Security menu to view and configure the following:</p> <ul style="list-style-type: none"> • General • Certificates • Access Policies • RADIUS • RADIUS CoA • SSH • IPSec • TACACS+ • IKE • EAPoL • Advanced Filters such as ACE and ACLs • MACsec

Table 23: Navigation Pane Folders (continued)

Menu	Description
QOS	Use the QOS menu to view and configure the following: <ul style="list-style-type: none"> • Mapping Tables • Port QoS Config • CoS Queue Stats • Queue Profile
Serviceability	Use the Serviceability menu to run diagnostics, and to enable, configure, or view features including but not limited to the following: <ul style="list-style-type: none"> • RMON • Sflow • Application Telemetry • RESTCONF • Virtual Service • ExtremeCloud IQ Agent • Node Alias

Menu Bar

The menu bar is above the content pane and consists of two rows of tabs.

- The top row displays the tabs you can open through the navigation pane. The system displays these primary tabs in the sequence in which you open them.
- After you select a primary tab, the secondary tabs associated with it display in the bottom row. Select a secondary tab to display it in the content pane.

In both the top and bottom rows of the menu bar, if the number of tabs exceeds the viewable space, the system displays left- and right-pointing arrows. Select an arrow to scroll to the required tab.

To reduce the number of tabs on the top row, you can select the X on the right corner of a tab to remove it from the row. The following figure shows a sample menu bar.

**Figure 13: Menu bar**

Toolbar

The toolbar buttons provide quick access to commonly used operational commands. The system displays the buttons that vary depending on the tab you select. However,

the Apply, Refresh, and Help buttons are on almost every screen. Other common buttons are Insert and Delete. The following list details the common toolbar buttons.

- **Apply**—Use this to apply all edits that you make.
- **Refresh**—Use this to refresh all data on the screen.
- **Help**—Use this to display online help that is context sensitive to the current tab.
- **Insert**—Use this to display a secondary dialog box related to the selected tab. After you edit the configurable parameters, select **Insert** in the dialog box. This causes a new entry to display in the dialog box of the selected tab.
- **Delete**—Use this to delete a selected entry.

The following figure shows a sample toolbar.




Figure 14: Toolbar

Content Pane

The content pane is the main area on the right side of the window that displays the configuration tabs and dialog boxes. Use the content pane to view or configure parameters on the switch.



Note

You can view valid ranges for all configurable parameters on EDM tabs.

The following figure is a sample that shows the content pane for the Port 1/3 General, Interface tab. If you want to compare the information in two tabs, you can undock one, then open another tab. For more information about undocking a tab, see [Undock and dock tabs](#) on page 219.




Figure 15: Content pane

EDM user session extension

If the EDM user session remains unused for a duration of ten minutes, the system displays the following message:

Your session will expire in about 5 minute(s). Would you like to extend the session?

If you do not respond, EDM automatically ends the session with the following message:
Your session has expired.

You can log on again if you want to continue to use EDM.

EDM Interface Procedures

This section contains procedures for starting and using Enterprise Device Manager (EDM). The software is built-in to the switch.

Configure the Web Management Interface

Before You Begin

- Enable the web server.

About This Task

Configure the web management interface to change the user names and passwords for management access to the switch using a web browser.

HTTP, FTP, and TFTP server supports both IPv4 and IPv6 addresses, with no difference in functionality or configuration.

You can also use the CLI interface for creating users.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Web** tab.
4. Complete **WebRWAUserName** and **WebRWAUserPassword** to specify the user name and password for access to the web interface.
This user will have full permission.
5. To enable the RO user for the web server, select **WebROEnable**.
6. Complete **WebROUserName** and **WebROUserPassword** to specify the user name and password for access to the web interface.
This user will have read only permission.
7. Select **Apply**.

Web Field Descriptions

Use the data in the following table to use the **Web** tab.

Name	Description
WebRWAUserName	Specifies the RWA user name. The default is admin.
WebRWAUserPassword	Specifies the password for the RWA user account. The default is password.
WebROEnable	Enables the web server read-only (RO) user, which is disabled by default.
WebEncryptionType	Specifies the ciphers for preset version of TLS for the web server.
WebCertSubjectName	Specifies the digital certificate subject Name used as identity certificate in the web server.
WebCertCAName	Specifies the digital certificate CA trustpoint name used for the certificate in the web server.
WebROUserName	Specifies the RO user name. The default is user.
WebROUserPassword	Specifies the password for the RO user account. The default is 12345678.
MinimumPasswordLength	Configures the minimum password length. By default, the minimum password length is 8 characters.
HttpPort	Specifies the HTTP port for web access. The default value is 80.
HttpsPort	Specifies the HTTPS port for web access. The default value is 443.
SecureOnly	Controls whether the secure-only option is enabled. The default is enabled.
InactivityTimeout	Specifies the idle time (in seconds) to wait before the EDM login session expires. The default value is 900 seconds (15 minutes).
TlsMinimumVersion	Configures the minimum version of the TLS protocol supported by the web-server. You can select from the following options: <ul style="list-style-type: none"> • tlsv10 – Configures the version to TLS 1.0. • tlsv11 – Configures the version to TLS 1.1. • tlsv12 – Configures the version to TLS 1.2 The default TLS minimum version is tlsv12.
InUseCertType	Specifies whether the type of certificate in use is either user installed or self assigned.

Name	Description
SSLRenegotiation	Specifies whether SSL renegotiation is enabled or disabled. The default is enabled.
HelpTftp/Ftp_SourceDir	Configures the TFTP or FTP directory for Help files, in one of the following formats: a.b.c.d:/ peer:/ [<dir>]. The path can use 0–256 characters. The following example paths illustrate the correct format: <ul style="list-style-type: none"> • 192.0.2.1:/Help • 192.0.2.1:/
DefaultDisplayRows	Configures the number of rows each page displays. The default is 30.
LastChange	Shows the last web-browser initiated configuration change.
NumHits	Shows the number of hits to the web server.
NumAccessChecks	Shows the number of access checks performed by the web server.
NumAccessBlocks	Shows the number of access attempts blocked by the web server.
LastHostAccessBlockedAddressType	Shows the address type, either IPv4 or IPv6, of the last host access blocked by the web server.
LastHostAccessBlockedAddress	Shows the IP address of the last host access blocked by the web server.
NumRxErrors	Shows the number of receive errors the web server encounters.
NumTxErrors	Shows the number of transmit errors the web server encounters.
NumSetRequest	Shows the number of set-requests sent to the web server.

Use the chassis shortcut menu

About This Task

Perform the following procedure to display the chassis shortcut menu.

Procedure

1. In the Device Physical View, select the chassis.
2. Right-click the chassis.

Chassis Shortcut Menu Field Descriptions

Use the data in the following table to use the shortcut menu.

Name	Description
Edit	Edits chassis parameters.
Graph	Graphs chassis statistics.
Refresh Status	Refreshes the status of the chassis and MDAs.
Refresh Port Tooltips	Refreshes the port tooltip data of the system. The port tooltip data contains the following variables: Slot/Port, PortName, and PortOperSpeed.

Use the port shortcut menu

About This Task

Perform this procedure to display the port shortcut menu.

Procedure

1. In the Device Physical View, select a port.
2. Right-click the selected port.

Port Shortcut Menu Field Descriptions

Use the data in the following table to use the port shortcut menu.

Name	Description
Edit General	Configures the general options for the port.
Edit IP	Configures the IP options for the port.
Edit IPv6	Configures the IPv6 options for the port.
Channelization Enable	Enables channelization for the port.
Channelization Disable	Disables channelization for the port.
Graph	Displays the statistics for the port.
Enable	Enables the port.
Disable	Disables the port.

NEW! Use the VIM shortcut menu

About This Task

Perform this procedure to display the VIM shortcut menu.

Procedure

1. In the Device Physical View, select the VIM.
2. Right-click the VIM.

NEW! VIM Shortcut Menu Field Descriptions

Use the data in the following table to use the VIM shortcut menu.

Name	Description
Edit	Configures the general options for the VIM.
Enable	Enables the power in the VIM slot.
Disable	Disables the power in the VIM slot.

Use a table-based tab

About This Task

Change an existing configuration using a table-based tab. You cannot edit grey-shaded fields in the table. The following procedure is an illustration on how to use a table-based tab.



Note


You can expand the appropriate folders for any feature you configure and select a table-based tab.

Procedure

1. In the Device Physical View, select multiple ports.
2. In the navigation pane, expand **Configuration > Edit > Port > General**.
3. Select the **VLAN** tab.
The system displays a table-based tab with the VLAN information.
4. Select a table-based tab.
5. Double-click a white-shaded field to edit the value.
6. Select the arrow in the list field to view the options, and then choose the appropriate value.

Index	PerformTagging	VlanIdList	DiscardTaggedFrames	DiscardUntaggedFrames	UntagDefaultVlan	DefaultVlanId	LoopId
2-9	false		false	false	false	0	false
224	false		false	false	false	0	false
226	false		false	false	false	0	false
228	false		false	false	false	0	false

- Double-click a text-entry field and edit the value.



- Select **Apply** to save the configuration changes.

Monitor Multiple Ports and Configuration Support

About This Task

You can monitor or apply the same configuration changes to more than one port by using the multiple port selection function. You can use the standard menu or the shortcut menu to edit the configuration settings for multiple ports.



Tip

A selected port shows a yellow outline around the port.

Procedure

- Select the **Device Physical View** tab.
- To select multiple ports, press **Ctrl** (Control), and then select the required ports.

Open Folders and Tabs

About This Task

Perform this procedure to navigate in EDM.

Procedure

- In the navigation pane, expand **Configuration**.
- Select a subfolder to expand it and see the list of options, for example, the **VLAN** folder.
- In a folder or subfolder menu, select an option to open the related tabs.

Undock and dock tabs

About This Task

Perform this procedure to undock a tab. You can undock tabs to have more than one tab visible at a time.

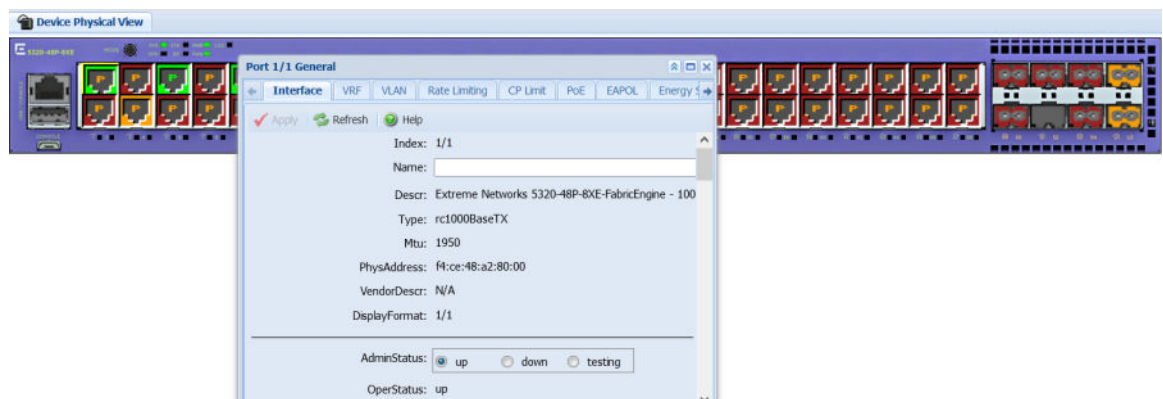
Procedure

1. In the navigation pane, select a tab.
2. In the menu bar, select and drag a tab to undock it.
3. In the top right corner of the tab, select **pages** to dock the tab.

Example of Undocking and Docking Tabs

Procedure


1. Select the **Device Physical View** tab.
2. In the Device Physical View, select a port. In this example, right-click port 1.
3. In the **Port** shortcut menu, select **Edit General**.
4. Select and drag the **Port 1/1 General** tab wherever you want on the screen as shown in the following figure.



5. To reposition the tab anywhere on the screen, Select and drag the title bar.
6. To manipulate the tab, Select on the buttons in the top-right of the dialog box.



7. Select the up arrowhead to minimize the tab as shown in the following figure.



8. Select the down arrowhead to restore the tab to its original size.
9. Select the pages to dock the tab back into the menu bar.
10. Select the X to close the tab.

Install EDM Help Files

While the EDM GUI is bundled with the switch software, the associated EDM help files are not. To access the help files from the EDM GUI, you must install the EDM help files on a TFTP or FTP server in the network.

Use the following procedure to install the EDM help files on a TFTP or FTP server, and configure EDM to use the help files

Before You Begin

If you use an FTP server to store the help files, ensure that you configure the switch with the host user name and password.

Procedure

1. Download the EDM help file.
2. On a TFTP or FTP server reachable from the switch, create a directory called **Help**.



Tip

You can name the directory anything that helps you remember its purpose.

3. Unzip the EDM help zip file into the directory created in the preceding step.
4. In the EDM navigation pane, expand **Configuration > Security > Control Path**.
5. Select **General**.
6. Select **Web**.
7. In **HelpTftp/Ftp_SourceDir**, type the path, including the IP address, to the help files.
Example: 192.0.2.15:/home/Help/.
8. Select **Apply**.

File Management in EDM

This section contains procedures for managing files with Enterprise Device Manager (EDM).

Use the File System tab to perform the following tasks:

- Copy a file.
- Check the amount of memory used and the number of files stored in the internal flash memory.
- Verify the name, size, and storage date of each file present in the internal flash memory.
- Display USB file information.

Copy a File

About This Task

Copy files on the internal flash.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **File System**.
3. Select the **Copy File** tab.
4. Edit the fields as required.
5. Select **Apply**.

Copy File *Field Descriptions*

Use the data in the following table to use the **Copy File** tab.

Name	Description
Source	Identifies the device and file name to copy. You must specify the full path and file name, for example, <deviceip-ftp server>:/<file name> Note: For certain switches in enhanced secure mode, sensitive files and paths are protected.
Destination	Identifies the location to which to copy the source file with the file name, for example, /intflash/<file name>. Note: For certain switches in enhanced secure mode, sensitive files and paths are protected.
Action	Starts or stops the copy process.
Result	Specifies the result of the copy process: <ul style="list-style-type: none"> • none • inProgress • success • fail • invalidSource • invalidDestination • outOfMemory • outOfSpace • fileNotFound

Display Storage Use

About This Task

Display the amount of memory used, memory available, and the number of files for internal flash memory.



Note

5320 Series and 5420 Series support 512 MB of flash memory but only 390 MB is available for use.

Procedure

1. In the navigation pane, expand **Configuration** > **Edit**.
2. Select **File System**.
3. Select the **Storage usage** tab.

Storage Usage Field Descriptions

Use the data in the following table to use the **Storage Usage** tab.

Name	Description
IntflashBytesUsed	Specifies the number of bytes used in internal flash memory.
IntflashBytesFree	Specifies the number of bytes available for use in internal flash memory.
IntflashNumFiles	Specifies the number of files in internal flash memory.
UsbBytesUsed	Specifies the number of bytes used in USB device.
UsbBytesFree	Specifies the number of bytes available for use in USB device.
UsbNumFiles	Specifies the number of files in USB device.

Display Internal Flash File Information

About This Task

Display information about the files in internal flash memory on this device.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **File System**.
3. Select the **Flash Files** tab.

Flash Files Field Descriptions

Use the data in the following table to use the **Flash Files** tab.

Name	Description
Slot	Specifies the slot number.
Name	Specifies the directory name of the file.
Date	Specifies the creation or modification date of the file.
Size	Specifies the size of the file.

Display USB File Information

About This Task

Display information about the files on a USB device to view general file information.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **File System**.
3. Select the **USB Files** tab.

USB Files *Field Descriptions*

Use the data in the following table to use the **USB Files** tab.

Name	Description
Slot	Specifies the slot number of the device.
Name	Specifies the directory name of the file.
Date	Specifies the creation or modification date of the file.
Size	Specifies the size of the file.



Image Management

- [Image Upgrades](#) on page 225
- [Image Naming Conventions](#) on page 226
- [Interfaces](#) on page 227
- [File Storage Options](#) on page 227
- [Boot Loader Image on Universal Hardware](#) on page 228
- [Before You Upgrade](#) on page 229
- [Calculate and Verify the MD5 Checksum for a File on the Switch](#) on page 230
- [Calculate and Verify the MD5 Checksum for a File on a Client Workstation](#) on page 231
- [Save the Configuration](#) on page 231
- [Upgrade the Software](#) on page 233
- [Verify the Upgrade](#) on page 237
- [Commit an Upgrade](#) on page 237
- [Downgrade the Software](#) on page 238
- [Remove a Software Build](#) on page 241
- [Update the Complex Programmable Logic Device \(CPLD\) Image](#) on page 241
- [LRM/MACsec Adapter](#) on page 243
- [Upgrade the LRM/MACsec Adapter PHY Firmware](#) on page 244

The topics in this section detail what you must know to manage the software image on the switch.

Image Upgrades

Install new software upgrades to add functionality to the switch. Major and minor upgrades are released depending on how many features the upgrade adds or modifies.

Upgrade Time Requirements

Image upgrades take less than 30 minutes to complete. The switch continues to operate during the image download process. A service interruption occurs during the installation and subsequent reset of the device. The system returns to an operational state after a successful installation of the new software and device reset.

Before You Upgrade the Software Image

Before you upgrade the switch, ensure that you read the entire upgrading procedure.



Important

You must keep a copy of the previous configuration file in case you need to return to the previous version. The upgrade process automatically converts, but does not save, the existing configuration file to a format that is compatible with the new software release. The new configuration file may not be backward compatible.

Image Naming Conventions

The switch software use a standardized dot notation format.

Software Images

Software image names use the following number formats to identify release and maintenance values:

Product Name.Major Release.Minor Release.Maintenance Release.Maintenance Release Update.voss

For example, the image file name `5320.8.6.1.0.voss` denotes a software image for the 5320 Series product with a major release version of 8, a minor release version of 6, a maintenance release version of 1 and a maintenance release update version of 0. VOSS is the file extension.

Firmware Update And Verification With Digital Signed Certificate

Software images are cryptographic signed. Code signing is the process of digitally signing executables and scripts to confirm the software author and guarantee that the code has not been altered or corrupted since it was signed. This process employs the use of a cryptographic hash to validate authenticity and integrity.

The **show software** command displays information about the software image:

```
Switch:1#show software

=====
                        software releases in /intflash/release/
=====
8.8.0.0.GA                (Backup Release)                (Signed Release)
9.0.0.0.GA                (Primary Release)                (Signed Release)
-----
Auto Commit       : enabled
Commit Timeout   : 7 minutes
```

Operational Considerations

The following section describes operational considerations:

- You not required to provide additional input.

- You can use unsigned images; however, this is not recommended. To use an unsigned image, downgrade to a pre-VOSS 8.5 software image, and then load a debug image.
- You cannot enter Enhanced Secure Mode with an unsigned image. Enhanced Secure Mode requires a signed image.

Interfaces

You can apply upgrades to the switch using the Command Line Interface (CLI).

For more information about CLI, see [Command Line Interface](#) on page 178.

File Storage Options

This section details what you must know about the internal boot and system flash memory and Universal Serial Bus (USB) mass-storage device, which you can use to store the files that start and operate the switch.

The switch file system uses long file names.

Internal Flash

The switch has two internal flash memory devices: the boot flash memory and the system flash memory.

Boot flash memory is split into two banks that each contain a different copy of the boot image files. Only the Image Management feature can make changes to the boot flash.

The system flash memory stores configuration files, runtime images, the system log, and other files. You can access files on the internal flash through the `/intflash/` folder.

During the software upgrade process on 4220 Series, 5320 Series, and 5420 Series the switch automatically creates two archives for logs and message files:

- `archive.log.tgz`—includes all `log*` files present in `/intflash/shared`
- `archive.messages.tgz`—includes all `messages*` files present in `/inflash/var/log`

USB Device

The switch can use a USB device for additional storage or configuration files, release images, and other files. The USB device provides a convenient, removable mechanical to copy files between a computer and a switch, or between switches. In cases where network connectivity has not yet been established, or network file transfer is not feasible, you can use a USB device to upgrade the configuration and image files on the switch.

File Transfer Protocol

You can use File Transfer Protocol (FTP) to load the software directly to the switch, or to download the software to the internal flash memory or to an installed USB device.

The switch can act as an FTP server or client. If you enable the FTP daemon (ftpd), you can use a standards-based FTP client to connect to the switch by using the CLI log on parameters. Copy the files from the client to either the internal flash memory or USB device.

TFTP Block Number Rollover

The software uses block number rollover functionality to transfer files larger than 32 MB using TFTP. With this functionality, when the block number reaches 65,535, it resets to 0 while it maintains the standard block size of 512 bytes and allows the transfer to continue seamlessly.



Note

The best practice is to use FTP, SFTP, or USB devices to transfer the file with the new software release to the switch.

Boot Loader Image on Universal Hardware



Note

Not all products support the ability for the Fabric Engine software image to upgrade the uboot image. The output of the **show sys-info uboot** command displays only the installed image on these switches and no alternate image version.

On universal hardware products, new software activations automatically detect the uboot image, check if it is valid, and then compare the version of the uboot image with the version currently on the system. If the system started with the default uboot image, indicated by `Boot Version Used` in the **show sys-info uboot** command output, and the uboot image in the activated software release is newer, image synchronization performs a default uboot upgrade. After the default uboot upgrade is complete, the **show sys-info uboot** command output displays the uboot version from the system boot time, so it still shows the previous version, but it also indicates that the default uboot was upgraded and that a system reboot is required.

After you reboot the system for a default uboot upgrade, if a temporary default uboot upgrade file is present on the system, if the system started with the default uboot image, and if the default and alternate uboot versions are not the same, the alternate uboot is upgraded. You must restart the system for the alternate uboot upgrade to take effect.

Before You Upgrade

This section provides important feature impacts you need to understand before you upgrade the switch software.

DHCP Vendor Options Configuration Change

Prior to Release 9.1, if you configured two IP addresses when you defined the DHCP Server vendor options, only one IP address was included in Option 43. The workaround was to include two backslashes (\\) in the configuration.

For example:

```
ip dhcp-server vendor-option-data 1 value 192.0.2.2\\,192.0.2.5
```

This workaround is no longer required in Release 9.1. You must remove the backslashes (\\) in the configuration before you upgrade to Release 9.1 or later.

IS-IS Route Tagging Upgrade Implications



Caution

To use IS-IS Route Tagging on GRT IS-IS routes, you must also configure the metric-type as external. If you want to use IS-IS tags on GRT as internal routes, all Fabric nodes must be above a minimum software version. Any switch in the SPB Fabric that runs earlier software versions triggers an exception if you use metric type internal. To ensure this does not occur, if you attempt to configure a tag and the metric-type is not external, the switch reminds you to upgrade the software on all devices. You must ensure all devices in the network run the minimum required software.

Table 24: Minimum software required

NOS	Minimum software versions
Fabric Engine	8.10.6.1 and later 9.0.5.1 and later 9.1 and later
VOSS	8.10.6.1 and later 9.0.5.1 and later 9.1 and later
VSP 8600 Series	8.1.7 and later

Digital Certificate Upgrade Considerations

CA Field in Root and Intermediate Certificates

The system checks basic constraints prior to checking the certificates. Ensure the CA field is True for every root and intermediate certificate in the certificate chain, including

older certificates. If the CA field is False for the root certificate in the certificate chain, RESTCONF TLS server state is down.

Calculate and Verify the MD5 Checksum for a File on the Switch

Perform this procedure to verify that the software files are downloaded properly to the switch. The MD5 checksum for each release is available from the [Extreme Networks Support portal](#).

Before You Begin

- Download the MD5 checksum to an intermediate workstation or server where you can open and view the contents.
- Download the image file to the switch.

About This Task

Calculate and verify the MD5 checksum after you download software files.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the list of files:
`ls *.voss`
3. Calculate the MD5 checksum for the file:
`file-checksum md5 WORD<1-99>`
4. Compare the number generated for the file on the switch with the number that displays in the MD5 checksum on the workstation or server. Ensure that the MD5 checksum of the software suite matches the system output generated from calculating the MD5 checksum from the downloaded file.

Example

View the contents of the MD5 checksum on the workstation or server:

```
9e8e4356ec44661514cd142e933dca59 5420.8.4.0.0_edoc.tar
4602bf5adbc0c31c0e5f5d38209d7b87 5420.8.4.0.0_mib_sup.txt
25be1902e954e3140d1072562c926d4b 5420.8.4.0.0_mib.txt
09d06a9dbbfdaa07cabd2000556d5d2b 5420.8.4.0.0_mib.zip
8e23c6632b5b7a9ec134418348baebea 5420.8.4.0.0_oss-notice.html
d83cd61e0b115c83aca0c52c9dd5eb0e 5420.8.4.0.0.voss
eb86e5991d8ce2a69b21510b993e0727 restconf_yang.tgz
561a5b566c42b68cc4b1fc3b1e71cf70 VOSSv840_HELP_EDM_gzip.zip
```

Calculate the MD5 checksum for the file on the switch:

```
Switch:1>ls *.voss
Listing Directory /intflash:
-rw-r--r--  1 0      0      94195267 Dec 13 14:42  5420.8.4.0.0.voss
Switch:1>file-checksum md5 5420.8.4.0.0.voss
MD5 (5420.8.4.0.0.voss) = d83cd61e0b115c83aca0c52c9dd5eb0e
```

Variable Definitions

The following table defines parameters for the **file-checksum md5** command:

Variable	Value
<i>WORD<1-99></i>	Specifies the file name.

Calculate and Verify the MD5 Checksum for a File on a Client Workstation

About This Task

Perform this procedure on a Unix or Linux machine to verify that the software files downloaded properly. The MD5 checksum for each release is available from [Extreme Networks Support portal](#).

Procedure

1. Calculate the MD5 checksum of the downloaded file:

```
$ /usr/bin/md5sum <downloaded software file name>
```

Typically, downloaded software files are in the form of .voss files.

2. Verify the MD5 checksum of the software suite:
\$ more <md5-checksum output file>
3. Compare the output that displays on the screen. Ensure that the MD5 checksum of the software suite matches the system output generated from calculating the MD5 checksum from the downloaded file.

Example

Calculate the MD5 checksum of the downloaded file:

```
$ /usr/bin/md5sum 5420.8.4.0.0.voss
d83cd61e0b115c83aca0c52c9dd5eb0e 5420.8.4.0.0.voss
```

View the MD5 checksum of the software suite:

```
$ more 5420.8.4.0.0.md5
9e8e4356ec44661514cd142e933dca59 5420.8.4.0.0_edoc.tar
4602bf5adbc0c31c0e5f5d38209d7b87 5420.8.4.0.0_mib_sup.txt
25be1902e954e3140d1072562c926d4b 5420.8.4.0.0_mib.txt
09d06a9dbbfdaa07cabd2000556d5d2b 5420.8.4.0.0_mib.zip
8e23c6632b5b7a9ec134418348baebee 5420.8.4.0.0_oss-notice.html
d83cd61e0b115c83aca0c52c9dd5eb0e 5420.8.4.0.0.voss
eb86e5991d8ce2a69b21510b993e0727 restconf_yang.tgz
561a5b566c42b68cc4b1fc3b1e71cf70 VOSSv840_HELP_EDM_gzip.zip
```

Save the Configuration

Save the configuration

- When you make a change to the configuration.
- To create a backup configuration file before you upgrade the software on the switch.

After you change the configuration, you must save the changes on the device. Save the configuration to a file to retain the configuration settings.

Note that not all CLI commands are included in configuration files. Typical examples include, but are not limited to some operational and security-related commands.

**Note**

When loading large configuration files or large sections of a configuration file, avoid copying and pasting of the files into the console or terminal window as it can lead to the loss of configuration. You must either source the file or boot to the intended configuration file. Sourcing and booting allow for the debug and verification of the configuration file using the boot config flags. For more information about booting, sourcing, and debugging or verification using boot flags, see *Fabric Engine Command References*.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Save the running configuration:

```
save config [backup WORD<1-99>] [file WORD<1-99>] [verbose]
```

Example

Save the configuration to the default location:

```
Switch:1>enable  
Switch:1#save config
```

Identify the file as a backup file and designate a location to save the file:

```
Switch:1#save config backup /usb/PreUpgradeBackup.cfg
```

Variable Definitions

The following table defines parameters for the **save config** command.

Variable	Value
<i>backup</i> WORD<1-99>	Saves the specified file name and identifies the file as a backup file. WORD<1-99> uses one of the following format: <ul style="list-style-type: none"> • a.b.c.d:<file> • /intflash/<file> • /usb/<file> The file name, including the directory structure, can include up to 99 characters.
<i>file</i> WORD<1-99>	Specifies the file name in one of the following format: <ul style="list-style-type: none"> • a.b.c.d:<file> • /intflash/<file> • /usb/<file> The file name, including the directory structure, can include up to 99 characters.
<i>verbose</i>	Saves the default and current configuration. If you omit this parameter, the command saves only parameters you change.

Upgrade the Software



Important

Upgrades from some releases require release-specific steps. For more information, see *Fabric Engine Release Notes*.

Perform this procedure to upgrade the software on the switch. This procedure shows how to upgrade the software using the internal flash memory as the file storage location.

You can store up to six software releases on some switches. If you have six releases already stored on the switch, you are prompted to remove one release before you can proceed to add and activate a new software release.

On the 4220 Series, 5320 Series, and 5420 Series, you can store a maximum of two software releases. If you attempt to add a third software release, the software displays a confirmation to overwrite the non-primary release. You can also use the **software add <filename> -y** to bypass the confirmation question and automatically overwrite the non-primary release.

For information about how to remove a software release, see [Remove a Software Build](#) on page 241.

Before You Begin

- To obtain the new software, go to the [Extreme Networks Support portal](#). You need a valid user or site ID and password.
- Back up the configuration files.
- Use an FTP or SFTP application or USB device to transfer the file with the new software release to the switch.
- Ensure that you have not configured a VLAN above 4059. If you have, you must port all configuration on this VLAN to another VLAN, before you begin the upgrade.



Caution

Only VLAN range 2 to 4059 is supported. All configuration on a higher numbered VLAN from earlier releases will be lost after the upgrade.



Note

Software upgrade configurations are case-sensitive.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. If you are using the USB port to transfer files, go to the next step. If you are using FTP or SFTP to download the files, start the FTP daemon on the switch and enable the `ftpd` flag for FTP or `sshd` flag for SFTP:



Note

Start an FTP session from your computer to the switch using the same user name and password used to Telnet or SSH to the switch. Upload or copy the image to the switch.

```
boot config flag <ftpd | sshd>  
  
end
```

3. Download the files to the switch through FTP or SFTP, or transfer them to the switch through the USB port.
4. Enter Privileged EXEC configuration mode by exiting the Global Configuration mode.
`exit`

5. Extract the release distribution files to the `/intflash/release/` directory:

```
software add WORD<1-99> [-y]
```

After you enter this command, the switch verifies it has sufficient space to install the new image. If space is not insufficient, the switch checks the total space of the log files. If the total log file space is sufficient for the new image, the switch prompts you to remove log files.

**Note**

The image file is removed after the **software add** command runs.

6. (Optional) Respond if you see the following prompt:

```
Log files need to be removed to have enough space to install the new image.
```

```
Do you want to remove all log files except the current one?
```

If your response is affirmative, the switch removes all log files except the currently active file and displays the following message for each file removed: `Log file %s was removed`

7. Install the image:

```
software activate WORD<1-99>
```

8. Restart the switch:

```
reset
```

**Important**

After you restart the system, you have the amount of time configured for the commit timer to verify the upgrade and commit the software to gold. If you do not commit the software to gold and auto-commit is not enabled, the system restarts with the last known working version after the commit timer has expired. This feature ensures you can regain control of the system if an upgrade fails. By default, auto-commit is enabled.

9. After you restart the switch, enter Privileged EXEC configuration mode:

```
enable
```

10. Confirm the software is upgraded:

```
show software
```

11. Commit the software:

```
software commit
```

**Important**

If you disable the auto-commit feature, you must run the **software commit** command manually before the commit timer expires to commit the new software version, otherwise the system restarts automatically to the previous (committed) version. By default, auto-commit is enabled.

Example



Note

The image file name is switch dependent. See *Fabric Engine Release Notes* for information about file names.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#boot config flags ftpd
Switch:1(config)#end
Switch:1#copy /usb/5520.9.0.0.0.voss /intflash/5520.9.0.0.0.voss
Switch:1#software add 5520.9.0.0.0.voss -y
Switch:1#software activate 5520.9.0.0.0.GA
Switch:1#reset -y
Switch:1>show software

=====
                        software releases in /intflash/release/
=====
8.10.0.0.GA                (Backup Release)                (Signed Release)
9.0.0.0.GA                 (Primary Release)                (Signed Release)
-----

Auto Commit      : enabled
Commit Timeout   : 10 minutes

Remaining time until software auto-commit is 8 minutes 59 seconds

Switch:1>show software detail

=====
                        software releases in /intflash/release/
=====
9.0.0.0.GA                (Primary Release) (Signed Release)
  SSIO
    UBOOT                  2.2.1.4
    APP_FS                 9.0.0.0.GA
  AVAILABLE ENCRYPTION MODULES
    No Modules Added

8.10.0.0.GA                (Backup Release) (Signed Release)
  SSIO
    UBOOT                  2.2.1.14
    APP_FS                 8.10.0.0.GA
  AVAILABLE ENCRYPTION MODULES
    No Modules Added
```

Variable Definitions

The following table defines parameters for the **software** command.

Variable	Value
<code>activate WORD<1-99></code>	<p>Copies the software version to the boot flash file.</p> <p>When you use the software activate command, the system checks for hardware dependencies and prevents a downgrade if it detects a dependency. For example, if a hardware component has a minimum software version dependency, you cannot downgrade to an incompatible software version or install the hardware component in a chassis that runs an incompatible software version.</p>
<code>add WORD<1-99></code>	<p>Unpacks a software release <version>.</p> <p>After you enter this command, the switch verifies it has sufficient space to install the new image. If space is not insufficient, the switch checks the total space of the log files. If the total log file space is sufficient for the new image, the switch prompts you to remove all log files except the currently active file. If your response is affirmative, the switch displays the following message for each file removed: Log file %s was removed</p>
<code>-y</code> Note: Exception: not supported on 5520 Series.	<p>Suppresses the confirmation message to automatically overwrite the non-primary image. If you omit this parameter, you must confirm the action to overwrite the non-primary image.</p>

Verify the Upgrade

Verify your upgrade to ensure proper switch operation.

Procedure

1. Check for alarms or unexpected errors:

```
show logging file tail
```
2. Verify all modules and slots are online:

```
show sys-info
```

Commit an Upgrade

About This Task

The software commit functionality for software upgrades allows maximum time set by the commit timer (the default is 10 minutes) to ensure that the upgrade is successful.

If you enable the auto-commit option, the system automatically commits to the new software version after the commit timer expires.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. (Optional) Configure the timer to activate the software:
`sys software commit-time <1-60>`
3. (Optional) Extend or reduce the time to commit the software:
`software reset-commit-time [<1-60>]`
4. Commit the upgrade:
`software commit`



Important

If you disable the auto-commit feature, you must run the **software commit** command manually before the commit timer expires to commit the new software version, otherwise the system restarts automatically to the previous (committed) version. By default, auto-commit is enabled.

Downgrade the Software

Perform this procedure to downgrade the switch from the current trusted version to a previous release.



Important

MACsec connectivity association (CA) configurations fail during downgrade. If you plan to downgrade MACsec to an earlier version, delete the MACsec CA entries, perform the downgrade, and then reconfigure the MACsec CA entries. This applies to both 2AN and 4AN modes.

Before You Begin

- Ensure that you have a previous version uploaded to the switch.
- Ensure you have a compatible configuration file if you downgrade to an earlier release. New releases contain significant enhancements, which cannot be used in previous software versions. For information about feature support, see *Fabric Engine and VOSS Feature Matrix*.
- If you need to downgrade the image on ExtremeCloud IQ Managed Switches to release 9.0.0.0, from 9.0.2.0, or later, you must remove the file `.telegraf.csv` from the `/intflash` directory if it exists. Failure to do so can cause the switch to crash and revert to 9.0.2.0. For more information, see *Fabric Engine Release Notes*.

About This Task



Note

The image file name is switch dependent. See *Fabric Engine Release Notes* for information about file names.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Extract the release distribution files to the `/intflash/release/` directory:

```
software add WORD<1-99> [-y]
```

After you enter this command, the switch verifies it has sufficient space to install the new image. If space is not insufficient, the switch checks the total space of the log files. If the total log file space is sufficient for the new image, the switch prompts you to remove log files.



Note

The image file is removed after the **software add** command runs.

3. (Optional) Respond if you see the following prompt:

```
Log files need to be removed to have enough space to install the new image.
```

```
Do you want to remove all log files except the current one?
```

If your response is affirmative, the switch removes all log files except the currently active file and displays the following message for each file removed: `Log file %s was removed`

4. Activate a prior version of the software:

```
software activate WORD<1-99>
```

5. Restart the switch:

```
reset
```



Important

After you restart the system, you have the amount of time configured for the commit timer to verify the software change and commit the software to gold. If you do not commit the software to gold and auto-commit is not enabled, the system restarts with the last known working version after the commit timer expires. This feature ensures you can regain control of the system if an upgrade fails. By default, auto-commit is enabled.

6. Commit the software change:

```
software commit
```

**Important**

If you disable the auto-commit feature, you must run the **software commit** command manually before the commit timer expires to commit the new software version, otherwise the system restarts automatically to the previous (committed) version. By default, auto-commit is enabled.

7. Verify the downgrade:

- Check for alarms or unexpected errors using the `show logging file tail` command.
- Verify all modules and slots are online using the `show sys-info` command.

8. (Optional) Remove unused software:

```
software remove WORD<1-99>
```

Variable Definitions

The following table defines parameters for the `software` command.

Variable	Value
<code>activate WORD<1-99></code>	<p>Copies the software version to the boot flash file.</p> <p>When you use the <code>software activate</code> command, the system checks for hardware dependencies and prevents a downgrade if it detects a dependency. For example, if a hardware component has a minimum software version dependency, you cannot downgrade to an incompatible software version or install the hardware component in a chassis that runs an incompatible software version.</p>
<code>add WORD<1-99></code>	<p>Unpacks a software release <version>.</p> <p>After you enter this command, the switch verifies it has sufficient space to install the new image. If space is not insufficient, the switch checks the total space of the log files. If the total log file space is sufficient for the new image, the switch prompts you to remove all log files except the currently active file. If your response is affirmative, the switch displays the following message for each file removed: Log file %s was removed</p>

Variable	Value
<code>remove WORD<1-99></code>	Specifies the path and version of the compressed software release archive file to remove.
<code>-y</code> Note: Exception: not supported on 5520 Series.	Suppresses the confirmation message to automatically overwrite the non-primary image. If you omit this parameter, you must confirm the action to overwrite the non-primary image.

Remove a Software Build

Use the following procedure to remove a software build for the switch.



Note

You can install a maximum of 6 software releases on some switches. When the limit is reached, you are prompted to remove one release before you can proceed with adding and activating a new software release.

On the 4220 Series, 5320 Series, and 5420 Series, you can store a maximum of 2 software releases. When the limit is reached, the software displays a confirmation to overwrite the non-primary release before you can install a new software release.

For more information, see [Upgrade the Software](#) on page 233.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Remove the software build:
`software remove WORD<1-99>`

Example

Remove the software build:

```
Switch:1>enable
Switch:1#software remove w.x.y.z
```

Update the Complex Programmable Logic Device (CPLD) Image

During the device bootup, if an older version of a CPLD module is detected, the system displays a log message to upgrade the CPLD module image.

You can also use **show sys-info cpld** command to check the current version of the CPLD module on the device.

Before You Begin

Upgrade the software on the switch to the latest build.

About This Task

The **cpld-install** command compares the image version of the modules with the current version on the device:

- If the versions are the same, the command exits.
- If the current version is an earlier version, you must update the image version of the specific module.

The device automatically restarts after successful installation of the specific module.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Update one of the following CPLDs:

- APP:

```
cpld-install app [WORD<1-99>]
```

- BOOT:

```
cpld-install boot [WORD<1-99>]
```

- Main CPLD module:

```
cpld-install main [WORD<1-99>]
```

- Port:

```
cpld-install port [WORD<1-99>]
```

3. When prompted, type *y* to continue with the CPLD update.

Variable Definitions

The following table defines parameters for the **cpld-install** command.

Variable	Value
<i>app</i>	Updates the APP module.
<i>boot</i>	Updates the BOOT module.
<i>main</i>	Updates the main CPLD module.
<i>port</i>	Updates the Port module.
<i>WORD<1-99></i>	Specifies the image file name. Note: This parameter is optional. If you do not specify the file name, the command checks the image file for the image from the running file system.

LRM/MACsec Adapter

Table 25: LRM/MACsec Adapter product support

Feature	Product	Release introduced
LRM support Note: MACsec functionality of the adapter is not supported.	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.9 Supported on 5720-VIM-6YE
	7520 Series	Not Supported
	7720 Series	Not Supported
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	Not Supported

The LRM/MACsec Adapter (Part number: 10695) is a two-port external adapter that can be used in conjunction with a switch to enable the use of 10 Gigabit Long Reach Multimode (LRM) optics and/or MACsec link encryption on the switch. The adapter effectively enables either LRM or MACsec support on SFP+ switch ports that do not otherwise support these capabilities.

The adapter includes four physical SFP+ ports: two for direct connection to the host switch and two for external links. Host ports use a proprietary 10 Gbps direct attach cable (DAC) to connect to the host. The adapter includes two host cables. Fully managed from the host switch, the adapter does not require setup using a local console cable or Ethernet port.



Note

You must connect both host ports to trigger PHY initialization and you must insert the DAC in the adapter before you insert the DAC in the switch.

The adapter provides the necessary hardware to drive up to two LRM external links when connected to the host switch's 10 Gbps SFP+ ports. The adapter can simultaneously support MACsec encryption across the same LRM links. MACsec-only operation is also possible on host switch ports when LRM is not required.



Note

MACsec encryption is not supported with the adapter when used with Universal switches.

Both host ports and an auxiliary USB port provide power to the adapter. As a best practice, always use the USB port to provide power.

If the PHY firmware on the adapter is out-of-date, you must upgrade the adapter firmware. To understand if you need to upgrade, see [Upgrade the LRM/MACsec Adapter PHY Firmware](#) on page 244.

For more information about the adapter, see the following resources:

- [Two-Port External LRM/MACsec Adapter](#)
- *5720 Series Installation Guide*

Upgrade the LRM/MACsec Adapter PHY Firmware



Note

Only 5720-VIM-6YE supports the LRM/MACsec Adapter.

Update the PHY firmware if the running version is out-of-date. The latest version of firmware is bundled with the Network Operating System (NOS) image.

About This Task

The upgrade process can take up to 10 minutes to complete. After completion, the adapter restarts and reinitializes.

As with regular switch ports, configuration is based on the optics detected. For information about how to configure VIM port speeds, see [Configure Port Speeds for All VIM Ports](#) on page 730 (CLI) or [Configure Port Speed for All VIM Ports](#) on page 795 (EDM).

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Confirm if an upgrade is required:

```
show sys lrm-adapter
```



Note

After you first connect the adapter, the switch also displays a log message if the firmware is out-of-date.

3. Upgrade the firmware:

```
sys lrm-adapter-upgrade <0-32>
```

Examples

The following example shows the console message that indicates an upgrade is required, and then the command to perform the upgrade. After you connect the adapter, the adapter ID displays in the log message that indicates you must perform a firmware upgrade. Use this adapter ID as part of the upgrade command.

```
5720-24MXW-FabricEngine:1>enable
5720-24MXW-FabricEngine:1#configure terminal
```

```

Enter configuration commands, one per line. End with CNTL/Z.
1 2022-09-20T16:31:07.340Z 5720-24MXW-FabricEngine IO1 - 0x000e0610 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 port 2/2 has PHY firmware 0xd006. Min
required is 0xd008. Please run "sys lrm-adapter-upgrade <adapter_id>" to upgrade.
1 2022-09-20T16:31:07.340Z 5720-24MXW-FabricEngine IO1 - 0x000e0609 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 initialization completed.
5720-24MXW-FabricEngine:1(config)#sys lrm-adapter-upgrade 0
WARNING: Running this command will upgrade the LRM/MACSec Adapter PHY firmware. The
operation will take few minutes to complete. During this time the LRM/MACSec adapter will
restart and reinitialize.
Do you want to continue (y/n) ? y

1 2022-09-20T16:31:50.566Z 5720-24MXW-FabricEngine IO1 - 0x000e0612 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 PHY firmware upgrade is in progress.
1 2022-09-20T16:31:53.810Z 5720-24MXW-FabricEngine IO1 - 0x000e0607 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 initialization started.
1 2022-09-20T16:31:53.964Z 5720-24MXW-FabricEngine CP1 - 0x0000c5e7 - 00300001.256
DYNAMIC SET GlobalRouter HW INFO Link Down(2/1)
1 2022-09-20T16:31:53.971Z 5720-24MXW-FabricEngine CP1 - 0x0000c5e7 - 00300001.257
DYNAMIC SET GlobalRouter HW INFO Link Down(2/2)
1 2022-09-20T16:31:54.816Z 5720-24MXW-FabricEngine IO1 - 0x000e0608 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 initialization in progress (25%).
1 2022-09-20T16:43:43.304Z 5720-24MXW-FabricEngine IO1 - 0x000e0608 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 initialization in progress (50%).
1 2022-09-20T16:44:09.406Z 5720-24MXW-FabricEngine IO1 - 0x000e0608 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 initialization in progress (75%).
1 2022-09-20T16:44:27.686Z 5720-24MXW-FabricEngine IO1 - 0x000e0613 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 PHY firmware upgrade completed.
1 2022-09-20T16:44:27.686Z 5720-24MXW-FabricEngine IO1 - 0x000e0609 - 00000000
GlobalRouter HAL INFO LRM/MACSec Adapter #0 initialization completed.
1 2022-09-20T16:44:28.161Z 5720-24MXW-FabricEngine CP1 - 0x0000c5ec - 00300001.257
DYNAMIC CLEAR GlobalRouter HW INFO Link Up(2/2)
1 2022-09-20T16:44:31.040Z 5720-24MXW-FabricEngine CP1 - 0x000145f2 - 00000000
GlobalRouter BRIDGE WARNING Spanning Tree: Port 2/2 - Received 2 TCs in 1 minute(s)
1 2022-09-20T16:44:33.037Z 5720-24MXW-FabricEngine CP1 - 0x002d0609 - 00000000
GlobalRouter LLDP INFO New LLDP Neighbor Discovered on interface 2/2

```

The following example displays a connected adapter with out-of-date firmware.

```

Switch:1>show sys lrm-adapter
=====
                        LRM/MACSec Adapter General Info
=====
ADAPTER      PORT      PHY FW      UPGRADE
ID           NUM       VERSION    NEEDED
-----
0            2/1    2/2        D006        Yes
=====

```



Managing the Switch

[Fundamentals](#) on page 246
[Basic Configuration Using CLI](#) on page 284
[System Access Configuration Using EDM](#) on page 309
[Verification](#) on page 320
[Enhanced Secure Mode Configuration Using CLI](#) on page 323
[Enhanced Secure Mode Configuration Using EDM](#) on page 337
[Basic Administration Procedures using CLI](#) on page 338
[Basic Administration Procedures Using EDM](#) on page 354
[Boot Parameter Configuration Using the CLI](#) on page 356
[Run-Time Process Management Using CLI](#) on page 371
[Run-Time Process Management Using EDM](#) on page 377
[Hardware Status Using EDM](#) on page 379

The following topics provide instructions to perform basic configuration of, and administrative tasks for, the switch and software.

Examples and network illustrations may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Fundamentals

This section includes the minimum, but essential, configuration steps to:

- provide a default, starting point configuration
- establish basic security on the node

For more information about hardware specifications and installation procedures, see the hardware documentation.

System Connections

Connect the serial console interface (an RJ-45 jack) to a PC or terminal to monitor and configure the switch. The port uses a RJ-45 connector that operates as data terminal equipment (DTE). Some switches also provide a USB port or micro USB port for serial console interface connectivity. See your hardware documentation for available ports.

The default communication protocol settings for the console port are:

- Baud rate is 115200
- 8 data bits
- 1 stop bit
- No parity
- No flow control

To use the console port, you need a terminal or teletypewriter (TTY)-compatible terminal, or a portable computer with a serial port and terminal-emulation software.

Boot Sequence

Table 26: Linux kernel version product support

Feature	Product	Release introduced
Linux kernel version	4220 Series	5.10 as of Fabric Engine 9.2
	5320 Series	5.10 as of Fabric Engine 9.1
	5420 Series	5.10 as of Fabric Engine 9.1
	5520 Series	5.10 as of Fabric Engine 9.1
	5720 Series	5.10 as of Fabric Engine 9.1
	7520 Series	5.10 as of Fabric Engine 9.1
	7720 Series	5.10 as of Fabric Engine 9.1
	7830 Series	5.10 as of Fabric Engine 9.3
	VSP 4900 Series	5.10 as of VOSS 9.1
	VSP 7400 Series	5.10 as of VOSS 9.1

The switch goes through a boot sequence before it becomes fully operational. After you turn on power to the switch, the system starts.

The boot sequence consists of the following stages:

- [Loading Linux](#) on page 248
- [Loading the Primary Release](#) on page 249
- [Deploying Zero Touch](#) on page 249 or [Loading the Configuration File](#) on page 249



Note

Extreme Networks offers universal hardware products that support more than one Network Operating System (NOS) personality. The first time you start a universal hardware product, the boot sequence can be different from what is documented in this section. The boot sequence documented in this section assumes a NOS selection of Fabric Engine is already established. For more information on NOS personalities, see [Network Operating System Personalities](#) on page 2717.

The following figure shows a summary of the boot sequence.

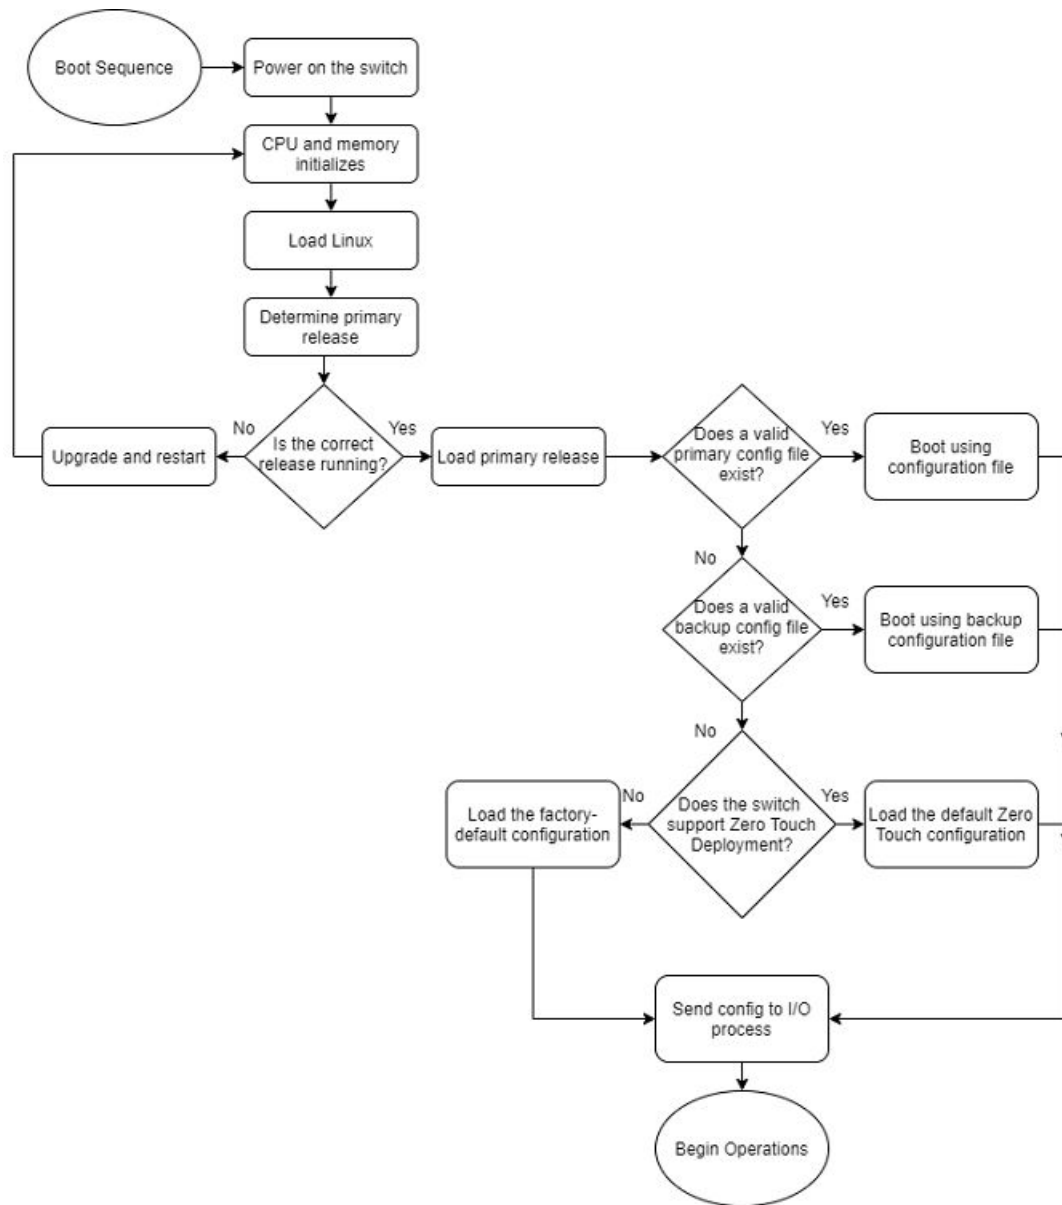


Figure 16: Boot sequence

Loading Linux

Depending on the Linux kernel used, the boot image is stored either in a boot flash partition, Secure Digital (SD), or Solid State Drive (SSD) flash card. The boot image includes the boot loader, and the Linux kernel and applications.

The boot location contains two versions of the boot image: a committed version (the primary release) and a backup version. A committed version is one that is marked as good (if you can start the system using that version). The system automatically uses the backup version if the system fails the first time you start with a new version.

Loading the Primary Release

You can store up to six software releases on the 5520 Series, 5720 Series, 7520 Series, and 7720 Series. If you have six releases already stored on the switch, you are prompted to remove one release before you can proceed to add and activate a new software release.

You can store a maximum of two software releases on the 4220 Series, 5320 Series, and 5420 Series. If you attempt to add a third software release, the software displays a confirmation to overwrite the non-primary release. You can also use the **software add <filename> -y** to bypass the confirmation question and automatically overwrite the non-primary release.

The system saves software image files to the `/intflash/release/` directory.

After loading the primary release, the CPU and basic system devices, such as the console port, initialize. At this stage, the I/O ports are not available; the system does not initialize the I/O ports until the port sends configuration data.

Deploying Zero Touch



Important

Zero Touch Deployment does not function if primary or secondary configuration files exist.

The switch automatically deploys when the system loads the primary release and the switch is in a Zero Touch Deployment-ready configuration mode.



Note

The switch does not automatically deploy if the option **reset-all-files** or **zero-touch** is selected before reboot.

For more information, see [Zero Touch Deployment](#) on page 91.

Loading the Configuration File

After the system loads the primary release, it identifies the location and file name of the primary configuration file. You can save this file in internal flash.

If the primary configuration file does not exist, the system looks for the backup configuration file, as identified by `version.cfg`. If this file does not exist, the system initiates Zero Touch functionality on the switch that enables Zero Touch Fabric Configuration. For more information, see [Zero Touch Fabric Configuration](#) on page 101.

The switch configuration consists of higher-level functionality, including:

- Chassis configuration
- Port configuration
- Virtual LAN (VLAN) configuration
- Routing configuration

- IP address assignments
- Remote monitoring (RMON) configuration

The default switch configuration in Zero Touch Fabric Configuration mode includes the following:

- Shortest Path Bridging MAC (SPBM) instance is created.
- Intermediate System-to-Intermediate System (IS-IS) is enabled.
- All ports are enabled and operating in Auto-sense mode.
- The switch issues DHCP requests on the out-of-band (OOB) management port and the management VLAN.

The default switch configuration in factory default mode includes the following:

- A single, port-based default VLAN with a VLAN identification number of 1
- No interface assigned IP addresses
- Traffic priority for all ports configured to normal priority
- All ports as untagged ports
- Default communication protocol settings for the console port. For more information about these protocol settings, see [System Connections](#) on page 246.

Configuration File Statements

In the configuration file, statements preceded by both the number sign (#) and exclamation point (!) load prior to the general configuration parameters. Statements preceded by only the number sign are comments meant to add clarity to the configuration; they do not load configuration parameters. The following table illustrates the difference between these two statement formats.

Table 27: Configuration file statements

Sample statement	Action
# software version : 8.6.0.0	Adds clarity to the configuration by identifying the software version.
#!no boot config flags sshd	Configures the flag to the false condition, prior to loading the general configuration.

Boot Sequence Modification

You can change the boot sequence in the following ways:

- Change the primary designations for file sources.
- Change the file names from the default values. You can store several versions of the configuration file and specify a particular one by file name. The specified configuration file only gets loaded when the chassis starts. To load a new configuration file, you need to restart the system.
- Start the system without loading an existing configuration file so that the system uses the factory default configuration. You can do this by running the **boot config flags factorydefaults** command. For more information, see [Enable Factory Default Behavior](#) on page 363.

The factorydefaults boot flag removes the runtime, primary, and backup configuration files, resets all local default user account passwords, and removes all digital certificates. The Radsec, IPsec, IKE, OSPF, SNMP, SSL, SSH, OVSDDB, and NTP files are also removed. The CLI displays a warning that the configurations, passwords, and files will be reset, and the system logs an informational message. The configuration and file removals occur during the next boot sequence when the factorydefaults boot flag is enabled. After the switch reboots, the security mode setting is retained. To enable Zero Touch Onboarding after a factorydefaults boot, reboot the switch again without saving a configuration.

- Start the system in Zero Touch Deployment mode, which includes Zero Touch Fabric Configuration. For more information, see [Zero Touch Deployment](#) on page 91.

Runtime

After the switch is operational, you can use the runtime commands to perform configuration and management functions necessary to manage the system. These functions include the following

- Resetting or restarting the switch
- Adding, deleting, and displaying address resolution protocol (ARP) table entries
- Pinging another network device
- Viewing and configuring variables for the entire system and for individual ports
- Configuring and displaying MultiLink Trunking (MLT) parameters
- Creating and managing port-based VLANs or policy-based VLANs

To access the runtime environment you need a connection from a PC or terminal to the switch. You can use a direct connection to the switch through the console port or remotely through Telnet or Secure Shell (SSH) sessions.



Important

Before you attempt to access the switch using one of the preceding methods, ensure you first enable the corresponding daemon flags.

System Logon

After the platform boot sequence is complete, the system opens the logon prompt.



Note

With enhanced secure mode enabled, the person in the role-based authentication level of administrator configures the login and password values for the other role-based authentication levels. The administrator initially logs on to the switch using the default login of `admin` and the default password of `admin`. After the initial login, the switch prompts the administrator to create a new password. The administrator then configures default logins and passwords for the other users based on the role-based authentication levels of the user.

The following table shows the default values for logon and password for console and Telnet sessions.

Table 28: Access levels and default logon values

Access level	Description	Default logon	Default password
Read-only	Permits view-only configuration and status information. Is equivalent to Simple Network Management Protocol (SNMP) read-only community access.	ro	ro
Layer 1 read/write	View most switch configuration and status information and change physical port settings.	l1	l1
Layer 2 read/write	View and change configuration and status information for Layer 2 (bridging and switching) functions.	l2	l2
Layer 3 read/write	View and change configuration and status information for Layer 2 and Layer 3 (routing) functions.	l3	l3
Read/write	View and change configuration and status information across the switch. You cannot change security and password settings. This access level is equivalent to SNMP read/write community access.	rw	rw
Read/write/all	Permits all the rights of read/write access and the ability to change security settings, including CLI and web-based management user names and passwords and the SNMP community strings.	rwa	rwa

You can enable or disable users with particular access levels, eliminating the need to maintain large numbers of access levels and passwords for each user.

The system denies access to a user with a disabled access level who attempts to log on. The system displays the following error message after a user attempts to log on with a blocked access level:

```
CPU1 [mm/dd/yy hh:mm:ss] 0x0019bfff GlobalRouter CLI WARNING Slot 1: Blocked unauthorized cli access
```

The system logs the following message to the log file:

```
User <user-name> tried to connect with blocked access level <access-level> from <src-ipaddress> via <login type>.
```

The system logs the following message for the console port:

```
User <user-name> tried to connect with blocked access level <access-level> from console port.
```

RADIUS Authentication

Remote Authentication Dial-in User Service (RADIUS) authentication takes precedence over the local configuration. If you enable RADIUS authentication on the switch, the user can access the switch even if you block an access level on the switch.



Important

When you enable RADIUS on the switch and configure a RADIUS server to be used by CLI or EDM, the server authenticates the connection, whether it is FTP, HTTPS, SSH, or TELNET. However, in the event that the RADIUS server is unresponsive or is unreachable, the switch will fall back to the local authentication, so that you can access the switch using your local login credentials.

If you disable an access level, all running sessions, except FTP sessions, with that access level to the switch terminate.



Important

Only the RWA user can disable an access level on the switch. You cannot disable the RWA access level on the switch.
The system preserves these configurations across restarts.

hsecure Mode Boot Configuration Flag

The switch supports a configurable flag called high secure (hsecure). Use the hsecure flag to enable the following password features:

- 10 character enforcement
- aging time
- limitation of failed login attempts
- protection mechanism to filter designated IP addresses

If you activate the **hsecure** flag, the software enforces the 10-character rule for all passwords. The password must contain a minimum of two uppercase characters, two lowercase characters, two numbers, and two special characters.

If you enable hsecure for the first time and the password file does not exist, then the device creates a normal default user name (rwa) and password (rwa). In this case, the password does not meet the minimum requirements for hsecure and as a result the system prompts you to change the password.

For more information about the hsecure flag, see [hsecure Mode](#) on page 3369.

Enhanced Secure Mode

If you enable enhanced secure mode, the system uses different authentication levels. Enhanced secure mode allows the system to:

- Provide role-based access levels
- Stronger password requirements
- Stronger rules on password length
- Stronger rules on password complexity
- Stronger rules on password change intervals
- Stronger rules on password reuse
- Stronger password maximum age use

For more information on enhanced secure mode, see [Enhanced Secure Mode Configuration Using CLI](#) on page 323.

Default Web-Server Behavior

The default switch configuration enforces the following restrictions for web-server access:

- The web-server password must be a minimum of 8 characters.
- Secure communications with the web server use Transport Layer Security (TLS) version 1.2 and later.
- The switch does not support the RC4 cipher. The switch supports the following ciphers:
 - TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
 - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
 - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
 - TLS_RSA_WITH_AES_128_CBC_SHA
 - TLS_RSA_WITH_AES_128_CBC_SHA256
 - TLS_RSA_WITH_AES_256_CBC_SHA
 - TLS_RSA_WITH_AES_256_CBC_SHA256

For information about how to enable and configure the web server, see [Configure the Web Server](#) on page 195 or [Configure the Web Management Interface](#) on page 214. For information about supported browser versions, see [Supported Browsers](#) on page 204.

Secure and Nonsecure Protocols

The following table describes the secure and nonsecure protocols that the switch supports.

Table 29: Secure and nonsecure protocols for IPv4 and IPv6

Nonsecure protocols	Default status	Equivalent secure protocols	Default status
FTP and Trivial FTP	Disabled	Secure Copy (SCP) and Secure File Transfer Protocol (SFTP)	Disabled
Note: File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP) support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.			
Telnet	Disabled	Secure Shell version 2 (SSHv2)	Disabled
SNMPv1, SNMPv2	Enabled	SNMPv3	Enabled
HTTP	Disabled	HTTPS Important: Take the appropriate security precautions within the network if you use HTTP. You must use the web-server enable command in CLI before you can access EDM.	Enabled

Client and Server Support

Table 30: Client and Server product support

Feature	Product	Release introduced
File Transfer Protocol (FTP) server and client (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 30: Client and Server product support (continued)

Feature	Product	Release introduced
File Transfer Protocol (FTP) server and client (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 30: Client and Server product support (continued)

Feature	Product	Release introduced
Secure Copy (SCP) Note: The switch does not support the WinSCP client.	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Secure File Transfer Protocol (SFTP) server (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Secure File Transfer Protocol (SFTP) server (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 30: Client and Server product support (continued)

Feature	Product	Release introduced
Telnet server and client (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Telnet server and client (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Trivial File Transfer Protocol (TFTP) server and client (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 30: Client and Server product support (continued)

Feature	Product	Release introduced
TFTP server (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
TFTP client (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 31: Secure Shell product support

Feature	Product	Release introduced
Secure Shell (SSH) server (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	Fabric Engine 9.3

Table 31: Secure Shell product support (continued)

Feature	Product	Release introduced
Secure Shell (SSH) client (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Secure Sockets Layer (SSL) certificate management	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
SSH server (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 31: Secure Shell product support (continued)

Feature	Product	Release introduced
SSH client (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
SSH client disable	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
SSH rekey	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The client-server model partitions tasks between servers that provide a service and clients that request a service.

For active CLI clients, users initiate a client connection from the switch to another device.

**Note**

Both FTP and TFTP clients are supported by the switch. The switch does not launch FTP and TFTP clients explicitly as a separate command; you can launch them through the CLI **copy** command. If you have configured the user name through the **boot config host** command, the FTP client is used to transfer files to and from the switch using the CLI **copy** command; If you have not configured the user name, the TFTP client is used to transfer files to and from the switch using the CLI **copy** command.

Configuring the `boot config flags ftpd` or `boot config flags tftpd` enables the FTP or TFTP Servers on the switch.

For non-active clients, the client exists on the switch and the switch console initiates the request, with no intervention from users after the initial setup. For instance, Network Time Protocol (NTP) is a non active client. The switch initiates the client request to the central server to obtain the up-to-date time.

Password Encryption

The platform stores passwords in encrypted format and not in the configuration file.

**Important**

For security reasons, configure the passwords to values other than the factory defaults.

CLI Passwords

The switch ships with default passwords configured for access to CLI through a console or Telnet session. If you possess read-write-all access authority, and you use SNMPv3, then you can change passwords in encrypted format. If you use EDM, then you can also specify the number of allowed Telnet sessions.

**Important**

Be aware that the default passwords and community strings are documented and well known. Change the default passwords and community strings immediately after the first login.

For security, if you fail to log on correctly in three consecutive instances, then the device locks for 60 seconds.

The switch stores passwords in encrypted format and not in the configuration file.

Subscriber or Administrative Interaction

As a network administrator, you can configure the RADIUS server for user authentication to override user access to commands. You must still provide access

based on the existing access levels in the switch, but you can customize user access by allowing and denying specific commands.

You must configure the following three returnable attributes for each user:

- Access priority (single instance)—the access levels currently available on the switch (ro, l1, l2, l3, rw, rwa)
- Command access (single instance)—indicates whether the user has access to the commands on the RADIUS server
- CLI commands (multiple instances)—the list of commands that the user can or cannot use

Multiple CLI Users for Each Role

Table 32: Multiple CLI Users product support

Feature	Product	Release introduced
Multiple CLI users per role	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

You can create up to a maximum of 10 CLI users for each role, which includes:

- 3 default users (rwa, rw, and ro)—User Type = default
- 7 user defined users (rwa or rw or ro)—User Type = userDefined

User names for default users (rwa, rw, and ro) can be changed; however, user names for user defined users cannot be changed. User names support the following character types:

- alphanumeric
- hyphen (-)
- underscore (_)

Users require a user name and password to connect to the switch. Users can log on through the local serial port, Telnet, SSH, or ftp. When a user is created, authentication is enabled, by default.

For security reasons, if a login attempt fails, the error feedback does not indicate if the failed login is due to an invalid user name or an invalid password. Response times for

invalid user name and invalid user name and password pair are identical to prevent identification of which of the two failed.

**Note**

Multiple CLI users for each role does not apply in enhanced secure mode.

Enterprise Device Manager

The switch includes Enterprise Device Manager (EDM), an embedded graphical user interface (GUI) that you can use to manage and monitor the platform through web-based access without additional installations.

For more information about EDM, see [Enterprise Device Manager](#) on page 203.

Enterprise Device Manager Access

To access EDM, enter one of the following addresses in your web browser:

- `http://<A.B.C.D>`
- `https://<A.B.C.D>`

Where <A.B.C.D> is the device IP address.

Ensure you use a supported browser version. For more information about supported browsers, see [Supported Browsers](#) on page 204.

**Important**

For security reasons, EDM is disabled by default.

You must enable the web server from CLI (see [Configure the Web Server](#) on page 195) to enable HTTP access to EDM. For HTTP access to the device, you must also disable the web server secure-only option. The web server secure-only option, allowing for HTTPS access to the device, is enabled by default. As a best practice, take the appropriate security precautions within the network if you use HTTP

If you experience issues while connecting to EDM, check the proxy settings. Proxy settings can affect EDM connectivity to the switch. Clear the browser cache and do not use proxy when connecting to the device.

For information about user names and passwords, see [Default User Names and Passwords for EDM](#) on page 204.

*TLS Server for Secure HTTPS***Table 33: TLS server for secure HTTPS product support**

Feature	Product	Release introduced
TLS server for secure HTTPS	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

This feature enhances communications security by using the Transport Layer Security (TLS) cryptographic protocol.

The following are the key properties of Secure web server with TLS:

- This feature can be implemented on a maximum of only 10 concurrent client connections.
- The switch supports version TLS 1.2 and later, by default. You can explicitly configure TLS 1.0 and TLS 1.1 version support using CLI or EDM.
- This feature replaces SSL 3.0 with TLS. SSL 3.0 is not supported.
- The TLS server does not support RC4, DES, TDES, and MD5 based cipher suites.
- The minimum password length for the web server is 8 characters, by default. You can change this using CLI or EDM.

For information about the certificate order priority when the Transport Layer Security (TLS) server and switch connect, see [Certificate Order Priority](#) on page 3384.

Enhanced Secure Mode

Table 34: Enhanced Secure Mode product support

Feature	Product	Release introduced
Enhanced Secure mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Enhanced Secure mode for JITC and non-JITC sub-modes.	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Enhanced Secure mode - sensitive file protection	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

Authentication Levels

After you enable enhanced secure mode with the **boot config flags** **enhancedsecure-mode** command, the switch supports role-based authentication

levels. With enhanced secure mode enabled, the switch supports the following authentication access levels for local authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+) authentication:

- Administrator
- Privilege
- Operator
- Auditor
- Security

Each user name is associated with a certain role in the product and appropriate authorization rights for viewing and executing commands are available for that role.

With enhanced secure mode enabled, the person in the role-based authentication level of administrator configures the login and password values for the other role-based authentication levels. The administrator access level cannot be disabled on VOSS switches.

The administrator initially logs on to the switch using the default login of `admin` and the default password of `admin`. After the initial login, the switch prompts the administrator to create a new password.

The following displays an example of the initial login to the switch by the administrator after enhanced secure mode is enabled.

```

Login: admin
Password: *****
      This is an initial attempt using the default user name and password.
      Please change the user name and password to continue.
Enter the new name : rwa
Enter the New password : *****
Re-enter the New password : *****
Password changed successfully
      Last Successful Login:Wed Oct 14 12:20:42 2015
      Unsuccessful Login attempts from last login is: 0

```

The administrator then configures default logins and passwords for the other users based on the role-based authentication levels of the user.

Access Level and Login Details

Access level	Description	Login location
Administrator	The administrator access level permits all read-write access, and can change security settings. The administrator access level can configure CLI and web-based management user names, passwords, and the SNMP community strings. The administrator access level can also view audit logs.	SSH/Telnet (in band/mgmt)/console
Privilege	The privilege access level has the same access permission as the administrator; however, the privilege access level cannot use RADIUS or TACACS+ authentication.	Console

Access level	Description	Login location
Operator	The operator access level can view most switch configurations and status information. The operator access level can change physical port settings at layer 2 and layer 3. The operator access level cannot access audit logs or security settings.	SSH/Telnet(in band/mgmt)/console/
Auditor	The auditor access level can view configuration information, status information, and audit logs.	SSH/Telnet(in band/mgmt)/console/
Security	The security access level can change security settings only. The security access level also has permission to view configuration and status information.	SSH/Telnet(in band/mgmt)/console/

Password Requirements

After you enable enhanced secure mode on the switch, the password requirements are stronger. The individual in the administrator access level role configures and provides a temporary user name and password. After you log in for the first time with the temporary user name and temporary password, the system forces you to change the temporary password. After you change the temporary password, you cannot use the password again in subsequent sessions.

The following topics discuss the enhanced password requirements.

Password Complexity Rule

After you enable enhanced secure mode, the system checks each password change request to ensure the new password meets the password complexity required.

The default for the password complexity rule includes the following:

- Two uppercase characters, from the range: ABCDEFGHIJKLMNOPQRSTUVWXYZ
- Two lowercase characters, from the range: abcdefghijklmnopqrstuvwxyz
- Two numeric characters, from the range: 1234567890
- Two special characters, from the range: `~!@#\$%^&*()_-=+{[]|\:;'"<,.>?/

Password Length Rule

The system enforces a minimum password length of 15 characters after you enable enhanced secure mode.

If the password is not long enough, the system displays the following message:

```
Password change aborted. The new password does not meet the minimum
complexity requirement. Please select another password that meets the
change interval, length, complexity, no consecutive repeating characters
or history requirements of the domain.
```

Password Change Interval Rule

The system enforces a minimum password change interval, which defines the minimum amount of time before you can change to a new password. By default, the minimum change interval is 24 hours between changing from one password to

a new password. When you attempt to change your password, the system checks the timestamp for your password to determine if enough time has passed to enable you to change the password.

If you attempt to change the password and not enough time has passed, the system rejects the request, and the system informs you that the password was recently changed. Any password change outside of the enforced interval requires the Administrator to approve the change.

If you try to change the password before the change interval allows, the system displays the following message:

```
Password change aborted. The new password does not meet the minimum
complexity requirement. Please select another password that meets the
change interval, length, complexity, no consecutive repeating characters
or history requirements of the domain.
```

Minimum Password Change

The system enforces a minimum password change requirement, which defines that 8 characters must differ within the same position from the old password.

If the new password does not have at least 8 characters changed within the same position from the old password, the system rejects the password and displays the following message:

```
The password change failed, less than eight characters were changed.
```

Password Reuse Rule

After you enable enhanced secure mode, the administrator access level can define the number of old passwords that cannot be reused. The password reuse rule ensures that recently used passwords are not reused immediately, which reduces the risk of someone unlawfully gaining access to the system. The default number of prohibited recently used passwords is 3, but you can define up to 99.

The system saves the password history and stores the history in an encrypted format, along with the user name and date of change. If a particular user attempts to change a password, the system checks the new password against the stored passwords the user has previously used. If the password is on the list of previously used passwords, the system rejects the password change, and displays the following message:

```
Old password not allowed.
```

Password Maximum Age Rule

The system enforces automatic password renewal and password lockout after the expiration period because long-term usage of the same password can cause the system to be vulnerable to hacking.

You can configure the password expiration period to a range of 1 to 365 days. The default password expiration period is 90 days.

Password Max-Session

The password max-sessions value indicates the maximum number of times a particular type of role-based user can log in to the switch through the SSH session at the same time. The max-sessions value applies only for SSH sessions, and only with enhanced secure mode enabled.

After the maximum session number is reached that particular type of user cannot login. For example, if the max-sessions for an auditor user is configured as 5, then the auditor user can log in to only five SSH sessions at the same time. The default is 3.

Password Pre-Notification Interval and Post-Notification Interval Rule

After enhanced secure mode is enabled, the switch enforces password expiry. To ensure a user does not lose access, the switch offers pre- and post-notification messages explaining when the password will expire.

The administrator can define pre- and post-notification intervals to between one to 99 days.

The system maintains the password with a time stamp for when the password expiration. When you log in, the system checks the password time stamp and the notification timer values. If the administrator configures the pre-notification to 30 days, when you log in, the system checks the time stamp and notification timer values, and if the password expiry is due in 30 days, the system displays the first notification.

The pre-notification intervals provide messages to warn users that their passwords will expire within a particular time frame:

- interval 1—By default, interval 1 is 30 days.
- interval 2—By default, interval 2 is 7 days.
- interval 3—By default, interval 3 is 1 day.

The post-notification intervals provide notification to users that their passwords have expired within a particular time frame:

- interval 1—By default, interval 1 is 1 day.
- interval 2—By default, interval 2 is 7 days.
- interval 3—By default, interval 3 is 30 days.

If you do not change the password before the expiry date, the system locks your account. When your account is locked, only the administrator can unlock the account. The administrator creates a temporary password, and then you can login with the temporary password. If the administrator password expires and the configured notification interval lapses, access to SSH/Telnet/FTP connections is denied. The administrator must connect to the console port using a serial connection to change the password.

Sensitive File Protection

For certain switches in enhanced secure mode, sensitive files and paths are protected. The home directory for enhanced secure mode is `/intflash/shared`. You cannot access sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. You

cannot access sensitive files using CLI commands. Files transferred on the switch through the default path are saved in `/intflash/shared`.

The following sensitive files and paths are protected in enhanced secure mode:

- `/intflash/.cert`
- `/intflash/.ike_psk.txt`
- `/intflash/.isis_md5key.txt`
- `/intflash/.isis_sha2key.txt`
- `/intflash/.isis_simplekey.tx`
- `/intflash/.ntp_keys.txt`
- `/intflash/.radsec/conf`
- `/intflash/.radsec/profile`
- `/intflash/.shadov.txt`
- `/intflash/.shadovfedmoc.txt`
- `/intflash/.snmp_comm_moc_fed.txt`
- `/intflash/.snmp_usm_moc_fed.txt`
- `/intflash/.ssh/id_dsa_*`
- `/intflash/.ssh/id_rsa_*`
- `/intflash/.ssh/moc_sshc_dsa_file`
- `/intflash/.ssh/moc_sshc_rsa_file`
- `/intflash/.ssh/ssh_dss.key`
- `/intflash/.ssh/ssh_rsa.key`
- `/intflash/.sys/is_esm`
- `/intflash/.sys/is_sha512`
- `/intflash/app/restweb/certs/cert.pem`
- `/intflash/app/restweb/certs/privkey.pem`
- `/intflash/app/restweb/certs/server.pem`
- `/intflash/app/server.pem`
- `/intflash/app/slamon/certs/trustcerts.txt`
- `/intflash/ospf_key.txt`
- `/intflash/ospf_vrfif_key.txt`
- `/intflash/ospf_vrfvif_key.txt`
- `/intflash/ovsdb/keys/privatekey.pem`
- `/intflash/ovsdb/keys/sc-cert.pem`
- `/intflash/server.pem`
- `/intflash/snmp_comm_moc.txt`
- `/intflash/snmp_usm_moc.txt`

Sensitive files cannot be accessed using the following CLI commands:

- **cd**
- **copy**

- **move**
- **move**
- **rename**
- **remove**
- **delete**
- **more**
- **edit**
- **dir**
- **mkdir**

Operational Considerations

- If you attempt to change the value of the `enhancedsecure-mode` flag when the system is running, you are prompted to continue or to cancel the action. If you decide to continue, all sensitive files are deleted. You must save the current configuration and then reset the switch for the change to take effect.
- When you upgrade the software from versions earlier than VOSS 8.5, even if the `enhancedsecure-mode` flag is changed, no sensitive files are deleted.

If you upgrade the software to VOSS 8.5, if the `enhancedsecure-mode` flag is changed, all sensitive files are deleted.

- If you enable the `boot config flags factorydefaults` configuration flag to return an existing switch to factory default configuration, in enhanced secure mode, all sensitive files are deleted. When you enable the factory default optional parameters with default configuration, it preserves the user security files and security mode. For more information, see [Enable Factory Default Behavior](#) on page 363.
- When the switch resets after you enable enhanced secure mode or you upgrade the switch software, all configuration files, such as the runtime, primary, and backup configuration files are copied to `/intflash/shared`. The system displays the following message:

```
GlobalRouter SW INFO The runtime config file /intflash/config.cfg is
copied to /intflash/shared/config.cfg due to changing ESM mode/upgrade
version.
```

When the switch resets after you disable enhanced secure mode or after you downgrade the switch software, the switch uses the same configuration files as it used before the switch reset, either from `/intflash/shared` or from `/intflash`.

System Flags

After you enable or disable certain modes and functions, you need to save the configuration, and then restart the switch for your change to take effect. This section lists parameters and indicates if they require a switch restart.

The following table lists parameters you configure in the CLI using the **boot config flags** command. For information on system flags and their configuration, see [Boot Flag Configuration using CLI](#) on page 360.

Table 35: Boot configuration flags

Flag	Restart
<i>advanced-feature-bandwidth-reservation</i> Note: For product support, see Advanced Feature Bandwidth Reservation on page 275.	Yes
<i>block-snmp</i>	No
<i>debug-config</i>	Yes
<i>debugmode</i>	Yes
<i>dvr-leaf-mode</i> Note: Exception: does not apply to 4220 Series.	No Note: The switch does not require a restart when you enable the dvr-leaf-mode flag, as long as there is no unsupported configuration on the switch.
<i>enhancedsecure-mode</i>	Yes
<i>factorydefaults [config-only reset-all-files zero-touch zero-touch-config-only]</i>	Yes
<i>flow-control-mode</i> Note: Exception: does not apply to 7830 Series.	Yes
<i>ftpd</i>	No
<i>hsecure</i>	Yes
<i>ipv6-egress-filter</i> Note: Exception: does not apply to 4220 Series.	Yes
<i>ipv6-mode</i> Note: Exception: does not apply to 4220 Series and 7830 Series.	Yes
<i>logging</i>	No

Table 35: Boot configuration flags (continued)

Flag	Restart
<i>macsec</i> Note: Exception: Only required for MACsec supporting models of 5320 Series and 5420 Series.	Yes
<i>nni-mstp</i> Note: Exception: does not apply to 4220 Series.	Yes
<i>reboot</i>	No
<i>spanning-tree-mode</i> Note: Only MSTP mode is compatible with Fabric SPBM. If you configure RSTP mode, the switch cannot participate in an SPBM network.	Yes
<i>spbm-config-mode</i>	Yes
<i>spbm-node-scaling</i> Note: Exception: only applies to 5320 Series (not 5320-16P-2MXT-2X or 5320-24T-4X-XT) and 5420 Series.	Yes
<i>sshd</i>	No
<i>telnetd</i>	No
<i>tftpd</i>	No
<i>trace-logging</i>	No
<i>urpf-mode</i>	Yes
<i>verify-config</i>	Yes

Table 35: Boot configuration flags (continued)

Flag	Restart
<i>vrf-scaling</i> Note: Exception: Does not apply to switch models that support a single active VRF with IP configuration, or to 5320-16P-2MXT-2X and 5320-24T-4X-XT.	Yes
<i>vxlان-gw-full-interworking-mode</i> Note: Exception: Only supported on 7520 Series and 7720 Series.	Yes

Advanced Feature Bandwidth Reservation

Table 36: Advanced Feature Bandwidth Reservation product support

Feature	Product	Release introduced
Advanced Feature Bandwidth Reservation Note: If the switch does not support this boot flag, it is because the hardware reserves the bandwidth automatically with no user interaction.	4220 Series	Fabric Engine 9.2 Supported on 4220-4MW-8P-4X, 4220-4MW-20P-4X, and 4220-8X
	5320 Series	Fabric Engine 8.6 Supported on 5320-24P-8XE, 5320-24T-8XE, 5320-48P-8XE, and 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3 Supported on 7830-VIM-16CE (ports 3/13-3/16 only) and 7830-VIM-8DE (ports 3/7-3/8 only)
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.0

The switch enables this boot configuration flag by default to use advanced features on the switch. If you disable this flag and attempt to enable an advanced feature the switch displays an error message to explain why the advanced feature failed to start, and to remind you that you must enable this boot configuration flag for that advanced feature.



Important

If you change the configuration, you must save the configuration, and then reboot the switch for the change to take effect.

If you disable this feature and save the configuration, any configuration for advanced features remains saved in the configuration file but is not used.

When you disable this feature, you can use all ports for Layer 2 or Layer 3 forwarding of standard unicast and multicast features. Use this mode if you are not configuring advanced features.

When you enable this feature, also known as Full Feature mode, the switch supports advanced features by reassigning some of the front-panel ports to be loopback ports.

Advanced Features

The following tables identifies the features that use Advanced Feature Bandwidth Reservation.



Note

Full Feature mode does not support PIM.

Feature support can vary by hardware. For high-level feature support information, see *Fabric Engine and VOSS Feature Matrix*.



Important

Advanced Bandwidth Feature Reservation must be enabled when using SPBM even if features requiring loopback bandwidth are not used.

Table 37: Advanced Feature Bandwidth Reservation features

Feature	4220 Series	5320 Series	5420 Series	5520 Series	5720 Series	7520 Series	7720 Series	7830 Series
Application Telemetry	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Bridged and Routed Layer 2 VSNs (C-VLAN UNIs)	Yes	Yes	Yes					
Layer 2 VSN	Yes	Yes	Yes					
Layer 2 VSN (NNI-UNI MiM Multicast only)				Yes	Yes	Yes	Yes	Yes
E-TREE	Yes	Yes	Yes					

Table 37: Advanced Feature Bandwidth Reservation features (continued)

Feature	4220 Series	5320 Series	5420 Series	5520 Series	5720 Series	7520 Series	7720 Series	7830 Series
Fabric Extend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
I-SID Mirroring	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Inter-I-SID Routing	Yes	Yes	Yes					
IPFIX		Yes	Yes	Yes	Yes	Yes	Yes	
IS-IS Accept Policies (Inter-VRF Routing)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Layer 3 VSNs	Yes	Yes	Yes					
Multicast over SPB - UNI to NNI	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Segmented Management CLIP	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Simplified Virtual IST			Yes	Yes	Yes	Yes	Yes	
VXLAN Gateway						Yes	Yes	

**Note**

For 7520 Series and 7720 Series, you can consider using low bandwidth mode instead of high bandwidth mode to free up two reserved ports for use as regular ports.

Configuration Modes

Depending on the hardware, this feature supports three modes:

- low
- high
- vim

When you select *low* or *high*, the corresponding bandwidth of the reserved Ethernet ports is allocated to the loopback required for advanced feature functionality. When

you select *vim*, the bandwidth is allocated from the unused VIM slot capacity; you cannot select this parameter if the VIM is present.



Note

The following switch models do not require Advanced Bandwidth Feature Reservation boot flag configuration to support advanced features, and so do not reserve front-panel ports. For feature support information, see [Table 36](#) on page 275.

- 4220-12P-4X
- 4220-12T-4X
- 4220-24P-4X
- 4220-24T-4X
- 5320-16P-2MXT-2X
- 5320-16P-4XE and 5320-16P-4XE-DC
- 5320-24T-4X-XT
- 5320-24T-24S-4XE-XT

The following table shows the supported modes and the ports reserved as loopback ports by each mode. An empty cell indicates the mode (parameter) is not supported.

Table 38: Reserved loopback ports

Platform	Modes supported	Reserved ports by mode		
		low	high	vim
4220-4MW-8P-4X 4220-4MW-20P-4X 4220-8X	low high (default)	none - all front-panel ports available. UNI-NNI and NNI-UNI bandwidth is limited to 1 Gbps.	last 2 SFP+ ports	
5320-24P-8XE 5320-24T-8XE	low	1/25-1/27		
5320-48P-8XE 5320-48T-8XE	low	1/49-1/51		
5420F-24T-4XE 5420F-8W-16P-4XE 5420F-24P-4XE 5420F-24S-4XE 5420M-24T-4YE 5420M-24W-4YE	low	Universal Ethernet ports 1/29 and 1/30 Note: To understand restrictions on using reserved loopback ports as front-panel ports, see 5420 Series on page 4220.		

Table 38: Reserved loopback ports (continued)

Platform	Modes supported	Reserved ports by mode		
		low	high	vim
5420F-48T-4XE 5420F-16MW-32P-4XE 5420F-16W-32P-4XE 5420F-48P-4XE 5420F-48P-4XL 5420M-48T-4YE 5420M-48W-4YE 5420M-16MW-32P-4YE 5420M-24W-24S-4YE	low	Universal Ethernet ports 1/53 and 1/54 Note: To understand restrictions on using reserved loopback ports as front-panel ports, see 5420 Series on page 4220.		
5520-24T 5520-24W 5520-24X	low vim (default if vim is not installed, otherwise low)	Universal Ethernet ports 1/25 and 1/26		VIM slot reserved
5520-12MW-36W 5520-48SE 5520-48T 5520-48W	low vim (default if vim is not installed, otherwise low)	Universal Ethernet ports 1/49 and 1/50		VIM slot reserved
5720-24MW 5720-24MXW	high vim (default if vim is not installed, otherwise high)		none - all front-panel ports available	VIM slot reserved Note: This mode uses less internal bandwidth speed than high mode.
5720-48MW 5720-48MXW	high vim (default if vim is not installed, otherwise high)		Universal Ethernet ports 1/49 and 1/50	VIM slot reserved
7520-48XT-6C	low (default) high	All front-panel ports are available	with VXLAN Full Interworking Mode disabled: 1/54 with VXLAN Full Interworking Mode enabled: 1/53 and 1/54	

Table 38: Reserved loopback ports (continued)

Platform	Modes supported	Reserved ports by mode		
		low	high	vim
7520-48Y-8C 7520-48YE-8CE	low (default) high	1/55 and 1/56	with VXLAN Full Interworking Mode disabled: 1/54, 1/55, and 1/56 with VXLAN Full Interworking Mode enabled: 1/53, 1/54, 1/55, and 1/56	
7720-32C	low (default) high	1/31 and 1/32	with VXLAN Full Interworking Mode disabled: 1/30, 1/31, and 1/32 with VXLAN Full Interworking Mode enabled: 1/29, 1/30, 1/31, and 1/32	
7830-32CE-8DE	low	3/13-3/16 if you install on 7830-VIM-16CE in slot 3 3/7-3/8 if you install on 7830-VIM-8DE in slot 3. Otherwise, no ports are reserved.		

**Note**

To understand restrictions on using reserved loopback ports as front-panel ports on the 5420 Series, see [5420 Series](#) on page 4220.

Reserved Ports

After the switch reserves the appropriate ports to become loopback ports, the ports are no longer visible in the output when you enter **show interfaces gigabitEthernet**.

By default, for the 5520 Series and 5720 Series switches, if you do not install a VIM in the switch, the switch uses the vim parameter. In this configuration, the VIM ports are used as loopback ports and the Universal Ethernet ports are used as regular uplink ports. When used as regular uplink ports, the port speed for the Universal Ethernet ports is 40 Gbps as a single channel port. Although the maximum supported single channel

port speed is 40 Gbps, the ports can be channelized to operate as four 10 or 25 Gbps channels.

If a VIM is already installed in the 5520 Series, the switch uses the low parameter that uses Universal Ethernet ports as loopback ports and VIM ports as regular uplink ports. For the 5720 Series, the switch uses the high parameter that uses Universal Ethernet ports as loopback ports and VIM ports as regular uplink ports.



Important

You must ensure your configuration does not include reserved ports before you enable this feature. If the configuration includes reserved ports, after you enable this feature and restart the switch, the switch stops loading the configuration.

spbm-config-mode boot flag

Table 39: spbm-config-mode product support

Feature	Product	Release introduced
spbm-config-mode (boot config flags spbm-config-mode)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Shortest Path Bridging (SPB) and Protocol Independent Multicast (PIM) cannot interoperate with each other on the switch at the same time. To ensure that SPB and PIM stay mutually exclusive, the software uses a boot flag called **boot config flags spbm-config-mode**.

- The **boot config flags spbm-config-mode** flag is enabled by default. This enables you to configure SPB and IS-IS, but you cannot configure PIM and IGMP either globally or on an interface.

- If you disable the boot flag, save the configuration, and then reboot with the saved configuration. When the flag is disabled, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.

**Note**

You cannot disable the boot flag on the 7830 Series.

**Important**

After you change the **boot config flags spbm-config-mode** flag, you must save the configuration, and then reboot the switch for the change to take effect.

For more information about this boot flag and Simplified vIST, see [IP Multicast](#) on page 1696.

nni-mstp boot config flag

Table 40: nni-mstp boot flag product support

Feature	Product	Release introduced
nni-mstp boot flag (boot config flags nni-mstp)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The nni-mstp boot flag changes the default behavior of the MSTP on SPBM network-to-network interface (NNI) ports. The Common and Internal Spanning Tree (CIST) is disabled automatically on the NNI, and the NNI ports can only be members of backbone VLANs (B-VLAN).

- During startup, if you have non-B-VLAN on SPBM NNI ports in your configuration file, the system sets the nni-mstp flag to true (if it was not already set to true) and enables MTSP on SPBM NNI ports, and all other configurations remain the same. Save your configuration file. If you do not save your configuration, you continue to see the following message on reboot:

Warning

Detected brouter and/or vlans other than BVLANS on NNI ports. Setting the boot config


```
flag nni-mstp to true. Saving configuration avoids repetition of this warning on
reboot.
```

**Note**

When the nni-mstp flag is set to true, only MSTI 62 is disabled on the SPBM NNI ports. You can add the SPBM NNI ports to any VLAN.

- If you configure the nni-mstp boot configuration flag to false (default), the system checks to make sure that the SPBM NNI ports do not have brouter (IPv4 or IPv6) or non-SPBM VLANs configured. The nni-mstp flag is then set to false. Save your configuration file, and reboot the switch for the configuration change to take effect.

**Note**

Ensure that all SPBM NNI ports in non-B-VLAN are removed prior to setting the nni-mstp flag to false.

Example: Configuring nni-mstp to true

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags nni-mstp
Warning: Please save the configuration and reboot the switch for this configuration to
take effect.
Switch:1(config)#
```

Access Policies for Services

You can control access to the switch by creating an access policy. An access policy specifies the hosts or networks that can access the switch through various services, such as Telnet, Simple Network Management Protocol (SNMP), Hypertext Transfer Protocol (HTTP), and Secure Shell version 2 (SSHv2). You can enable or disable access services by configuring flags.

Use access policies for in-band management to secure access to the switch. When configuring an access policy, a lower precedence takes higher priority if you use multiple policies. For example, preference 120 has priority over preference 128.

You can define network stations that can access the switch or stations that cannot access the switch. For each service you can also specify the level of access, such as read-only or read-write-all.

**Important**

A third-party security scan shows the switch service ports open and in the listen state. No connections are accepted on these ports unless you enable the particular daemon. The switch does not dynamically start and stop the daemons at runtime and needs to keep them running from system startup.

When you configure access policies, you can perform either of the following actions:

- Globally enable the access policy feature, and then create and enable individual policies. Each policy takes effect immediately after you enable it.

- Create and enable individual access policies, and then globally enable the access policy feature to activate all the policies at the same time.

HTTP and SSH support IPv4 and IPv6 with no difference in configuration or functionality.

Managing the System using Different VRF Contexts

You can manage the system using different Virtual Router Forwarding (VRF) contexts:

- Using the GlobalRouter (VRF 0), you can manage the entire system. GlobalRouter is the default EDM view at log in.
- Using a VRF context other than the GlobalRouter (VRF 0), you have limited functionality to manage the system. For instance, you can only manage the ports assigned to the specified VRF instance.

In CLI, configure a specific VRF using the VRF Router Configuration Mode. For more information, see [CLI Command Modes](#) on page 178. For more information about using EDM to configure a specific VRF, see [Select and Launch a VRF Context View](#) on page 4100.

You can use the context names (SNMPv3) and community strings (SNMPv1/v2) to assign different VRFs to manage selected components, such as ports and VLANs. For more information about context names and community strings, see [Simple Network Management Protocol \(SNMP\)](#) on page 3467.

Basic Configuration Using CLI

This section provides CLI procedures to connect to a switch and configure initial access and user accounts.

Connect a Terminal

Before You Begin

- To use the console port, you need the following equipment:
 - A terminal or Teletypewriter (TTY)-compatible terminal, or a portable computer with a serial port and terminal-emulation software.
 - A specific cable with an RJ-45 or USB connector for the console port on the switch. The other end of the cable must use a connector appropriate to the serial port on the computer or terminal.
- To comply with emissions regulations and requirements, you must shield the cable that connects to the console port.



Note

If you are using the USB console port with a terminal running Windows 10, you must install the CP210x USB to UART Bridge Virtual COM Port (VCP) driver from Silicon Labs before you connect to the terminal.

About This Task

Connect a terminal to the serial console interface to monitor and configure the system directly.

Procedure

1. Configure the terminal protocol as follows:
 - 115200 baud
 - 8 data bits
 - 1 stop bit
 - No parity
 - No flow control
2. Connect the RJ-45 or USB cable to the console port on the switch.
3. Connect the other end of the cable to the terminal or computer serial port.
4. Turn on the terminal.
5. Log on to the switch.

Use Telnet to Log on to the Device

About This Task

Use Telnet to log on to the device and remotely manage the switch.

Procedure

1. From a PC or terminal, start a Telnet session:
`telnet <ipv4 or ipv6 address>`
2. Enter the logon and password when prompted.

Example

```
C:\Users\jsmith>telnet 192.0.2.40
Connecting to 192.0.2.40.....
Login:rwa
Password:rwa
```

Change the Default CLI User Names

You can change user names for default users (rwa, rw, and ro) for the specified access level. You cannot change user names for user-defined users.

About This Task



Note

Only the rwa user can change the default user name.

Perform this procedure to change the default user name for the specified access level.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Change the default user name:

```
username WORD <1-20> level {l1|l2|l3|ro|rw|rwa}
```

The system displays the following message, in this case for the ro access level:

```
Do you want to change username for the default RO user ?
```

3. Type y.
4. Enter the old password.
5. Enter the new password.
6. Re-enter the new password.

Example

Change the default user name for the RO access to CLI:

```
Switch:1>enable
Switch:1#config t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#username EXUser level ro

Do you want to change username for the default RO user ? (y/n) ? y

Enter the old password : **
Enter the New password : *****
Re-enter the New password : *****
```

Variable Definitions

The following table defines parameters for the **username** command.

Variable	Value
<i>WORD</i> <1-20>	Specifies the user logon name.
<i>level</i> l1 l2 l3 ro rw rwa	Specifies the access level.

Change Passwords for CLI Accounts

Before You Begin

- You must use an account with read-write-all privileges to change passwords. For security, the switch saves passwords to a hidden file.

About This Task

Configure new passwords for each access level, or change the logon or password for the different access levels of the switch. After you receive the switch, use default passwords to initially access CLI. If you use Simple Network Management Protocol version 3 (SNMPv3), you can change encrypted passwords.

If you enable the `hsecure` flag, after the aging time expires, the system prompts you to change your password. If you do not configure the aging time, the default is 90 days.

If you enable enhanced secure mode with the `boot config flags enhancedsecure-mode` command, you enable new access levels, along with stronger password complexity, length, and minimum change intervals. For more information, see [Enhanced Secure Mode](#) on page 266.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Change a password:

```
cli password WORD<1-20> {layer1|layer2|layer3|read-only|read-write|read-write-all}
```

3. Enter the old password.

4. Enter the new password.

5. Re-enter the new password.

6. Configure password options:

```
password [access-level WORD<2-8>] [aging-time day <1-365>] [default-lockout-time <60-65000>] [lockout WORD<0-46> time <60-65000>] [min-passwd-len <10-20>] [password-history <3-32>]
```

Example

Change a password for an access level of read-write-all and configure the expiration period for the password to 60 days:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#cli password rwa read-write-all
Enter the old password: ***
Enter the new password: ***
Re-enter the new password: ***
Switch:1(config)#password access-level rwa aging-time 60
```

Variable Definitions

Use the data in the following table to use the `cli password` command.

Variable	Value
<code>layer1 layer2 layer3 read-only read-write read-write-all</code>	Changes the password for the specific access level.
<code>WORD<1-20></code>	Specifies the user logon name.

Use the data in the following table to use the **password** command.

Variable	Value
<code>access-level WORD<2-8></code>	Permits or blocks this access level. The available access level values are as follows: <ul style="list-style-type: none"> • layer1 • layer2 • layer3 • read-only • read-write • read-write-all
<code>aging-time day <1-365></code>	Configures the expiration period for passwords in days, from 1–365. The default is 90 days.
<code>default-lockout-time <60-65000></code>	Changes the default lockout time after three invalid attempts. Configures the lockout time, in seconds, and is in the 60–65000 range. The default is 60 seconds.
<code>lockout WORD<0-46> time <60-65000></code>	Configures the host lockout time. <ul style="list-style-type: none"> • <code>WORD<0-46></code> is the host IPv4 or IPv6 address. • <code><60-65000></code> is the lockout-out time, in seconds, in the 60–65000 range. The default is 60 seconds.
<code>min-passwd-len <10-20></code>	Configures the minimum length for passwords in high-secure mode. The default is 10 characters.
<code>password-history <3-32></code>	Specifies the number of previous passwords the switch stores. You cannot reuse a password that is stored in the password history. The default is 3.

Enable CLI Access Levels

Enable CLI access levels to control the configuration actions of various users.

About This Task



Important

Only the RWA user can disable an access level on the switch. You cannot disable the RWA access level on the switch.

The system preserves these configurations across restarts.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable an access level:
`password access-level WORD<2-8>`

Example

Block CLI access to Layer 1:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no password access-level 11
```

Variable Definitions

The following table defines parameters for the **password access-level** command.

Variable	Value
<i>WORD<2-8></i>	<p>Permits or blocks this access level. The available access level values are as follows:</p> <ul style="list-style-type: none">• 11 — Specifies Layer 1.• 12 — Specifies Layer 2.• 13 — Specifies Layer 3.• ro — Specifies read-only.• rw — Specifies read-write.• rwa — Specifies read-write-all. <p>By default, the system permits all access levels.</p>

Configure System Identification

Configure system identification to specify the system name, contact person, and location of the switch.

Procedure

1. Log on as rwa.
2. Enter Global Configuration mode:
`enable`
`configure terminal`
3. Change the system name:
`sys name WORD<0-255>`
4. Configure the system contact:
`snmp-server contact WORD<0-255>`

5. Configure the system location:

```
snmp-server location WORD<0-255>
```

Example

Change the system name, configure the system contact, and configure the system location:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#sys name Floor3Lab2
Floor3Lab2:1(config)#snmp-server contact http://companyname.com
Floor3Lab2:1(config)#snmp-server location "12 Street, City, State, Zip"
```

Variable Definitions

Use the data in the following table to use the system-level commands.

Variable	Value
<i>contact</i> WORD<0-255>	Identifies the contact person who manages the node. To include blank spaces in the contact, use quotation marks (") around the text.
<i>location</i> WORD<0-255>	Identifies the physical location of the node. To include blank spaces in the location, use quotation marks (") around the text.
<i>name</i> WORD<0-255>	Configures the system or root level prompt name for the switch. WORD<0-255> is an ASCII string from 1–255 characters (for example, LabSC7 or Closet4).

Configure the CLI Banner

Configure the logon banner to display a message to users before authentication and configure a system login message-of-the-day in the form of a text banner that displays after each successful logon.

About This Task

You can use the custom logon banner to display company information, such as company name and contact information. For security, you can change the default logon banner of the switch, which contains specific system information, including platform type and software release.

Use the custom message-of-the-day to update users on a configuration change, a system update or maintenance schedule. For security purposes, you can also create a message-of-the-day with a warning message to users that, "Unauthorized access to the system is forbidden."

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure the switch to use a custom banner or use the default banner:
`banner <custom|static>`
3. Create a custom banner:
`banner WORD<1-80>`



Note

To enter multiple lines for a message, use the **banner** command before each new line of the message. To provide a string with spaces, include the text in quotation marks.

4. Create the message-of-the-day:
`banner motd WORD<1-1516>`



Note

To enter multiple lines for a message, use the **banner motd** command before each new line of the message. To provide a string with spaces, include the text in quotation marks.

5. Enable the custom message-of-the-day:
`banner displaymotd`
6. Save the configuration:
`save config`
7. Display the banner information:
`show banner`
8. Log on again to verify the configuration.

Example

Configure the custom banner to “Company, www.Companyname.com.” and configure the message of the day to “Unauthorized access to this system is forbidden. Please logout now.”

```
Switch:1> enable
Switch:1#configure terminal
Switch:1(config)# banner custom
Switch:1(config)# banner Company
Switch:1(config)# banner www.Companyname.com
Switch:1(config)# banner motd "Unauthorized access to this system is forbidden"
Switch:1(config)# banner motd "Please logout now"
Switch:1(config)#banner displaymotd
Switch:1(config)#show banner
Company
www.company.com
                        defaultbanner : false
                        custom banner :
```

```

displaymotd : true
custom motd :
Unauthorized access to this system is forbidden
Please logout now

```

Variable definitions

Use the data in the following table to use the **banner** command.

Variable	Value
<i>custom</i>	Disables the use of the default banner.
<i>static</i>	Activates the use of the default banner.
<i>WORD <1-80></i>	Adds lines of text to the CLI logon banner.
<i>motd WORD<1-1516></i>	Create the message of the day. To provide a string with spaces, include the text in quotation marks (").
<i>displaymotd</i>	Enable the custom message of the day.

Configure the Date

Configure the calendar time in the form of month, day, year, hour, minute, and second.

About This Task

Log on as rwa to perform this procedure.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Configure the date:

```
clock set <MMddyyyyhhmmss>
```

3. Verify the configuration:

```
show clock
```

Example

Configure the date and time, and then verify the configuration.

```

Switch:1>enable
Switch:1#clock set 01222024163312
Switch:1#show clock
Mon Jan 22 16:33:12 2024 EDT

```

Variable Definitions

Use the data in the following table to use the **clock set** command.

Variable	Value
<i>MMddyyyyhhmmss</i>	Specifies the date and time in the format month, day, year, hour, minute, and second.

Enable Remote Access Services

Enable the remote access service to provide multiple methods of remote access.

About This Task

File Transfer Protocol (FTP), Trivial File Transfer Protocol (TFTP) and Telnet server support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Enable the access service:

```
boot config flags {ftpd | sshd | telnetd | tftpd}
```
3. Repeat as necessary to activate the desired services.
4. Save the configuration.

Variable Definitions

The following table defines remote access parameters for the **boot config flags** command.

Variable	Value
<i>ftpd</i>	Activates or disables the FTP server on the switch. The default value is disabled. To enable FTP, ensure that the <i>tftpd</i> flag is disabled.
<i>sshd</i>	Activates or disables the SSHv2 server service. The default value is disabled.
<i>telnetd</i>	Activates or disables the Telnet server service. The default is disabled.
<i>tftpd</i>	Activates or disables Trivial File Transfer Protocol server service. The default value is disabled.

Configure an Access Policy

Configure an access policy to control access to the switch.

About This Task

You can permit network stations to access the switch or forbid network stations to access the switch.

For each service, you can also specify the level of access; for example, read-only or read-write-all.

If you configure the access policy mode to *deny*, the system checks the mode and service, and if they match the system, denies the connection. With the access policy mode configured to *deny*, the system does not check *accesslevel* and *access-strict* information.

If you configure the access policy mode to *allow*, the system continues to check the *accesslevel* and *access-strict* information.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create an access policy by assigning it a number:

```
access-policy <1-65535>
```
3. Restrict the access to a specific level:

```
access-policy <1-65535> access-strict
```
4. Configure access for an access policy:

```
access-policy <1-65535> accesslevel <ro|rwa|rw>
```
5. Configure the access policy mode, network, and precedence:

```
access-policy <1-65535> [mode <allow|deny>] [precedence <1-128>]  
[network <A.B.C.D> <A.B.C.D>]
```
6. (Optional) Configure access protocols for an access policy:

```
access-policy <1-65535> [ftp] [http] [ssh] [telnet] [tftp]
```
7. (Optional) Configure trusted access for an access policy:

```
access-policy <1-65535> host WORD<0-46>
```
8. (Optional) Configure SNMP parameters for an access policy:



Note

If you configure an access policy for SNMPv3, you must create an SNMP group and add that group as part of the access policy configuration. For more information about how to create an SNMP group, see [Create a New User Group](#) on page 3481.

- ```
access-policy <1-65535> [snmpv3]
```
  - ```
access-policy <1-65535> [snmp-group WORD<1-32> <snmpv1|snmpv2c|usm>]
```
9. Enable the access policy:

```
access-policy <1-65535> enable
```
 10. Enable access policies globally:

```
access-policy
```

Examples

Assuming no access policies exist, start with policy 3 and name the policy policy3. Add the read-write-all access level and the usm group group_example. Enable access strict, and finally, enable the policy.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#access-policy 3
Switch:1(config)#access-policy 3 name policy3
Switch:1(config)#access-policy 3 accesslevel rwa
Switch:1(config)#access-policy 3 snmp-group group_example usm
Switch:1(config)#access-policy 3 access-strict
Switch:1(config)#access-policy 3 enable
```

Create an SNMPv3 policy that references an existing SNMP group.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#access-policy 1
Switch:1(config)#access-policy 1 snmpv3
Switch:1(config)#access-policy 1 snmp-group Test usm
```

Variable Definitions

The following table defines parameters for the **access-policy** command.

Variable	Value
<i>access-strict</i>	Restraints access to criteria specified in the access policy. <ul style="list-style-type: none"> • true—The system accepts only the currently configured access level. • false—The system accepts access from the configured level up, denying access for lower access levels.
<i>accesslevel</i> <ro rwa rw>	Specifies the level of access if you configure the policy to allow access.
<i>enable</i>	Enables the access policy.
<i>ftp</i>	Activates or disables FTP for the specified policy. Because FTP derives its login and password from the CLI management filters, FTP works for read-write-all (rwa) and read-write (rw) access, but not for the read-only (ro) access.
<i>host</i> WORD<0-46>	For remote login access, specifies the trusted host address as an IP address. The switch supports access-policies over IPv4 and IPv6 with no difference in functionality or configuration.
<i>http</i>	Activates the HTTP and HTTPS for this access policy.

Variable	Value
<i>mode</i> <allow deny>	Specifies whether the designated network address is allowed access to the system through the specified access service. The default is allow. If you configure the access policy mode to <i>deny</i> , the system checks the mode and service, and if they match, the system denies the connection. With the access policy mode configured to <i>deny</i> , the system does not check <i>accesslevel</i> and <i>access-strict</i> information. If you configure the access policy mode to allow, the system continues to check the <i>accesslevel</i> and <i>access-strict</i> information.
<i>name</i> WORD<0-15>	Specifies the access policy name.
<i>network</i> <A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask for IPv4, or the IP address and prefix for IPv6, that can access the system through the specified access service. The switch supports access-policies over IPv4 and IPv6 with no difference in functionality or configuration.
<i>precedence</i> <1-128>	Specifies a precedence value for a policy, expressed as a number from 1–128. The precedence value determines which policy the system uses if multiple policies apply. Lower numbers take higher precedence. The default value is 10.
<i>snmp-group</i> WORD<1-32> <snmpv1 snmpv2c usm>	Adds an SNMP version 3 group under the access policy. WORD<1-32> is the SNMP version 3 group name consisting of 1–32 characters. <snmpv1 snmpv2c usm> is the security model; either snmpv1, snmpv2c, or usm.
<i>snmpv3</i>	Activates SNMP version 3 for the access policy.
<i>ssh</i>	Activates SSH for the access policy.
<i>telnet</i>	Activates Telnet for the access policy.
<i>tftp</i>	Activates the Trivial File Transfer Protocol (TFTP) for this access policy.

Specify a Name for an Access Policy

Before You Begin

The policy must exist before you can name it.

About This Task

Assign a name to an existing access policy to uniquely identify the policy.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Assign a name to the access policy:
`access-policy <1-65535> name WORD<0-15>`

Example

Assign a name to an access policy:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#access-policy 10 name useraccounts
```

Variable Definitions

The following table defines parameters for the **access-policy** command.

Variable	Value
<code>name WORD<0-15></code>	Specifies a name expressed as a string from 0–15 characters.

Allow a Network Access to the Switch

About This Task

Specify the network to which you want to allow access.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Specify the network:
`access-policy <1-65535> [mode <allow|deny>] [network <A.B.C.D> <A.B.C.D>]`

Example

Specify the network to which you want to allow access:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#access-policy 5 mode allow network 192.192.192.0 24
```

Variable Definitions

The following table defines parameters for the **access-policy** command.

Variable	Value
<i>mode</i> <allow deny>	Specifies whether a designated network address is allowed or denied access through the specified access service. The default is allow.
<i>network</i> <A.B.C.D> <A.B.C.D>	Specifies the IPv4 address and subnet mask, or the IPv6 address and prefix-length, permitted or denied access through the specified access service.

Configure Access Policies by MAC Address

About This Task

Configure access-policies by MAC address to allow or deny local MAC addresses on the network management port after an access policy is activated. If the source MAC does not match a configured entry, the default action is taken. For connections coming in from a different subnet, the source MAC of the last hop is used in decision making. Configuring access-policies by MAC address does not perform MAC or Forwarding Database (FDB) filtering on data ports.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Add the MAC address and configure the action for the policy:

```
access-policy by-mac <0x00:0x00:0x00:0x00:0x00:0x00> <allow|deny>
```
3. Specify the action for a MAC address that does not match the policy:

```
access-policy by-mac action <allow|deny>
```

Example

Add the MAC address:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#access-policy by-mac 00-C0-D0-86-BB-E7 allow
```


Variable Definitions

The following table defines parameters for the **access-policy by-mac** command.

Variable	Value
<0x00:0x00:0x00:0x00:0x00:0x00>	Adds a MAC address to the policy. Enter the MAC address in hexadecimal format.
<allow deny>	Specifies the action to take for the MAC address.

Create Multiple CLI Users

You can create up to seven new CLI users on the switch, in addition to the three default CLI users. The user name must be unique. If you enable the hsecure flag, password complexity rules apply to all users.

Before You Begin

You must use an account with read-write-all privileges to create new CLI users.

About This Task



Note

When a new CLI user is created, the specified user name and access level cannot be changed later.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Create a new CLI user:
`username add {<WORD 1-20> level [ro|rw|rwa] enable}`
3. Enter a password.
4. Enter the password a second time.



Note

The password can be up to 80 characters long and cannot contain the user name in either forward or reverse order.

Example

Create a new CLI user:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#username add mar1 level rw enable
Enter password : *****
Re-enter password : *****
```

Variable Definitions

The following table defines parameters for the **username** command.

Variable	Value
<i>add WORD<1-20></i>	Specifies the user name to create.
<i>enable</i>	Enables the new CLI user.
<i>level <ro rw rwa></i>	Specifies the level assigned to the new CLI user: <ul style="list-style-type: none"> • ro: Read-only level • rw: Read-write level • rwa: Read-write-all level

Delete a CLI User Name

Before You Begin

You must use an account with read-write-all privileges to delete a user.

About This Task

Use this task to delete a user name. You cannot delete default ro, rw, and rwa users.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Delete the user name:

```
no username <WORD 1-20>
```

Example

Delete user name "smith":

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no username smith
The specified username will be deleted! Continue (y/n) ? Y
Switch:1(config)#show cli username smith
Username does not exist
```

Variable Definitions

The following table defines parameters for the **no username** command.

Variable	Value
<i>WORD <1-20></i>	Specifies the user name to delete.
<i>enable</i>	Disables the user name.

Enable a CLI User

About This Task

Perform this procedure to enable a user.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable the user:

```
username WORD<1-20> enable
```

3. Verify the user the configuration:

```
show cli username
```

Example

Enable the CLI user:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#username rw enable
Switch:1(config)#username m3 enable
Switch:1(config)#username m2 enable
```

Display CLI user information:

```
Switch:1#show cli username
=====
UserName          AccessLevel    State    Type
=====
ro                 ro             enable   default
rw                 rw             enable   default
rwa                rwa            NA       default
mar1               rw             enable   userDefined
mar2               ro             enable   userDefined
m3                 13             enable   default
m2                 12             enable   default
l1                 11             enable   default
```

Display user name information:

```
Switch:1(config)#show cli username
=====
UserName          AccessLevel    State    Type
=====
ro                 ro             enable   default
rw                 rw             disable  default
rwa                rwa            NA       default
mar1               rw             enable   userDefined
mar2               ro             enable   userDefined
m3                 13             disable  default
m2                 12             disable  default
l1                 11             enable   default
```

Variable Definitions

The following table defines parameters for the **username** command.

Variable	Value
<i>WORD</i> <1-20>	Specifies the user name to enable or disable.
<i>enable</i>	Enables or disables the user name.

Disable a CLI User

About This Task

Perform the following steps to disable users, such as rw, ro, l1, l2, and l3.

You must use an account with read-write-all privileges to disable a user.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Disable the user:

```
no username WORD<1-20> enable
```
3. Verify the user is disabled:

```
show cli username
```

Example

Disable a CLI user:

```
Switch:1(config)#no username rw enable
Switch:1(config)#no username m3 enable
Switch:1(config)#no username m2 enable
```

Display user name information:

```
Switch:1(config)#show cli username
=====
UserName          AccessLevel    State      Type
=====
ro                 ro             enable     default
rw                 rw             disable    default
rwa                rwa            NA         default
mar1               rw             enable     userDefined
mar2               ro             enable     userDefined
m3                 l3             disable    default
m2                 l2             disable    default
l1                 l1             enable     default
```

Display CLI Usernames and Roles

About This Task

Use this task to display CLI usernames and roles.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display CLI usernames and roles:

```
show cli username
```

Example

```
Switch:1#show cli username
=====
UserName          AccessLevel      State      Type
=====
ro                 ro               enable     default
rw                 rw               enable     default
rwa                rwa              NA          default
mar1               rw               enable     userDefined
mar2               ro               enable     userDefined
m3                 l3               enable     default
m2                 l2               enable     default
l1                 l1               enable     default
```

Enable the Web Management Interface

About This Task

Enable the web management interface to provide management access to the switch using a web browser. This procedure assumes this is the first time you enable the web server on a switch with default passwords.

The system prompts you to change the admin and read-only user default passwords when you use the **web-server enable** command to enable the web management interface.



Important

To enable HTTP access to the device, you must disable the web server secure-only option. To enable HTTPS access to the device, the web server secure-only option is enabled by default.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Enable the web server:

```
web-server enable
```
3. (Optional) Enter and confirm the passwords for the RWA (admin) and read-only (user) accounts.

4. (Optional) Disable the secure-only option:

```
no web-server secure-only
```

5. (Optional) Enable read-only user:

```
web-server read-only-user enable
```

6. Save the configuration:

```
save config
```

7. Display the web server status:

```
show web-server
```

Examples

Enable the web server for the first time on a new switch:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#web-server enable
Default password is set for the admin account.
Please set a new password to enable WebServer.
Enter the New password : *****
Re-enter the New password : *****
Password changed.
Default password is set for the ro account.
Please set a new password to enable WebServer.
Enter the New password : *****
Re-enter the New password : *****

Password changed.

Switch:1(config)#web-server read-only-user enable
Switch:1>show web-server
Web Server Info :

      Status                               : on
      Secure-only                           : enabled
      TLS-minimum-version                   : tlsv12
      RO Username Status                    : enabled
      RO Username                           : user
      RO Password                           : *****
      RWA Username                           : admin
      RWA Password                           : *****
      Def-display-rows                      : 30
      Inactivity timeout                    : 900 sec
      Html help tftp source-dir             :
      HttpPort                             : 80
      HttpsPort                             : 443
      NumHits                              : 0
      NumAccessChecks                       : 0
      NumAccessBlocks                       : 0
      NumRxErrors                           : 0

      NumTxErrors                           : 0
      NumSetRequest                         : 1
      Minimum password length               : 8
      Last Host Access Blocked               : 0.0.0.0
      In use certificate                     : Self signed
      Certificate Truspoint CA Name          :
      Certificate with Subject Name          :

      Ciphers-Tls                           : TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
      TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
```

```

                                TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
                                TLS_RSA_WITH_AES_256_CBC_SHA256
TLS_RSA_WITH_AES_256_CBC_SHA
                                TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA
                                : enabled
SSL renegotiation

```

Variable Definitions

Use the data in the following table to use the **web-server** command.

Variable	Value
<i>def-display-rows</i> <10-100>	Configures the number of rows each page displays. The default is 30.
<i>enable</i>	Enables the web interface. The default is disabled.
<i>help-tftp</i> <WORD/0-256>	Configures the TFTP or FTP directory for Help files, in one of the following formats: a.b.c.d:/ peer:/ [<dir>]. The path can use 0-256 characters. The following example paths illustrate the correct format: <ul style="list-style-type: none"> • 192.0.2.1:/help • 192.0.2.1:/
<i>http-port</i> <80-49151>	Configures the web server HTTP port. The default port is 80.
<i>https-port</i> <443-49151>	Configure the web server HTTPS port. The default port is 443.
<i>inactivity-timeout</i> <30-65535>	Configures the web-server session inactivity timeout. The default is 900 seconds (15 minutes).
<i>password</i> {ro rwa} WORD<1-80>	Configures the user names and passwords for the web interface. The default user name for the RO account is user. The default user name for the RWA account is admin.
<i>password min-passwd-len</i> <1-80>	Configures the minimum password length. By default, the minimum password length is 8 characters.
<i>read-only-user</i>	Enables read-only user for the web server. The default is disabled.
<i>secure-only</i>	Enables secure-only access for the web server.

Variable	Value
<code>ssl-renegotiation</code>	Enables SSL renegotiation in the web server. The default is enabled.
<code>tls-min-ver<tlsv10 tlsv11 tlsv12></code>	Configures the minimum version of the TLS protocol supported by the web-server. You can select from the following: <ul style="list-style-type: none"> • <code>tlsv10</code> – Configures the version to TLS 1.0. <p>Note: <code>tlsv10</code> is not supported in enhanced secure mode.</p> <ul style="list-style-type: none"> • <code>tlsv11</code> – Configures the version to TLS 1.1. • <code>tlsv12</code> – Configures the version to TLS 1.2 The default is <code>tlsv12</code> .

Configure the TLS Protocol Version

The switch, by default, supports version TLS 1.2 and later. You can explicitly configure TLS 1.0 and TLS 1.1 version support using CLI.



Note

In enhanced secure mode, TLS 1.0 is available on 5420 Series and 5520 Series only.

Before You Begin

Disable the web server before changing the TLS version. By disabling the web server, other existing users with a connection to the web server are not affected from changing to a different version after you run the command.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Disable the web server:
`no web-server enable`
3. Set the TLS protocol version:
`web-server tls-min-ver [tlsv10 | tlsv11 | tlsv12]`
4. Enable the web server:
`web-server enable`
5. Verify the protocol version:
`show web-server`

Example

```
Switch> enable
Switch# configure terminal
Switch(config)# web-server tls-min-ver tlsv12
```

Verify the protocol version.

```
Switch:1>show web-server
Web Server Info :

      Status                               : on
      Secure-only                           : enabled
      TLS-minimum-version                   : tlsv12
      RO Username Status                    : enabled
      RO Username                           : user
      RO Password                           : *****
      RWA Username                           : admin
      RWA Password                           : *****
      Def-display-rows                      : 30
      Inactivity timeout                    : 900 sec
      Html help tftp source-dir             :
      HttpPort                             : 80
      HttpsPort                             : 443
      NumHits                              : 0
      NumAccessChecks                       : 0
      NumAccessBlocks                       : 0
      NumRxErrors                           : 0

      NumTxErrors                           : 0
      NumSetRequest                         : 1
      Minimum password length               : 8
      Last Host Access Blocked              : 0.0.0.0
      In use certificate                    : Self signed
      Certificate Truspoint CA Name         :
      Certificate with Subject Name         :

      Ciphers-Tls                           : TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
      TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
      TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
      TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
      TLS_RSA_WITH_AES_256_CBC_SHA256
      TLS_RSA_WITH_AES_128_CBC_SHA256
      TLS_RSA_WITH_AES_128_CBC_SHA
      SSL renegotiation                     : enabled
```

Variable Definitions

Use the data in the following table to use the **web-server** command.

Variable	Value
<i>def-display-rows</i> <10-100>	Configures the number of rows each page displays. The default is 30.
<i>enable</i>	Enables the web interface. The default is disabled.

Variable	Value
<i>help-tftp</i> <WORD/0-256>	Configures the TFTP or FTP directory for Help files, in one of the following formats: a.b.c.d:/ peer:/ [<dir>]. The path can use 0-256 characters. The following example paths illustrate the correct format: <ul style="list-style-type: none"> 192.0.2.1:/help 192.0.2.1:/
<i>http-port</i> <80-49151>	Configures the web server HTTP port. The default port is 80.
<i>https-port</i> <443-49151>	Configure the web server HTTPS port. The default port is 443.
<i>inactivity-timeout</i> <30-65535>	Configures the web-server session inactivity timeout. The default is 900 seconds (15 minutes).
<i>password</i> {ro rwa} WORD<1-80>	Configures the user names and passwords for the web interface. The default user name for the RO account is user. The default user name for the RWA account is admin.
<i>password min-passwd-len</i> <1-80>	Configures the minimum password length. By default, the minimum password length is 8 characters.
<i>read-only-user</i>	Enables read-only user for the web server. The default is disabled.
<i>secure-only</i>	Enables secure-only access for the web server.
<i>ssl-renegotiation</i>	Enables SSL renegotiation in the web server. The default is enabled.
<i>tls-min-ver</i> <tlsv10 tlsv11 tlsv12>	Configures the minimum version of the TLS protocol supported by the web-server. You can select from the following: <ul style="list-style-type: none"> tlsv10 – Configures the version to TLS 1.0. <p>Note: tlsv10 is not supported in enhanced secure mode.</p> <ul style="list-style-type: none"> tlsv11 – Configures the version to TLS 1.1. tlsv12 – Configures the version to TLS 1.2 The default is tlsv12.

Back Up Configuration Files

Before and after you upgrade your switch software, make copies of the configuration files. If an error occurs, use backup configuration files to return the switch to a previous state.

Before You Begin

- If you use File Transfer Protocol (FTP) or Trivial File Transfer Protocol (TFTP), ensure that you enable the FTP or TFTP server. File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP) support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.

About This Task

Keep several copies of backup files.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Determine the configuration file names:
`show boot config choice`
3. Save the configuration files. Assuming the files use the default file names, enter:
`save config`
4. Copy the files to a safe place:
`copy /intflash/config.cfg /intflash/config_backup.cfg`
`copy /intflash/config.cfg a.b.c.d:/dir/config_backup.cfg`

Example

Determine the configuration file names, save the configuration files, and copy the files to a safe place.

```
Switch:1>enable
Switch:1#show boot config choice
choice primary config-file "/intflash/config.cfg"
choice primary backup-config-file "/intflash/config.cfg"
Switch:1#save config
Switch:1#copy /intflash/config.cfg 00:11:f9:5b:10:42/dir/config_backup.cfg
Do you want to continue? (y/n)
y
```

System Access Configuration Using EDM

The section provides procedures you can use to manage system access by using Enterprise Device Manager (EDM). Procedures include configurations for usernames, passwords, and access policies.

Access the Switch Through the Web Interface

Before You Begin

- Note the IP address of the switch.
- Verify that you use a supported browser version.
- Enable the web server using CLI.

About This Task

Monitor the switch through a web browser from anywhere on the network. The web interface uses a 15-minute timeout period. If no activity occurs for 15 minutes, the system logs off the switch web interface, and you must re-enter the password information.

Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.



Note

By default the web server is configured with the secure-only option, which requires you to use HTTPS to access EDM. To access EDM using HTTP, you must disable the secure-only option. For more information about configuring the secure-only option, see [Enable the Web Management Interface](#) on page 303.

Procedure

1. Start your web browser.
2. In the address field of your web browser, enter the IP address of the system using one of the following formats: **https://<IP_address>** (default) or **http://<IP_address>**.
3. In **User Name**, type `admin` and in **Password**, type the password you configured when enabling the web server.
4. Select **Login**.

Configure the Minimum Version of the TLS Protocol

Use the following procedure to configure the minimum version of the TLS protocol.

Disable the web server before changing the TLS version. By disabling the web server, other existing users with a connection to the web server are not affected by changing to a different version.

The switch, by default, supports version TLS 1.2 and later. You can explicitly configure TLS 1.0 and TLS 1.1 version support.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Web** tab.

4. In **TlsMinimumVersion**, select the TLS version you want to configure as the minimum on the system.
5. Select **Apply**.

Web Field Descriptions

Use the data in the following table to use the **Web** tab.

Name	Description
WebRWAUserName	Specifies the RWA user name. The default is admin.
WebRWAUserPassword	Specifies the password for the RWA user account. The default is password.
WebROEnable	Enables the web server read-only (RO) user, which is disabled by default.
WebEncryptionType	Specifies the ciphers for preset version of TLS for the web server.
WebCertSubjectName	Specifies the digital certificate subject Name used as identity certificate in the web server.
WebCertCAName	Specifies the digital certificate CA trustpoint name used for the certificate in the web server.
WebROUserName	Specifies the RO user name. The default is user.
WebROUserPassword	Specifies the password for the RO user account. The default is 12345678.
MinimumPasswordLength	Configures the minimum password length. By default, the minimum password length is 8 characters.
HttpPort	Specifies the HTTP port for web access. The default value is 80.
HttpsPort	Specifies the HTTPS port for web access. The default value is 443.
SecureOnly	Controls whether the secure-only option is enabled. The default is enabled.
InactivityTimeout	Specifies the idle time (in seconds) to wait before the EDM login session expires. The default value is 900 seconds (15 minutes).
TlsMinimumVersion	Configures the minimum version of the TLS protocol supported by the web-server. You can select from the following options: <ul style="list-style-type: none"> • tlsv10 – Configures the version to TLS 1.0. • tlsv11 – Configures the version to TLS 1.1. • tlsv12 – Configures the version to TLS 1.2 The default TLS minimum version is tlsv12.

Name	Description
InUseCertType	Specifies whether the type of certificate in use is either user installed or self assigned.
SSLRenegotiation	Specifies whether SSL renegotiation is enabled or disabled. The default is enabled.
HelpTftp/Ftp_SourceDir	Configures the TFTP or FTP directory for Help files, in one of the following formats: a.b.c.d:/ peer:/ [<dir>]. The path can use 0–256 characters. The following example paths illustrate the correct format: <ul style="list-style-type: none"> • 192.0.2.1:/Help • 192.0.2.1:/
DefaultDisplayRows	Configures the number of rows each page displays. The default is 30.
LastChange	Shows the last web-browser initiated configuration change.
NumHits	Shows the number of hits to the web server.
NumAccessChecks	Shows the number of access checks performed by the web server.
NumAccessBlocks	Shows the number of access attempts blocked by the web server.
LastHostAccessBlockedAddressType	Shows the address type, either IPv4 or IPv6, of the last host access blocked by the web server.
LastHostAccessBlockedAddress	Shows the IP address of the last host access blocked by the web server.
NumRxErrors	Shows the number of receive errors the web server encounters.
NumTxErrors	Shows the number of transmit errors the web server encounters.
NumSetRequest	Shows the number of set-requests sent to the web server.

CLI Configuration Access using EDM

Use the following procedures to perform CLI access configuration tasks such as:

- Enable access levels
- [Change Passwords](#) on page 3428
- Configure the logon banner

Enable Access Levels

About This Task

Enable access levels to control the configuration actions of various users.



Important

Only the RWA user can disable an access level on the switch. You cannot disable the RWA access level on the switch.

The system preserves these configurations across restarts.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **CLI** tab.
4. Select the enable check box for the required access level.
5. Select **Apply**.

Configure the Logon Banner

About This Task

Configure the logon banner to display a warning message to users in the CLI before authentication.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **CLI** tab.
4. Type the banner text in **CustomBannerText**.
5. Select **CustomBannerEnable**.
6. Select **Apply**.

CLI Field Descriptions

The following table defines parameters for the **CLI** tab.

Name	Description
RWAUserName	Specifies the user name for the read-write-all CLI account.
RWAPassword	Specifies the password for the read-write-all CLI account.
RWEnable	Activates the read-write access. The default is enabled.
RWUserName	Specifies the user name for the read-write CLI account.
RWPassword	Specifies the password for the read-write CLI account.

Name	Description
RWL3Enable	Activates the read-write Layer 3 access. The default is enabled.
RWL3UserName	Specifies the user name for the Layer 3 read-write CLI account.
RWL3Password	Specifies the password for the Layer 3 read-write CLI account.
RWL2Enable	Activates the read-write Layer 2 access. The default is enabled.
RWL2UserName	Specifies the user name for the Layer 2 read-write CLI account.
RWL2Password	Specifies the password for the Layer 2 read-write CLI account.
RWL1Enable	Activates the read-write Layer 1 access. The default is enabled.
RWL1UserName	Specifies the user name for the Layer 1 read-write CLI account.
RWL1Password	Specifies the password for the Layer 1 read-write CLI account.
ROEnable	Activates the read-only CLI account. The default is enabled.
ROUserName	Specifies the user name for the read-only CLI account.
ROPassword	Specifies the password for the read-only CLI account.
MaxTelnetSessions	Specifies the maximum number of concurrent Telnet sessions in a range from 0–8. The default is 8.
Timeout	Specifies the number of seconds of inactivity for a Telnet before the system initiates automatic timeout and disconnect, expressed in a range from 30–65535. The default is 900 seconds.
NumAccessViolations	Indicates the number of CLI access violations detected by the system. This variable is a read-only field.
CustomBannerText	Specifies the text message that is displayed to users on the CLI before authentication. The message can be company information, such as company name and contact, or a warning message for the users of CLI. With character limitation from 1-1800, the text box displays 79 characters per line.
CustomBannerEnable	Specifies whether custom logon banner is enabled or disabled. The default is enabled.

Create an Access Policy

About This Task

Create an access policy to control access to the switch. An access policy specifies the hosts or networks that can access the switch through various services, such as Telnet, SNMP, HTTP, and SSH.

You can allow network stations to access the switch or forbid network stations to access the switch. For each service, you can also specify the level of access, such as read-only or read-write-all.

HTTP and HTTPS support IPv4 and IPv6 addresses.

If you configure the access policy mode to **deny**, the system checks the mode and service, and if they match the system denies the connection. With the access policy mode configured to **deny**, the system does not check **AccessLevel** and **AccessStrict** information. If you configure the access policy mode to allow, the system continues to check the **AccessLevel** and **AccessStrict** information.



Important

EDM does not provide SNMPv3 support for an access policy. If you modify an access policy with EDM, SNMPV3 is disabled.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Access Policies**.
3. Select the **Access Policies** tab.
4. Select **Insert**.
5. In **ID**, type the policy ID.
6. In **Name**, type the policy name.
7. Select **PolicyEnable**.
8. Select the **Mode** option to allow or deny a service.
9. From the **Service** options, select a service.
10. In **Precedence**, type a precedence number for the service (lower numbers mean higher precedence).
11. Select the **NetInetAddressType**.
12. In **NetInetAddress**, type an IP address.
13. In **NetInetAddressPrefixLen**, type the prefix length.
14. Select an **AccessLevel** for the service.
15. Select **AccessStrict**, if required.



Important

If you select **AccessStrict**, you specify that a user must use an access level identical to the one you select.

16. Select **Insert**.

Access Policies *Field Descriptions*

Use the data in the following table to use the **Access Policies** tab.

Name	Description
Id	Specifies the policy ID.
Name	Specifies the name of the policy.
PolicyEnable	Activates the access policy. The default is enabled.
Mode	Indicates whether a packet with a source IP address matching this entry is permitted to enter the device or is denied access. The default is allow. If you configure the access policy mode to deny , the system checks the mode and service, and if they match the system denies the connection. With the access policy mode configured to deny , the system does not check AccessLevel and AccessStrict information. If you configure the access policy mode to allow, the system continues to check the AccessLevel and AccessStrict information.
Service	Indicates the protocol to which this entry applies. The default is no service enabled.
Precedence	Indicates the precedence of the policy expressed in a range from 1–128. The lower the number, the higher the precedence. The default is 10.
NetInetAddrType	Indicates the source network Internet address type as one of the following. <ul style="list-style-type: none"> any IPv4 IPv6 IPv4 is expressed in the format a.b.c.d. Express IPv6 in the format x:x:x:x:x:x.
NetInetAddress	Indicates the source network Inet address (prefix/network). If the address type is IPv4, you must enter an IPv4 address and its mask length. If the type is IPv6, you must enter an IPv6 address. You do not need to provide this information if you select the NetInetAddrType of any .
NetInetAddrPrefixLen	Indicates the source network Inet address prefix-length/mask. If the type is IPv4, you must enter an IPv4 address and mask length. If the type is IPv6, you must enter an IPv6 address and prefix length. You do not need to provide this information if you select the NetInetAddrType of any .

Name	Description
TrustedHostInetAddress	<p>Indicates the trusted address of a host performing a login to the device. You do not need to provide this information if you select the NetInetAddressType of any.</p> <p>Important: You cannot use wild card entries in TrustedHostInetAddress.</p> <p>If the type is IPv4, you must enter an IPv4 address and mask length. If the type is IPv6, you must enter an IPv6 address and prefix length.</p>
AccessLevel	<p>Specifies the access level of the trusted host as one of the following:</p> <ul style="list-style-type: none"> • readOnly • readWrite • readWriteAll <p>The default is readOnly.</p>
Usage	Counts the number of times this access policy applies.
AccessStrict	<p>Activates or disables strict access criteria for remote users.</p> <p>If selected, a user must use an access level identical to the one you selected in the dialog box to use this service.</p> <ul style="list-style-type: none"> • selected (true): The system accepts only the currently configured access level. • cleared (false): The system accepts access up to the configured level. <p>Note: If you configure Mode to allow, the system checks AccessStrict information. If you configure Mode to deny, the system does not check AccessStrict information.</p> <p>Important: If you do not select true or false, user access is governed by criteria specified in the policy table. For example, a user with an rw access level specified for a policy ID in the policy table is granted rw access, and ro is denied access.</p> <p>The default is false (cleared).</p>

Enable an Access Policy

About This Task

Enable the access policy feature globally to control access across the switch.

You can create an access policy to control access to the switch. An access policy specifies the hosts or networks that can access the switch through access services; for example Telnet, SNMP, and Hypertext Transfer Protocol (HTTP).

Procedure

1. In the Device Physical View tab, select the Device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **System Flags** tab.
5. Select **EnableAccessPolicy**.
6. Select **Apply**.

Create Multiple CLI Users

You can create up to seven new CLI user roles on the switch, in addition to the three default CLI user roles. The user name must be unique. If you enable the hsecure flag, password complexity rules apply to all users.

Before You Begin

You must use an EDM account with read-write-all privileges to create new CLI users.

About This Task

Use this task to create multiple CLI users on the switch using EDM.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Multiple Users** tab.
4. Select **Insert**.
5. Type the ID.
6. Type a unique user name.
7. Type a password.
8. Select the access level.
9. Select **Enable** to activate the user account.
10. Select **Insert**.

Multiple Users Field Descriptions

Use the data in the following table to use the **Multiple Users** tab.

Name	Description
Id	Enter a unique integer between 1 and 10. The Id identifies this user on the switch
Name	Specifies the user name.

Name	Description
Password	Specifies the password the user must enter to authenticate. Note: The password can be up to 80 characters long and cannot contain the user name in either forward or reverse order.
UserLevel	Specifies the user access level. <ul style="list-style-type: none"> • ro—Read-only access • rw—Read/write access • rwa—Read/write all access
Enable	Check this check box to enables the user access on the switch.
Type	An autogenerated value that identifies the type of user as follows: <ul style="list-style-type: none"> • Default—The user is a system default user. • userDefined—The user was created by the admin user.

Modify CLI User Passwords

About This Task

Use this task to modify user account passwords using EDM.



Note

This procedure configures a temporary password that must be changed at first login. If you change your own password using this procedure, you will be automatically logged out.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Multiple Users** tab.
4. To change the user account password, double-click the **Password** to edit.
5. Select **Apply**.

Disable a CLI User Account

About This Task

Use this task to disable a user account using EDM.



Note

You cannot disable users with rwa access rights. Only users with ro and rw access rights can be disabled.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Multiple Users** tab.
4. In the **Enable** column, double-click the cell, and then select **false**.
5. Select **Apply**.

Delete a CLI User Account

About This Task

Use this task to delete a user account using EDM. You cannot delete default ro, rw, and rwa users.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Multiple Users** tab.
4. Select the row with the user account to delete, and then select **Delete**.
5. Select **Yes** to confirm.

Verification

Verify Boot Configuration Flags

Verify the boot configuration flags to verify boot configuration settings. Boot configuration settings only take effect after you reset the system. Verification of these parameters is essential to minimize system downtime and the resets to change them.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Verify the flags:
`show boot config flags`

Example



Note

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

Verify the Software Release

About This Task

Use CLI to verify your installed software. It is important to verify your software version before you place a device into a production environment.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Verify the software release:
show software detail

Example

The following is an example of the output of the `show software detail` command.

```
Switch:1>show software detail
=====
                        software releases in /intflash/release/
=====
5420.8.5.0.0.GA (Primary Release) (Signed Release)
SSIO
```

```

UBOOT                2.3.2.3
APP_FS               5420.8.5.0.0.GA
AVAILABLE ENCRYPTION MODULES
3DES
AES/DES

5420.8.4.0.0.GA (Backup Release) (Unsigned Release)
SSIO
UBOOT                2.3.2.1
APP_FS               5420.8.4.0.0.GA
AVAILABLE ENCRYPTION MODULES
3DES
AES/DES

```

Display Local Alarms

About This Task

View local alarms to monitor alarm conditions.

Local alarms are raised and cleared by applications running on the switch. Local alarms are an automatic mechanism run by the system that do not require any additional user configuration. The raising and clearing of local alarms also creates a log entry for each event. Check alarms occasionally to ensure no alarms require additional operator attention.

For more information, see [Alarm Database](#) on page 3797.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display local alarms:

```
show alarm database
```

Display Log Files

Use this procedure to display log files.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display log files:

```
show logging file
```

Example

Display log files:

```

Switch:1>show logging file
CP1  [02/05/15 12:35:28.690:UTC] 0x00270428 00000000 GlobalRouter SW INFO Lifecy
cle: Start
CP1  [02/05/15 12:35:29.906:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s sockserv started, pid:4950
CP1  [02/05/15 12:35:29.907:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom95 started, pid:4951
CP1  [02/05/15 12:35:29.907:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom90 started, pid:4952

```



```

CP1 [02/05/15 12:35:29.908:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s imgsync.x started, pid:4953
CP1 [02/05/15 12:35:30.346:UTC] 0x0026452f 00000000 GlobalRouter SW INFO No pat
ch set.
CP1 [02/05/15 12:35:30.909:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s logServer started, pid:4996
CP1 [02/05/15 12:35:30.910:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s trcServer started, pid:4997
CP1 [02/05/15 12:35:30.910:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oobServer started, pid:4998
CP1 [02/05/15 12:35:30.911:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s cbcp-main.x started, pid:4999
CP1 [02/05/15 12:35:30.912:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s rssServer started, pid:5000
CP1 [02/05/15 12:35:30.912:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s dbgServer started, pid:5001
CP1 [02/05/15 12:35:30.913:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s dbgShell started, pid:5002
CP1 [02/05/15 12:35:30.914:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s coreManager.x started, pid:5003
CP1 [02/05/15 12:35:30.914:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s ssio started, pid:5004
CP1 [02/05/15 12:35:30.915:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s hckServer started, pid:5005
CP1 [02/05/15 12:35:30.916:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s remCmdAgent.x started, pid:5006
CP1 [02/05/15 12:35:32.910:UTC] 0x000006cc 00000000 GlobalRouter SW INFO rcStar
t: FIPS Power Up Self Test SUCCESSFUL - 0
CP1 [02/05/15 12:35:32.911:UTC] 0x000006c2 00000000 GlobalRouter SW INFO rcStar
t: Security Stack Init SUCCESSFUL - 0
CP1 [02/05/15 12:35:32.911:UTC] 0x000006c3 00000000 GlobalRouter SW INFO rcStar
t: IPSEC Init SUCCESSFUL
CP1 [02/05/15 12:35:32.911:UTC] 0x000006bf 00000000 GlobalRouter SW INFO rcStar
t: Security Stack Log init SUCCESSFUL - 0
IO1 [02/05/15 12:35:35.177:UTC] 0x0011054a 00000000 GlobalRouter COP-SW INFO De
tected Master CP in slot 1

--More-- (q = quit)

```

Enhanced Secure Mode Configuration Using CLI

The section provides information on security enhancements after you enable enhanced secure mode.

Display the Boot Config Flags Status

Use the following procedure to display the boot config flags status.

If enhanced secure mode is enabled, the status displays whether the JITC or non-JITC sub-mode is enabled. If enhanced secure mode is disabled, the status displays as false.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. View the boot flags status:

```
show boot config flags
```

Example

In the following example, the status displays that enhanced secure mode is disabled.



Note

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

Enable Enhanced Secure Mode

Use the following procedure to enable enhanced secure mode. Enhanced secure mode is disabled by default.

About This Task



Note

When you migrate your switch from enhanced secure mode enabled to disabled, or from disabled to enabled, you must build a new configuration. Do not use a configuration created in either enhanced secure mode disabled or enabled, and expect it to transfer over to the new mode.

The configuration file cannot be guaranteed if you transfer between enhanced secure mode enabled to disabled, or from enhanced secure mode disabled to enabled.

After you enable the enhanced secure mode, the system provides role-based access levels, stronger password requirements, and stronger rules on password length, password complexity, password change intervals, password reuse, and password

maximum age use. The enhanced secure mode boot flag supports two sub-modes namely JITC and non-JITC.

After you disable enhanced secure mode, the authentication, access-level, and password requirements work similarly to any of the existing commercial releases.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable enhanced secure mode:

```
boot config flags enhancedsecure-mode [jitc | non-jitc]
```



Note

As a best practice, enable the enhanced secure mode in the non-JITC sub-mode, because the JITC sub-mode is more restrictive and prevents the use of some CLI commands that are commonly used for troubleshooting.

3. (Optional) Disable enhanced secure mode:

```
no boot config flags enhancedsecure-mode
```

4. (Optional) Configure the enhanced secure mode to the default value:

```
default boot config flags enhancedsecure-mode
```

5. Save the configuration:

```
save config
```



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

6. Restart the switch:

```
boot [config WORD<1-99>] [-y]
```



Note

If you enter the **boot** command with no arguments, you cause the switch to start using the current boot choices defined by the **boot config choice** command.

If you enter a boot command and the configuration file name without the directory, the device uses the configuration file from `/intflash/`.

Example

Enable the enhanced secure non-JITC sub-mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags enhancedsecure-mode non-jitc
```

```
Switch:1(config)#save config
Switch:1(config)#exit
Switch:1(config)#boot config /intflash/config.cfg -y
```

Enable the enhanced secure JITC sub-mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags enhancedsecure-mode jitc
Switch:1(config)#save config
Switch:1(config)#exit
Switch:1(config)#boot config /intflash/config.cfg -y
```

What to Do Next

If you upgrade a device that operated in non-Enhanced Secure Mode (ESM) to a newer release, and then enable ESM, the switch retains the previously saved configuration. After you enable ESM, the system continues to use these inherited settings unless you reset them.

For customers that operate in environments governed by strict security or compliance frameworks, such as FIPS, the switch must load the ESM-specific default settings. These defaults ensure all cryptographic components adhere to the stricter controls mandated in ESM.

Regardless of release version or when the switch has been upgraded in relation to enabling ESM, it is always a best practice to load the correct ESM-compliant defaults. After you enable ESM, reset several security-related parameters to ensure the switch applies the updated values. After you enable ESM, enter the following commands:

- **default ssh encryption-type**
- **default ssh key-exchange-method**
- **default syslog tls-min-ver**

Save the configuration to preserve the settings across reboots.

Variable definitions

Use the data in the following table to use the **boot config flags enhancedsecure-mode** command.

Variable	Value
<i>jitc</i>	Enables the JITC enhanced secure mode. The JITC mode is more restrictive and prevents the use of some CLI commands that are commonly used for troubleshooting.
<i>non-jitc</i>	Enables the non-JITC enhanced secure mode.

Create Accounts for Different Access Levels

Use the following procedure to create accounts for different access levels in enhanced secure mode. You must be the administrator to configure the different access levels.

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create accounts on the switch for different access levels:

```
password create-user {auditor|operator|privilege|security} WORD<1-255>
```

3. Save the configuration:

```
save config
```



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Create an account at the auditor level for jsmith:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password create-user auditor jsmith
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password create-user** command.

Variable	Value
{auditor operator privilege security}	Specifies the access level for the user.
WORD<1-255>	Specifies the user name.

Delete Accounts in Enhanced Secure Mode

Use the following procedure to delete accounts in enhanced secure mode.

Before You Begin

- You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.
- You must be an admin or privilege user to delete accounts.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Delete an account on the switch:
`password delete-user username WORD<1-255>`
3. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Delete an account for jsmith:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password delete-user user-name jsmith
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password delete-user** command.

Variable	Value
<code>user-name WORD<1-255></code>	Specifies the user name.

Configure a Password for a Specific User

Configure a new password for a user if the password has expired or locked. Only the administrator can configure a password for a user.

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

About This Task



Note

The **password set-password** command configures a temporary password that must be changed at first login. If you change your own password using this command, you will be automatically logged out.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create accounts on the switch for different access levels:
`password set-password user-name WORD<1-255>`
3. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure a password for jsmith:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password set-password user-name jsmith
Enter the New password : *****
Switch:1(config)#Password modified for user jsmith
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password set-password** command.

Variable	Value
<code>user-name WORD<1-255></code>	Specifies the user for which to configure the password.

Return the System to the Factory Defaults

To return the system to factory defaults, run the **boot config flags factorydefaults** command.

For more information about the factory defaults boot flag, see the following sections:

- [Boot Sequence](#) on page 247
- [Enable Factory Default Behavior](#) on page 363
- [Configure Boot Flags](#) on page 775

Configure the Password Complexity Rule

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

About This Task

Use the following procedure to configure the password complexity rule.

The password complexity rule default is to use at least two uppercase, two lowercase, two numeric, and two special character to meet the password criteria.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the password complexity rule:
`password password-rule <1-2> <1-2> <1-2> <1-2>`
3. (Optional) Configure the password complexity rule to the default:
`default password password-rule`
4. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure the password complexity rule to require two uppercase, two lowercase, two numeric and two special characters in each password:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password password-rule 2 2 2 2
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password password-rule** command.

Variable	Value
<1-2> <1-2> <1-2> <1-2>	Configures the minimum password rule. The first variable defines the number of uppercase characters required. The second <1-2> variable defines the number of lowercase characters required. The third <1-2> variable defines the number of numeric characters required. The fourth <1-2> variable defines the number of special characters required. The default for each of these is 2.

Configure the Password Length Rule

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

About This Task

Configure the password length rule after you enable enhanced secure mode. By default, the minimum password length is 15.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the password length rule option:

```
password min-passwd-len <8-32>
```
3. (Optional) Configure the password length rule to the default:

```
default password min-passwd-len
```
4. Save the configuration:

```
save config
```



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure the password length rule to 20:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password min-passwd-len 20
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password min-passwd-len** command.

Variable	Value
<8-32>	Configures the minimum character length required. The default is 15.

Configure the Change Interval Rule

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

About This Task

Use the following procedure to configure the change interval rule. The system enforces a minimum password change interval, which defines the minimum amount of time before you can change to a new password. By default, the minimum change interval is 24 hours between changing from one password to a new password.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the change interval rule option:
`password change-interval <1-999 hours>`
3. (Optional) Configures the change interval rule to the default:
`default password change-interval`
4. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure the change interval rule to 72 hours:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password change-interval 72
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password change-interval** command.

Variable	Value
<1-999>	Configures the minimum interval between consecutive password changes. The default is 24 hours.

Configure the Reuse Rule

Use the following procedure to configure the password reuse rule. The default password reuse rule is 3.

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the password reuse rule option:
`password password-history <3-32>`
3. (Optional) Configure the password reuse rule to the default:
`default password password-history`
4. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure the reuse rule to 30:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password password-history 30
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password password-history** command.

Variable	Value
<3-32>	Configures the minimum number of previous passwords to remember. The default is 3.

Configure the Maximum Number of Sessions

Use the following procedure to configure the maximum number of sessions on the switch. The `max-sessions` value configures the number of times a particular role-based user can log in to the switch through the SSH session at the same time. The default `max-sessions` value is 3.

The `max-sessions` value applies only for SSH sessions, and only with enhanced secure mode enabled.

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the maximum number of sessions:
`password max-sessions <1-8> user-name WORD<1-255>`
3. (Optional) Configure the password reuse rule to the default:
`default password max-sessions`
4. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure the reuse rule to 5:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password max-sessions 5 user-name jsmith
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password max-sessions** command.

Variable	Value
<code><1-8></code>	Specifies the maximum number of sessions. The default is 3.
<code>user-name WORD<1-255></code>	Specifies the user-name.

Configure the Maximum Age Rule

Use the following procedure to configure the maximum age rule.

If enhanced secure mode is enabled, the individual with the administrator access level role can configure the aging-time for each user. If you configure the aging time for each user, the aging time must be more than the global change interval value. The default is 90 days.

If you do not enable enhanced secure mode, the aging time is a global value for all users.

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the maximum age rule option:
`password aging-time day <1-365> [user WORD<1-255>]`
3. (Optional) Configure the maximum age rule to the default:
`default password aging-time [user WORD<1-255>]`
4. Save the configuration:
`save config`



Note

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Example

Configure the maximum age rule option to 100 days for user jsmith:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password aging-time day 100 user jsmith
Switch:1(config)#save config
```

Variable definitions

Use the data in the following table to use the **password aging-time** command.

Variable	Value
<i>day</i> <1-365>	Configures the password aging time in days. The default is 90 days.
<i>user</i> <i>WORD</i> <1-255>	Specifies a particular user.

Configure the Pre-Notification and Post-Notification Rule

Use the following procedure to configure the pre-notification and post-notification rule.

After enhanced secure mode is enabled, the switch enforces password expiry. To ensure a user does not lose access, the switch offers pre- and post-notification messages explaining when the password will expire.

The administrator can define pre- and post-notification intervals to between one to 99 days.

Before You Begin

You must enable enhanced secure mode in either the JITC or non-JITC sub-modes. As a best practice, use the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

About This Task

The pre-notification intervals provide messages to warn users that their passwords will expire within a particular time frame:

The post-notification intervals provide notification to users that their passwords have expired within a particular time frame:

- interval 1—By default, interval 1 is 1 day.
- interval 2—By default, interval 2 is 7 days.
- interval 3—By default, interval 3 is 30 days.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the pre-notification rule option:
`password pre-expiry-notification-interval <1-99> <1-99> <1-99>`
3. Configure post-notification rule option:
`password post-expiry-notification-interval <1-99> <1-99> <1-99>`

4. Save the configuration:

```
save config
```

**Note**

This command saves the configuration using the file name configured as the primary configuration. Use the command **show boot config choice** to view the current primary and backup configuration file names.

Variable definitions

Use the data in the following table to use the **pre-expiry-notification-interval** command.

Variable	Value
<1-99> <1-99> <1-99>	<p>Configure the pre-notification intervals to provide messages to warn the users that their passwords will expire within a particular timeframe.</p> <p>The first <1-99> variable specifies the first notification, the second <1-99> specifies the second notification, and the third <1-99> variable specifies the third interval.</p> <p>By default, the first interval is 30 days, the second interval is 7 days, and the third interval is 1 day.</p>

Use the data in the following table to use the **post-expiry-notification-interval** command.

Variable	Value
<1-99> <1-99> <1-99>	<p>Configure the post-notification intervals to provide notification to the users that their passwords have expired within a particular timeframe.</p> <p>The first <1-99> variable specifies the first notification, the second <1-99> specifies the second notification, and the third <1-99> variable specifies the third interval.</p> <p>By default, the first interval is 1 day, the second interval is 7 days, and the third interval is 30 days.</p>

Enhanced Secure Mode Configuration Using EDM

The section provides information to enable enhanced secure mode.

Enable Enhanced Secure Mode

About This Task

Use the following procedure to enable enhanced secure mode in either the JITC or non-JITC sub-modes.

The enhanced secure mode is disabled by default.

After you enable enhanced secure mode, the system can provide role-based access levels, stronger password requirements, and stronger rules on password length, password complexity, password change intervals, password reuse, and password maximum age use.

After you disable enhanced secure mode, the authentication, access-level, and password requirements work similarly to any of the existing commercial releases.

**Note**

You can use EDM to enable or disable enhanced secure mode. To configure the security enhancements this feature provides, you must use CLI.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **Boot Config** tab.
4. In **EnableEnhancedsecureMode**, select either **jitc** or **non-jitc** to enable the enhanced secure mode in one of these sub-modes. Select **disable** to disable the enhanced secure mode.

**Note**

As a best practice, enable the non-JITC sub-mode. The JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.

5. Select **Apply**.
6. Save the configuration, and restart the switch.

Basic Administration Procedures using CLI

The following section describes common procedures that you use while you configure and monitor the switch operations using the Command Line Interface (CLI).

**Note**

Unless otherwise stated, to perform the procedures in this section, you must log on to the Privileged EXEC mode in the CLI. For more information about how to use CLI, see [CLI Procedures](#) on page 193.

Restart the Platform

Before You Begin**Note**

The command mode is key for this command. If you are logged on to a different command mode, such as Global Configuration mode, rather than Privileged EXEC mode, the system displays different options for this command.

About This Task

Restart the switch to implement configuration changes or recover from a system failure. When you restart the system, you can specify the boot config file name. If you do not specify a boot source and file, the boot command uses the configuration files on the primary boot device defined by the `boot config choice` command.

After the switch restarts normally, it sends a cold trap within 45 seconds after the restart.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Restart the switch:
`boot [config WORD<1-99>] [-y]`



Important

If you enter the `boot` command with no arguments, you cause the switch to start using the current boot choices defined by the `boot config choice` command.

If you enter a boot command and the configuration file name without the directory, the device uses the configuration file from `/intflash/`.

Variable Definitions

The following table defines parameters for the **boot** command.

Variable	Value
<code>config WORD<1-99></code>	Specifies the software configuration device and file name in one of the following formats: <ul style="list-style-type: none"> • <code>/intflash/ <file></code> The file name, including the directory structure, can include up to 99 characters.
<code>-y</code>	Suppresses the confirmation message before the switch restarts. If you omit this parameter, you must confirm the action before the system restarts.

Reset the Platform

About This Task

Reset the platform to reload system parameters from the most recently saved configuration file.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Reset the switch using one of the following commands:
 - Reset the switch immediately: `reset [-y]`
 - Reset the switch after a delay: `reset -in WORD<1-7>`
 - Reset the switch after you create an ssio core file and a cbcx-main.x core file:
`reset -coredump [-y]`
3. Cancel a delayed reset:
`reset -cancel`



Note

To cancel the delayed reset, you must issue the **reset -cancel** command before the configured timer expires. You can use the **show software** command to view the remaining time before the switch resets.

Example

Reset the switch with a configured delay of 45 minutes:

```
Switch:1>enable
Switch:1#reset -in 45
Are you sure you want to perform a delayed reset? (y/n) ? y
```

Reset the switch with a configured delay of 3 hours and 44 minutes:

```
Switch:1>enable
Switch:1#reset -in 3:44
Are you sure you want to perform a delayed reset? (y/n) ? y
```

Reset the switch with a configured delay of 7 days, 1 hour, and 7 minutes:

```
Switch:1>enable
Switch:1#reset -in 7:01:07
Are you sure you want to perform a delayed reset? (y/n) ? y
```

Cancel the reset:

```
Switch:1#reset -cancel
Scheduled reset successfully canceled!
```

View the remaining time before the switch reboots:

```
Switch:1#show software
=====
                        software releases in /intflash/release/
=====
Switch.8.6.0.0.GA          (Backup Release)          (Signed Release)
Switch.9.0.0.0.GA          (Primary Release)          (Signed Release)
-----
Auto Commit      : enabled
Commit Timeout   : 10 minutes
Reset in        : 3 days, 2 hours, 1 minute and 57 seconds
```

Variable Definitions

The following table defines parameters for the **reset** command.

Variable	Value
<code>-cancel</code>	Cancels the delayed switch reset.
<code>WORD<1-7></code>	<p>Specifies the time interval, in days, hours, or minutes to wait before the switch resets. After the timer expires, the device restarts from the saved configuration file.</p> <p>The syntax for the timer is [d:] [hh:]mm where days and hours are optional. If you specify a single number, like 45, the switch interprets that as minutes. If you specify two numbers separated by a colon (3:15), the switch interprets that as hours and minutes. If you specify three numbers separated by colons (4:3:15), the switch interprets that as days, hours, and minutes.</p> <p>Values for hours and minutes must use two digits. For example, for 7 hours and 1 minute, use 07:01 and not 7:1. You can specify a value of 00 for hours or minutes; 2:00:00 indicates a reset in 2 days.</p> <p>The maximum value is 9:23:59 (9 days, 23 hours, and 59 minutes).</p>
<code>-y</code>	Suppresses the confirmation message before the switch resets. If you omit this parameter, you must confirm the action before the switch resets.

Shut Down the System

Use the following procedure to shut down the system.



Caution

Before you unplug the power cord, always perform the following shutdown procedure.

This procedure:

- Flushes any pending data to ensure data integrity.
- Ensures the completion of recent configuration save actions, thus preventing the system from inadvertently booting up with incorrect configuration.

Procedure

1. Enter Privileged EXEC mode:
`enable`

2. Shut down the system:

```
sys shutdown
```

3. Before you unplug the power cord, wait until you see the following message:

```
System Halted, OK to turn off power
```

Example

Shut down a running system.

```
Switch:1#sys shutdown
Are you sure you want shutdown the system? Y/N (y/n) ? y
CP1 [05/08/14 15:47:50.164] 0x00010813 00000000 GlobalRouter HW INFO System shutdown
initiated from CLI
CP1 [05/08/14 15:47:52.000] LifeCycle: INFO: Stopping all processes
CP1 [05/08/14 15:47:53.000] LifeCycle: INFO: All processes have stopped
CP1 [05/08/14 15:47:53.000] LifeCycle: INFO: All applications shutdown, starting power
down sequence
INIT: Sending processes the TERM signal
Stopping OpenBSD Secure Shell server: sshdno /usr/sbin/sshd found; none killed
Stopping vsp...Error, do this: mount -t proc none /proc
done
sed: /proc/mounts: No such file or directory
sed: /proc/mounts: No such file or directory
sed: /proc/mounts: No such file or directory
Deconfiguring network interfaces... done.
Stopping syslogd/klogd: no syslogd found; none killed
Sending all processes the TERM signal...
Sending all processes the KILL signal...
/etc/rc0.d/S25save-rtc.sh: line 5: /etc/timestamp: Read-only file system
Unmounting remote filesystems...
Stopping portmap daemon: portmap.
Deactivating swap...
Unmounting local filesystems...
[24481.722669] Power down.
[24481.751868] System Halted, OK to turn off power
```

Configure the Default Ping and Traceroute Context

About This Task

Ping commands and traceroute commands execute in Global Router (GRT) context by default. You can configure ping commands and traceroute commands to execute in management (mgmt) context or in Virtual Router Forwarding (vrf) context.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the default ping command and traceroute command context:

```
sys default-ping-context {grt | mgmt | vrf}
```

Variable Definitions

The following table defines parameters for the **sys default-ping-context** command.

Variable	Value
grt	Specifies Global Routing Table (grt) context as the default context for ping commands and traceroute commands. The default configuration is grt as the default context.
mgmt	Specifies management (mgmt) context as the default context for ping commands and traceroute commands. The default configuration is grt as the default context.
vrf	Specifies Virtual Router Forwarding (VRF) context as the default context for ping commands and traceroute commands. The default configuration is grt as the default context.

Calculate the File Checksum

About This Task

Perform the following procedure to calculate or compare the MD5 or SHA512 digest for a specific file. The **file-checksum** command calculates the MD5 or SHA512 digest for files on the internal flash and either shows the output on screen or stores the output in a file that you specify. The **file-checksum** command compares the calculated MD5 or SHA512 digest with that in a checksum file on flash, and the compared output displays on the screen. By verifying the MD5 or SHA512 checksum, you can verify that the file is transferred properly to the switch.



Important

- If the MD5 key file parameters change, you must remove the old file and create a new file.
- Use the **file-checksum** command with reserved files (for example, a password file) only if you possess sufficient permissions to access these files.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Calculate the file checksum:

```
file-checksum {md5 | sha512} WORD<1-99> [-a] [-c] [-f WORD<1-99>] [-r]
```

Example

```
Switch:1>file-checksum md5 password -a -f password.md5
```

Variable Definitions

The following table defines parameters for the **file-checksum** command.

Variable	Value
<i>md5</i>	Calculates or compares the MD5 digest for a specific file.
<i>sha512</i>	Calculates or compares the SHA512 digest for a specific file.
<i>-a</i>	Adds data to the output file instead of overwriting it. You cannot use the <i>-a</i> option with the <i>-c</i> option.
<i>-c</i>	Compares the checksum of the specified file with the MD5 checksum present in the checksum file name. You can specify the checksum file name using the <i>-f</i> option. If the checksum file name is not specified, the file <code>/intflash/checksum.md5</code> is used for comparison. If the supplied checksum file name and the default file are not available on flash, the system displays the following error message on the switch: Error: Checksum file <i><file name></i> not present. The <i>-c</i> option also <ul style="list-style-type: none"> calculates the checksum of the specified files compares the checksum with all keys in the checksum file, even if file names do not match displays the output of comparison
<i>-f</i>	Stores the result of MD5 checksum to a file on internal flash. If the output file specified with the <i>-f</i> option is reserved file names on the switch, the command fails with the error message: Error: Invalid operation. If the output file specified with the <i>-f</i> option is files for which to compute MD5 checksum, the command fails with the error message: <pre>Switch:1# md5 *.cfg -f config.cfg Error: Invalid operation on file <file name></pre> If the checksum file name specified by the <i>-f</i> option exists on the switch (and is not one of the reserved file names), the system displays the following message on the switch: File exists. Do you wish to overwrite? (y/n)

Variable	Value
<code>-r</code>	Reverses the output. Use with the <code>-f</code> option to store the output to a file. You cannot use the <code>-r</code> option with the <code>-c</code> option.
<code>WORD<1-99></code>	Specifies the file name.

Reset System Functions

About This Task

Reset system functions to reset all statistics counters on the console port.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Reset system functions:
`sys action reset {console|counters}`

Example

Reset the statistics counters:

```
Switch:1>enable
Switch:1#sys action reset counters
Are you sure you want to reset system counters (y/n)?
y
```

Variable Definitions

The following table defines parameters for the **sys action** command.

Variable	Value
<code>reset {console counters}</code>	Reinitializes the hardware universal asynchronous receiver transmitter (UART) drivers. Use this command only if the console connection does not respond. Resets all the statistics counters in the switch to zero. Resets the console port.

Source a Configuration

Source a configuration to merge a script file into the running configuration or verify the syntax of a configuration file.

About This Task

The **source** cli command is intended for use with a switch that is running with a factory default configuration to quick load a pre-existing configuration from a file. If you source a configuration file to merge that configuration into a running configuration, it can result in operational configuration loss if the sourced configuration file contains any

configuration that has dependencies on or conflicts with the running configuration. Use the `source` command to merge smaller portions of a configuration into the existing configuration.

Not all CLI commands are included in configuration files. Typical examples include, but are not limited to some operational and security-related commands. Ensure that you understand what configuration options are included or not included in a configuration file, when you use that file to build new configurations.

The operational modes in the boot configuration file must be configured for some features (for example, `spbm-config-mode true/false`). Before sourcing a configuration file, you need to configure the `boot config flag`, save the configuration, and reboot the system. After the reboot, you can source the configuration file without fail.



Important

Do not source a verbose configuration (`verbose.cfg`) with the `debug stop` option. The sourcing process cannot complete if you use these two options with a verbose configuration.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Source a configuration:
`source WORD<1-99> [debug] [stop] [syntax]`

Example

Debug the script output:

```
Switch:1>enable
Switch:1#source testing.cfg debug
```

Variable Definitions

The following table defines parameters for the `source` command.

Variable	Value
<code>debug</code>	Debugs the script by outputting the configuration commands to the screen.
<code>stop</code>	Stops the sourcing of a configuration if an error occurs.

Variable	Value
<i>syntax</i>	<p>Checks the syntax of the configuration file. This parameter does not load the configuration file; only verifies the syntax. If you use this parameter with the <i>stop</i> parameter (source WORD<1-99> stop syntax), the output displays on screen and verification stops if it encounters an error.</p> <p>If you use this parameter with the <i>debug</i> parameter (source WORD<1-99> debug syntax), the output does not stop if it encounters an error; you must review the on-screen output to verify if an error exists. If you use this parameter by itself, it does not output to the screen or stop on error; it shows an error message, <i>syntax errors in script</i>, to indicate if errors exist in the configuration file.</p>
<i>WORD<1-99></i>	<p>Specifies a file name and location in one of the following formats:</p> <ul style="list-style-type: none"> • a.b.c.d:<file> • /intflash/<file> <p><file> is a string.</p>

Using the USB Device

The following sections describe common procedures that you can use with the USB device.



Note

The switch supports USB devices formatted as FAT32 and NTFS. The switch does not support the exFAT format.

Save a File to an External USB Device

Use the following procedure to save the configuration file or log file to an external USB device.



Caution

Always use the `usb-stop` command to safely unplug the USB drive from the USB slot.

Procedure

1. Enter Privileged EXEC mode:
`enable`

2. Save the file to an external USB device:

- a. To save the configuration file to an external USB device, enter:

```
save config file WORD<1-99>
```

- b. To save the log file to an external USB device, enter:

```
save log file WORD<1-99>
```

Example

```
Switch:1#save config file /usb/test.cfg
CP-1: Save config to file /usb/test.cfg successful.
WARNING: Choice Primary Node Config file is "/intflash/soak.cfg".

Switch:1#
Switch:1#save log file /usb/test.log

Save log to file /usb/test.log successful.
Save log to file /usb/test.log successful.
Switch:1#
```

Variable definitions

The following table defines parameters for the `save` command.

Variable	Value
config file WORD<1-99>	Specifies the software configuration device and configuration file name in one of the following formats: <ul style="list-style-type: none"> a.b.c.d:<file> /intflash/<file> /usb/<file> The file name, including the directory structure, can include up to 99 characters.
log file WORD<1-99>	Specifies the software configuration device and log file name in one of the following formats: <ul style="list-style-type: none"> a.b.c.d:<file> /intflash/<file> /usb/<file> The file name, including the directory structure, can include up to 99 characters.

Back Up and Restore the Compact Flash to an External USB Device


Perform this procedure to back up and restore the contents of the internal compact flash to a USB flash device without entering multiple **copy** commands. This procedure is useful if you want to copy the complete compact flash contents to another chassis.



Caution

Always use the `usb-stop` command to safely unplug the USB drive from the USB slot.

Before You Begin

-  **Important**
Disable logging using the command: `no boot config logging`.
- You must have a USB storage device ready to use that is at least 2 GB. The switch supports USB 1 and 2.

About This Task

The system verifies that the USB flash device has enough available space to perform the backup operation. If the USB flash device does not have enough available space, the system displays an error message. The backup command uses the following filepath on the USB flash device: `/usb/intflash/intflashbackup_yyyymmddhhmmss.tgz`.

The backup action can take up to 10 minutes.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Backup the internal flash to USB:
`backup intflash`
3. Restore the data to the internal flash:
`restore intflash`

Example

Switch:1#`backup intflash`

```
Warning: Command will backup all data from /intflash to /usb/intflash.
It will take a few minutes and may cause high CPU utilization.

Are you sure you want to continue? (y/n) ? y

For file system /intflash:
  7252475904 total bytes on the filesystem
  990920704 used bytes on the filesystem
  6261555200 free bytes on the filesystem

For file system /usb:
  2021216256 total bytes on the filesystem
  12038144 used bytes on the filesystem
  2009178112 free bytes on the filesystem

cd /intflash ; /bin/tar -czvf /usb/intflash/intflashbackup_20140610074501.tgz *
; /bin/sync

Info: Backup /intflash to filename /usb/intflash/intflashbackup_20140610074501.tgz is
complete!

Do you want to stop the usb? (y/n) ? n
```

Copy Configuration and Log Files from a USB Device to Intflash

Copy configuration and log files from an external USB device to the internal Flash memory.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Copy configuration or log files from the USB device to Intflash:
`copy /usb/<srcfile> /intflash/<destfile>`

Variable Definitions

The following table defines parameters for the `copy` command.

Variable	Value
<destfile>	Specifies the name of the configuration or log file when copied to the internal Flash memory. The destination file name must be lower case and have a file extension of .cfg or .log. For example, test.cfg or test.log. The file name, including the directory structure, can include up to 255 characters.
<srcfile>	Specifies the name of the configuration or log file on the USB device. For example, test.cfg or test.log. The file name, including the directory structure, can include up to 255 characters.

Display the Contents of a USB File

Use the following procedure to view content of a USB file.



Caution

Always use the `usb-stop` command to safely unplug the USB drive from the USB slot.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display content of a USB file:
`more WORD<1-99>`

Variable definitions

The following table defines parameters for the `more` command.

Variable	Value
<code>WORD<1-99></code>	Specifies the file name in the following format: <ul style="list-style-type: none"> <code>/usb/<file></code> The file name, including the directory structure, can include up to 99 characters.

Move a File to or from a USB Device

Use the following procedure to move a file from the internal Flash memory (Intflash) to an external USB device, or from a USB device to Intflash.



Caution

Always use the `usb-stop` command to safely unplug the USB drive from the USB slot.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Move a file to a safe location:
 - a. To move a file from Intflash to a USB device:


```
mv /intflash/<srcfile> /usb/<destfile>
```
 - b. To move a file from a USB device to Intflash:


```
mv /usb/<srcfile> /intflash/<destfile>
```

Example

```
Switch:1#enable
Switch:1#mv /intflash/test.cfg /usb/test.cfg
Switch:1#mv /usb/test.cfg /intflash/test.cfg
```

Variable Definitions

The following table defines parameters for the `mv` command.

Variable	Value
<code><destfile></code>	Specifies the name of the configuration or log file when moved to the USB device. The destination file name must be lower case and have a file extension of <code>.cfg</code> or <code>.log</code> . For example, <code>test.cfg</code> or <code>test.log</code> . The file name, including the directory structure, can include up to 255 characters.
<code><srcfile></code>	Specifies the name of the configuration or log file on the internal flash memory. For example, <code>test.cfg</code> or <code>test.log</code> . The file name, including the directory structure, can include up to 255 characters.

Delete a file from a USB Device

Use the following procedure to delete a file from an external USB device.

**Caution**

Always use the `usb-stop` command to safely unplug the USB drive from the USB slot.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Delete a file from a USB device:
`delete WORD<1-255>`

Variable Definitions

The following table defines parameters for the `delete` command.

Variable	Value
<code>WORD<1-255></code>	Specifies the file name in the following format: • <code>/usb/<file></code>

Back Up Configuration Files to ZIP

Table 41: ExtremeCloud IQ Site Engine backup configuration ZIP file product support

Feature	Product	Release introduced
ExtremeCloud IQ Site Engine backup configuration ZIP file For more information, see ExtremeCloud IQ Site Engine documentation.	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

ExtremeCloud IQ Site Engine has a configuration backup feature with a requirement to be able to backup configuration related files.

**Note**

License files are not backed up.

*Back Up Configuration Files to a ZIP File***About This Task**

Use this procedure to back up configuration files.

**Important**

Only the RWA user can use the **backup** command.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Use the backup command:
`backup configure WORD<1-99>`

Example

```
Switch:1>enable
Switch:1#backup configure /intflash/backup02072018

Successfully backed up config /intflash to /intflash/backup02072018.tgz
```

*Restore Configuration Files from a ZIP File***Before You Begin**

- Download the backup file to the `/intflash` directory.
- If restoring the configuration files on a new switch, you must do one of the following:
 - Disable IS-IS on the old switch.
 - Power the old switch down.
 - Remove the old switch from the network.
- If restoring the configuration files on a different switch, use the **isis dup-detection-temp-disable** command on the new switch to suspend duplicate detection prior to its insertion into the existing SPBM topology.

**Important**

This must be done after the original unit has been completely removed or isolated from the SPBM topology.

About This Task

Use the following procedure to restore previously backed up configuration files.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Run the restore command to restore the configuration files.
`restore configure WORD<1-99>`

Example

```
Switch:1>enable
Switch:1#restore configure /intflash/backup02072018.tgz

Warning: Command will restore your backup setup and access files
The current files will be overwritten.

Are you sure you want to continue? (y/n) ?y

Restore /intflash from /intflash/backup02072018.tgz is complete!
Reboot is required for the new configuration to be effective
```

Basic Administration Procedures Using EDM

The following section describes common procedures that you use while you configure and monitor the switch operations using Enterprise Device Manager (EDM).

Reset the Platform

Reset the platform to reload system parameters from the most recently saved configuration file. Use the following procedure to reset the device using EDM.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System** tab.
4. From **ActionGroup1**, select **saveRuntimeConfig**.
5. Select **Apply**.
6. From **ActionGroup4**, select **softReset**.
7. Select **Apply**.

Show the MTU for the System

About This Task

Perform this procedure to show the MTU configured for the system.

Procedure

1. On the Device Physical View, select the device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **Chassis** tab.

5. Verify the selection for the MTU size.

Save the Configuration

About This Task

After you change the configuration, you must save the changes on the device. Save the configuration to a file to retain the configuration settings.



Note

When you log out of the EDM interface, a dialog box automatically prompts if you want to save the configuration. If you want to save the configuration, select **Yes**. If you want to close without saving the configuration, select **No**. If you no longer see the prompt, clear your browser cache, restart your browser, and then reconnect.

Procedure

1. In the Device Physical View tab, select the device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **System** tab.
5. (Optional) Specify a file name in **ConfigFileName**.
If you do not specify a file name, the system saves the information to the default file.
6. In **ActionGroup1**, select **saveRuntimeConfig**.
7. Select **Apply**.

View UBoot Version and Status

About This Task

Use this information to verify the integrity of the software and hardware. A representative from Customer Support will instruct you to obtain this information when required.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **UBoot** tab.

UBoot Field Descriptions

Use the data in the following table to use the **UBoot** tab.

Name	Description
Version	Displays the UBoot version for products that support only one UBoot image. These products do not display DefaultVersion , AlternateVersion , or VersionUsed .
DefaultVersion	Displays the default UBoot version.
AlternateVersion	Displays the alternate UBoot version.
VersionUsed	Displays the UBoot version in use, either the default or alternate.
TrustedDeliveryStatus	Displays the status of the chain of trust image signature validation steps.

Boot Parameter Configuration Using the CLI

Use the procedures in this section to configure and manage the boot process.

Modify the Boot Sequence

About This Task

Modify the boot sequence to prevent the switch from using the factory default settings or, conversely, to prevent loading a saved configuration file.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Perform one of the following actions:

- Bypass the loading of the switch configuration file and load the factory defaults:

```
boot config flags factorydefaults [config-only | reset-all-files |  
zero-touch | zero-touch-config-only]
```

- Use a configuration file and not the factory defaults:

```
no boot config flags factorydefaults
```



Important

If the switch fails to read and load a saved configuration file after it starts, check the log file to see if the log file indicates that the factorydefaults setting was enabled, before you investigate other options.

Example

Modifying the boot sequence with an optional behavior of the factory default boot flag.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags factorydefaults
Switch:1(config)#boot config flags factorydefaults config-only
Warning: If you try to reset with a new image, it will not be loaded.
        To undo this operation use "no boot config flags factorydefaults"
```

Configure the Remote Host Logon

Before You Begin

- The FTP server must support the FTP passive (PASV) command. If the FTP server does not support the passive command, the file transfer is aborted, and then the system logs an error message that indicates that the FTP server does not support the passive command.

About This Task

Configure the remote host logon to modify parameters for FTP and TFTP access. The defaults enable TFTP transfers. If you want to use FTP as the transfer mechanism, you need to change the password to a non-null value.



Important

tftp-debug should be used exclusively to transfer small files less than 1MB in size. Using it for larger files might cause unwanted behavior, such as transfer failure.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Define conditions for the remote host logon:

```
boot config host {ftp-debug|password WORD<0-16>|tftp-debug|tftp-hash|tftp-rexmit <1-120>|tftp-timeout <1-120>|user WORD<0-16>}
```
3. Save the changed configuration.

Example

Enable console tftp/tftpd debug messages:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config host tftp-debug
Switch:1(config)#save config
```

Change the Primary or Secondary Boot Configuration Files

About This Task

Change the primary or secondary boot configuration file to specify which configuration file the system uses to start.

Configure the primary boot choices.

You have a primary configuration file that specifies the full directory path and a secondary configuration file that also contains the full directory path.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Change the primary boot choice:

```
boot config choice primary {backup-config-file|config-file} WORD<0-255>
```
3. Save the changed configuration.
4. Restart the switch.

Example

Specify the configuration file in internal flash memory as the primary boot source:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config choice primary config-file /intflash/config.cfg
Switch:1(config)#save config
Switch:1(config)#reset
```

Variable Definitions

The following table defines parameters for the **boot config** command.

Variable	Value
<i>{backup-config-file config-file}</i>	Specifies that the boot source uses either the configuration file or a backup configuration file.
<i>WORD<0-255></i>	Identifies the configuration file. <i>WORD<0-255></i> is the device and file name, up to 255 characters including the path, in one of the following format: <ul style="list-style-type: none"> • a.b.c.d:<file> • /usb/<file> • /intflash/<file>

Display the Boot Configuration

About This Task

Display the configuration to view current or changed settings for the boot parameters.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the configuration:

```
show boot config <choice|flags|general|host|master|running-config  
[verbose]|sio>
```

Example

Show the current boot configuration. (If you omit verbose, the system only displays the values that you changed from their default value.):

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1#(config)#show boot config running-config
#
#Mon Feb 13 13:32:58 2017 EST
#
boot config flags debug-config file
boot config flags debugmode
boot config flags ftpd
no boot config flags spbm-config-mode
boot config flags sshd
boot config flags telnetd
boot config flags tftpd
no boot config flags verify-config
boot config choice primary backup-config-file "/intflash/config.cfg"
#boot config sio console baud 115200
```

Variable Definitions

The following table defines parameters for the **show boot config** command.

Variable	Value
<i>choice</i>	Shows the current boot configuration choices.
<i>flags</i>	Shows the current flag settings.
<i>general</i>	Shows system information.
<i>host</i>	Shows the current host configuration.
<i>master</i>	Shows the master information.

Variable	Value
<i>running-config</i> <i>[verbose]</i>	Shows the current boot configuration. If you use <i>verbose</i> , the system displays all possible information. If you omit <i>verbose</i> , the system displays only the values that you changed from their default value.
<i>sio</i>	Specifies the current configuration of the serial ports.

Configure Serial Port Devices

Configure the serial port devices to define connection settings for the console port.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. View the current baud rate configuration:
`show boot config sio`
3. Change the console baud rate:
`boot config sio console baud <115200>`
4. Save the changed configuration.
5. Restart the switch.

Variable Definitions

The following table defines parameters for the **boot config sio console** command.

Variable	Value
<i>baud <115200></i>	Configures the baud rate for the port. The default is 115200 and cannot be modified.

Boot Flag Configuration using CLI

Configure the boot configuration flags to enable specific services and functions for the chassis.



Note

Flag support can vary across hardware models.

The following tasks document boot config flags included in other sections of this document:

- [Enable Remote Access Services](#) on page 293
- [Configure a DvR Leaf](#) on page 948
- [Increase the Number of SPB Nodes per Area](#) on page 1253

- [Enable the IPv6–Mode Boot Config Flag](#) on page 2191
- [Enable hsecure](#) on page 3386
- [Enable urpf-mode Boot Flag](#) on page 3394
- [Block SNMP Access](#) on page 3487
- [Configure the Spanning Tree Mode](#) on page 3506
- [Enable Enhanced Secure Mode](#) on page 324
- [Enable IPv6 Egress Filters](#) on page 3719
- [Configure the Maximum Number of VRFs](#) on page 4089
- [Enable VXLAN Gateway in Full Interworking Mode](#) on page 4183

Reserve Bandwidth for Advanced Features

Enables the switch to support advanced features by reserving ports as loopback ports. When disabled, you can use all ports on the switch, but advanced features do not work.

When you enable this boot flag, you need to save and reboot with the new configuration.

Before You Begin

You must ensure your configuration does not include reserved ports before you enable this feature. If the configuration includes reserved ports after you enable this feature and restart the switch, the switch stops loading the configuration.

About This Task

The command parameters apply to this procedure as follows:

- `high` applies to 4220 Series, 5720 Series, 7520 Series, and 7720 Series
- `low` does not apply to 5720 Series
- `vim` is only supported on 5520 Series and 5720 Series

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable the boot configuration flag:

```
boot config flags advanced-feature-bandwidth-reservation [low | high | vim]
```

3. Save the configuration, and then reboot the switch.



Important

A change to the `advanced-feature-bandwidth-reservation` boot flag requires a reboot for the change to take effect.

4. Verify the boot flag configuration:

```
show boot config flags
```

5. Verify that the switch reserved the ports as loopback ports. Reserved ports are not visible in the output of the following command:

```
show interfaces gigabitEthernet
```

Example

Enable this feature to the low level.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags advanced-feature-bandwidth-reservation low
Warning: Please note that the configuration for the following ports 1/25-1/26
        will be removed from the configuration file.
Are you sure you want to continue (y/n) ? y
Warning: Please save the configuration and reboot the switch
        for this to take effect.
Flag advanced-feature-bandwidth-reservation is changed to enable (low).
```

Display Advanced Feature Bandwidth Reservation Ports

After you configure the **advanced-feature-bandwidth-reservation** boot flag and reboot with the new configuration, you can use the following procedure to verify that the switch reserved ports for configuring advanced features.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the Advanced Feature Bandwidth Reservation mode and reserved ports:

```
show sys-info
```

Example

```
Switch#show sys-info

General Info :

SysDescr      : Switch1 (w.x.y.z)   BoxType: Switch1
SysName       : Switch1
.
.
.

Advanced Feature Bandwidth Reservation:
-----

Reservation Mode : low
Port Usage Info  : 1/53 and 1/54 are not available to use
```

Enable Debugging of the Configuration File During System Start

Enables you to debug the configuration file during loading configuration at system boot up. The default is disabled. You do not have to restart the switch after you enable debug-config, unless you want to immediately debug the configuration. After you enable debug-config and save the configuration, the debug output either displays on the console or logs to an output file the next time the switch reboots.

About This Task

The options are:

- `debug-config [console]`—Displays the line-by-line configuration file processing and result of the execution on the console while the device loads the configuration file.
- `debug-config [file]`— Logs the line-by-line configuration file processing and result of the execution to the debug file while the device loads the configuration file. The system logs the debug config output to `/intflash/debugconfig_primary.txt` for the primary configuration file. The system logs the debug config output to `/intflash/debugconfig_backup.txt` for the backup configuration, if the backup configuration file loads.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable debugging of the configuration file:
`boot config flags debug-config [console] | [file]`
3. Save the changed configuration.

Enable Debug Mode

Enables a TRACE on any port by prompting the selection on the console during boot up. This allows the user start trace for debugging earlier on specified port. Works on console connection only. The default is disabled.



Important

Do not change this parameter unless directed by Technical Support.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable debug mode:
`boot config flags debugmode`
3. Save the changed configuration.
4. Restart the switch.

Enable Factory Default Behavior

Specifies whether the switch uses the factory default settings at startup. The default value is disabled. This flag is automatically reset to the default value after the CPU restarts. If you change this parameter, you must restart the switch.

About This Task

The factory default boot configuration flag supports the following behaviors:

Table 42: Factory default boot configuration flag parameters

Parameter	Action
no parameter used	Removes the configuration files, primary and secondary config file names, all user account passwords, digital certificates, IKE/OSPF/IS-IS keys, and SNMP communities. All ports are disabled and assigned to VLAN 1. License files are not removed.
<i>config-only</i>	Boots the switch with a blank configuration. This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities. All ports are disabled and assigned to VLAN 1. Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.
<i>reset-all-files</i>	Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset. If required, you must revoke the license file. The 30-day factory license does not apply to 4220 Series. Note: You can also use the unconfigure switch command to achieve the same behavior.

Table 42: Factory default boot configuration flag parameters (continued)

Parameter	Action
<code>zero-touch</code>	Boots the switch with a default configuration that enables Auto-sense. This parameter resets secure files but keeps the security mode and performs a ZTP+ configuration after reboot. License files are not removed.
<code>zero-touch-config-only</code>	<p>Boots the switch with a blank configuration that enables Auto-sense and performs a ZTP+ configuration. This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities.</p> <p>Note: Feature support can vary across platforms. For more information, see <i>Fabric Engine and VOSS Feature Matrix</i>.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p>

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable factory default behavior:

```
boot config flags factorydefaults [config-only | reset-all-files | zero-touch | zero-touch-config-only]
```
3. Save the changed configuration.
4. Restart the switch.

Examples

The following examples display the factory default flag behaviors.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags factorydefaults config-only
Warning: If you try to reset with a new image, it will not be loaded.
        To undo this operation use "no boot config flags factorydefaults".

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags factorydefaults reset-all-files
Warning: Factorydefaults flag will delete all config (*.cfg) files on /intflash/,
        as well as local password files, authentication keys, and certificates.
        Also, if you try to reset with a new image, it will not be loaded.
```

```

Enhanced Secure and HiSecure modes will also be reset.
Ports will be configured with auto-sense enabled.
License files will also be removed on non-Universal-Hardware.
To undo this operation use "no boot config flags factorydefaults".

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags factorydefaults zero-touch
Warning: Factorydefaults flag will delete runtime, primary and backup config files,
        as well as local password files, authentication keys, and certificates.
        Also, if you try to reset with a new image, it will not be loaded.
        Ports will be configured with auto-sense enabled.
        To undo this operation use "no boot config flags factorydefaults".

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags factorydefaults zero-touch-config-only
Warning: If you try to reset with a new image, it will not be loaded.
        Ports will be configured with auto-sense enabled.
        To undo this operation use "no boot config flags factorydefaults".

```

Return an Existing Configuration to Factory Defaults

Boots the switch with a factory default configuration.

Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset. If required, you must revoke the license file. The 30-day factory license does not apply to 4220 Series.



Note

unconfigure switch and **boot config flags factorydefaults reset-all-files** commands have the same behavior. For more information, see [Enable Factory Default Behavior](#) on page 363.

About This Task

When you use the **unconfigure switch** command, it changes the value of the factory default flag in the **show boot config flags** command output to **flags factorydefaults reset-all-files**.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable factory default behavior:

```
unconfigure switch
```
3. Save the changed configuration.
4. Restart the switch.

Example

Enable the factory default behavior using the **unconfigure switch** command.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#unconfigure switch
Warning: Factorydefaults flag will delete all config (*.cfg) files on /intflash/,
        as well as local password files, authentication keys, and certificates.
        Also, if you try to reset with a new image, it will not be loaded.
        Enhanced Secure and HiSecure modes will also be reset.
        Ports will be configured with auto-sense enabled.
        License files will also be removed on non-Universal-Hardware.
        To undo this operation use "no boot config flags factorydefaults".
```

Enable Flow Control

Enables or disables flow control globally. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames.

The default is disabled.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable flow control:
`boot config flags flow-control-mode`
3. Save the changed configuration.
4. Restart the switch.

Enable System Logging

About This Task

The system names log files according to the following:

- The system displays the file names in 8.3 (log.xxxxxxxx.sss) format.
- The first 6 characters of the file name contain the last three bytes of the chassis base MAC address.
- The next two characters in the file name specify the slot number of the CPU that generated the logs.
- The last three characters in the file name are the sequence number of the log file.

The system generates multiple sequence numbers for the same chassis and same slot if the system reaches the maximum log file size.

The default value is enabled.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable system logging:
`boot config flags logging`
3. Save the changed configuration.

Enable MACsec

Enables Media Access Control Security (MACsec) globally.



Note

Exception: only required for 5320 Series and 5420 Series.

About This Task

The **boot config flags ipv6-egress-filter** and **boot config flags macsec** commands are mutually exclusive.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable MACsec:
`boot config flags macsec`
3. Save the changed configuration.
4. Restart the switch.

Enable MSTP and VLAN Configuration on NNI ports

Enables MSTP and VLAN configuration on network-to-network interface (NNI) ports. The default is disabled.



Note

This procedure does not apply to 4220 Series.

About This Task



Note

Spanning Tree is disabled on all NNIs.

You cannot add an SPBM NNI port or MLT port to any non SPBM B-VLAN. You cannot add additional C-VLANs to a brouter port.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the configuration on NNI ports:
`boot config flags nni-mstp`
3. Save the changed configuration.
4. Restart the switch.

Enable Automatic Reboot

Activates or disables automatic reboot on a fatal error. The default value is activated.



Important

Do not change this parameter unless directed by technical support.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable automatic reboot:
`boot config flags reboot`
3. Save the changed configuration.

Enable SPBM and IS-IS Configuration on the Switch

Enables you to configure SPB and IS-IS, but you cannot configure PIM and IGMP, either globally or on an interface.

About This Task

Use the **no** operator if you need to configure PIM and IGMP.

The boot flag is enabled by default.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable SPBM and IS-IS configuration:
`boot config flags spbm-config-mode`
3. Save the changed configuration.
4. Restart the switch.

Configure the Syslog Output Format

Controls the format of the syslog output and logging. By default, the switch uses the RFC 5424 format. If the RFC-based format is disabled, an older format is used.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure the output format:
`boot config flags syslog-rfc5424-format`
3. Save the changed configuration.

Enable Trace Logs

Activates or disables the creation of trace logs. The default value is disabled.



Important

Do not change this parameter unless directed by technical support.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable trace logs:
`boot config flags trace-logging`
3. Save the changed configuration.

Disable Syntax Checking of the Configuration File

Disables syntax checking of the configuration file. The default is enabled.



Note

As a best practice, disable the verify-config flag.

About This Task

- Primary config behavior: When the verifyconfig flag is enabled, the primary config file is pre-checked for syntax errors. If the system finds an error, the primary config file is not loaded, instead the system loads the backup config file.

If the verify-config flag is disabled, the system does not pre-check syntax errors. When the verify-config flag is disabled, the system ignores any lines with errors during loading of the primary config file. If the primary config file is not present or cannot be found, the system tries to load the backup file.

- Backup config behavior: If the system loads the backup config file, the system does not check the backup file for syntax errors. It does not matter if the verify-config flag is disabled or enabled. With the backup config file, the system ignores any lines with errors during the loading of the backup config file.

If no backup config file exists, the system defaults to factory defaults.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Disable syntax checking of the configuration file:
`no boot config flags verify-config`
3. Save the changed configuration.
4. Restart the switch.

Run-Time Process Management Using CLI

Configure the Time Zone

About This Task

Configure the time zone to use an internal system clock to maintain accurate time. The time zone data in Linux includes daylight changes for all time zones up to the year 2100. You do not need to configure daylight savings.

The default time zone is Coordinated Universal Time (UTC).

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure the time zone by using the following command:
`clock time-zone WORD<1-10> WORD<1-20> WORD<1-20>`
3. Save the changed configuration.

Example

Configure the system to use the time zone data file for Vevay:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clock time-zone America Indiana Vevay
```

Variable Definitions

The following table defines parameters for the **clock time-zone** command.

Variable	Value
<i>WORD<1-10></i>	Specifies a directory name or a time zone name in /usr/share/zoneinfo, for example, Africa, Australia, Antarctica, or US. To see a list of options, enter <code>clock time-zone</code> at the command prompt without variables.
<i>WORD<1-20> WORD<1-20></i>	The first instance of <i>WORD<1-20></i> is the area within the timezone. The value represents a time zone data file in /usr/share/zoneinfo/ <i>WORD<1-10></i> /, for example, Shanghai in Asia. The second instance of <i>WORD<1-20></i> is the subarea. The value represents a time zone data file in /usr/share/zoneinfo/ <i>WORD<1-10></i> / <i>WORD<1-20></i> /, for example, Vevay in America/Indiana. To see a list of options, enter <code>clock time-zone</code> at the command prompt without variables.

Configure the CLI Run-time Environment

About This Task

Configure the run-time environment to define generic configuration for CLI sessions.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Change the login prompt:
`login-message WORD<1-1513>`
3. Change the password prompt:
`passwordprompt WORD<1-1510>`
4. Configure the number of supported inbound Telnet sessions:
`telnet-access sessions <0-8>`
5. Configure the idle timeout period before automatic logoff for CLI and Telnet sessions:
`cli timeout <30-65535>`
6. Configure the number of lines in the output display:
`terminal length <8-64>`
7. Configure scrolling for the output display:
`terminal more <disable|enable>`

Examples

```
Switch:1>enable
Switch:#configure terminal
```

Use the default option to enable use of the default logon string:

```
Switch:(config)#default login-message
```

Use the default option before this parameter to enable use of the default string:

```
Switch:(config)#default passwordprompt
```

Configure the allowable number of inbound Telnet sessions:

```
Switch:(config)#telnet-access sessions 8
```

Configure the timeout value, in seconds, to wait for a Telnet or CLI login session before terminating the connection:

```
Switch:(config)#cli timeout 900
```

Configure the number of lines in the output display for the current session:

```
Switch:(config)#terminal length 30
```

Configure scrolling for the output display:

```
Switch:(config)#terminal more disable
```

Variable Definitions

The following table defines parameters for the **login-message** command.

Variable	Value
<i>WORD</i> <1-1513>	Changes the CLI logon prompt.

The following table defines parameters for the **passwordprompt** command.

Variable	Value
<i>WORD</i> <1-1510>	Changes the CLI password prompt.

The following table defines parameters for the **telnet-access sessions** command.

Variable	Value
<0-8>	Configures the allowable number of inbound Telnet sessions. The default value is 8.

The following table defines parameters for the **cli time-out** command.

Variable	Value
<30-65535>	Configures the timeout value, in seconds, to wait for a Telnet or CLI login session before terminating the connection.

The following table defines parameters for the **terminal** command.

Variable	Value
<8-64>	Configures the number of lines in the output display for the current session. The default is value 23.
<i>disable enable</i>	Configures scrolling for the output display. The default is enabled.

Configure System Parameters

About This Task

Configure individual system-level switch parameters to configure global options for the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Change the system name:

```
sys name WORD<0-255>
```
3. Enable support for Jumbo frames:

```
sys mtu <sys-mtu>
```
4. Enable the User Datagram Protocol (UDP) checksum calculation:

```
udp checksum
```

Variable Definitions

The following table defines parameters for the **sys** command.

Variable	Value
<i>control tcp-timestamp</i>	Enables or disables TCP Timestamp.
<i>force-msg</i>	Adds forced message control pattern. <i>WORD<4-4></i> Enter force message pattern.
<i>msg-control</i>	Configures system message control feature.

Variable	Value
<i>mtu {1522 9416 9600}</i>	Specifies the Jumbo frame size for the data path. Note: The 9416 MTU size applies to 7830 Series only. The default is 1950 bytes.
<i>name WORD<0-255></i>	Configures the system, or root level, prompt name for the switch. <i>WORD<0-255></i> is an ASCII string from 0-255 characters (for example, LabSC7 or Closet4).
<i>power</i>	Enables power to specified slot(s).
<i>security-console</i>	Enables the security console.
<i>software</i>	Configures software configuration.
<i>priv-exec-password</i>	Enables authentication for the Privileged EXEC CLI command mode.

Configure System Message Control

Configure system message control to suppress duplicate error messages on the console, and to determine the action to take if they occur.

About This Task

To enable the message control feature, you must specify an action, control interval, and maximum message number. After you enable the feature, the log messages, which get repeated and cross the maximum message number in the control interval, trigger the force message feature. You can either suppress the message or send a trap notification, or both.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure system message control action:

```
sys msg-control action <both|send-trap|suppress-msg>
```
3. Configure the maximum number of messages:

```
sys msg-control max-msg-num <2-500>
```
4. Configure the interval:

```
sys msg-control control-interval <1-30>
```
5. Enable message control:

```
sys msg-control
```

Example

Configure system message control to suppress duplicate error messages on the console and send a trap notification. Configure the number of occurrences of a message after which the control action occurs. Configure the message control interval in minutes. Finally, enable message control.

```
Switch:1(config)#sys msg-control action suppress-msg
Switch:1(config)#sys msg-control max-msg-num 10
Switch:1(config)#sys msg-control control-interval 15
Switch:1(config)#sys msg-control
```

What to Do Next

[Extend System Message Control](#) on page 376

Variable Definitions

The following table defines parameters for the **sys msg-control** command.

Variable	Value
<i>action</i> <both send-trap suppress-msg>	Configures the message control action. You can either suppress the message or send a trap notification, or both. The default is suppress.
<i>control-interval</i> <1-30>	Configures the message control interval in minutes. The valid options are 1–30. The default is 5.
<i>max-msg-num</i> <2-500>	Configures the number of occurrences of a message after which the control action occurs. To configure the maximum number of occurrences, enter a value from 2–500. The default is 5.

Extend System Message Control

About This Task

Use the force message control option after you enable the message control feature. Force message control stops sending the log messages starting with a configured pattern. This happens regardless if message control parameters (control interval or maximum message number) are accomplished. The action to suppress the message, send a trap notification, or both, is indicated by system message control.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the force message control option:
`sys force-msg WORD<4-4>`

Example

Add a force message control pattern. If you use a wildcard pattern (****), all messages undergo message control.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#sys force-msg ****
```

Variable Definitions

The following table defines parameters for the **sys force-msg** command.

Variable	Value
<i>WORD</i> <4-4>	Adds a forced message control pattern, where <i>WORD</i> <4-4> is a string of 4 characters. You can add a four-byte pattern into the force-msg table. The software and the hardware log messages that use the first four bytes that match one of the patterns in the force-msg table undergo the configured message control action. You can specify up to 32 different patterns in the force-msg table, including a wildcard pattern (****) as well. If you specify the wildcard pattern, all messages undergo message control.

Run-Time Process Management Using EDM

Configure the Date and Time

Configure the date and time to correctly identify when events occur on the system.

Procedure

1. On the Device Physical View, select the device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **User Set Time** tab.
5. Type and select the correct details.
6. Select **Apply**.

User Set Time Field Descriptions

Use the data in the following table to use the **User Set Time** tab.

Name	Description
Year	Specifies the year (1998–2100). The default is 1998.
Month	Specifies the month. The default is 1.
Date	Specifies the day (integer 1–31). The default is 1.

Name	Description
Hour	Specifies the hour (12am–11pm). The default is 0.
Minute	Specifies the minute (integer 0–59). The default is 0.
Second	Specifies the second (integer 0–59). The default is 0.
Time Zone	Specifies the time zone. The default time zone is Coordinated Universal Time (UTC).

NEW Configure System Message Control

Configure system message control to suppress duplicate error messages on the console, and to determine the action to take if they occur.

About This Task

To enable the message control feature, you must specify an action, control interval, and maximum message number. After you enable the feature, the log messages, which get repeated and cross the maximum message number in the control interval, trigger the force message feature. You can either suppress the message or send a trap notification, or both.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **System Log**.
3. Select the **Msg Control** tab.
4. (Optional) Select the action: either suppress the message or send a trap notification, or both.
5. (Optional) Configure the interval.
6. (Optional) Configure the maximum number of messages.
7. Enable message control.
8. Select **Apply**.

What to Do Next

[Configure a Forced Message Control Pattern](#) on page 379

Msg Control Field Descriptions

Use the data in the following table to use the **Msg Control** tab.

Name	Description
Action	Configures the message control action. You can either suppress the message or send a trap notification, or both. The default is suppress (suppressMsg).
Interval	Configures the message control interval in minutes. The default is 5.

Name	Description
MaxMsgNum	Configures the maximum number of occurrences of a message after which the control action occurs. The default is 5.
Enable	Enables system message control if duplicate error messages occur on the console. The default is disabled.

Configure a Forced Message Control Pattern

About This Task

Use the force message control option after you enable the message control feature. Force message control stops sending the log messages starting with a configured pattern. This happens regardless if message control parameters (control interval or maximum message number) are accomplished. The action to suppress the message, send a trap notification, or both, is indicated by system message control.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Force Msg Patterns** tab.
3. Select **Insert**.
4. In **PatternId**, type a pattern ID number.
5. In **Pattern**, type a message control pattern.
6. Select **Insert**.

Force Msg Patterns Field Descriptions

Use the data in the following table to use the **Force Msg Patterns** tab.

Name	Description
PatternId	Specifies a pattern identification number in the range 1–32.
Pattern	Specifies a forced message control pattern of 4 characters. The software and the hardware log messages that use the first four bytes matching one of the patterns in the force-msg table undergo the configured message control action. You can specify up to 32 different patterns in the force-msg table, including a wildcard pattern (***). If you specify the wildcard pattern, all messages undergo message control.

Hardware Status Using EDM

This section provides methods to check the status of basic hardware in the chassis using Enterprise Device Manager (EDM).

Configure Polling Intervals

About This Task

Enable and configure polling intervals to determine how frequently EDM polls for port and LED status changes, or detects the hot swap of installed ports.

Procedure

1. In the navigation pane, expand **Configuration > Device**.
2. Select **Preference Setting**.
3. Enable polling or hot swap detection.
4. Configure the frequency to poll the device.
5. Select **Apply**.

Preference Setting *Field Descriptions*

Use the data in the following table to use the **Preference Setting** tab.

Name	Description
Enable	Enables polling for port and LED status changes. The default is disabled.
Poll Interval	Specifies the polling interval, if enabled. The default is 60 seconds.
Enable	Detects the hot swap of installed ports. The default is disabled.
Detection per Status Poll Intervals	Specifies the number of poll intervals for detection, if enabled. The default is 2 intervals.

View Module Information

About This Task

This tab is not available for hardware platforms with only fixed configuration; it is available for platforms where you can install modules in slots, for example, a Versatile Interface Module (VIM).

Procedure

1. In the Device Physical View tab, select the module slot.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Card**.
4. Select the **Card** tab.

Card Field Descriptions

Use the data in the following table to use the **Card** tab.

Name	Description
CardType	Displays the model number of the module.
SerialNum	Displays the serial number for the installed module.
PartNumber	Displays the part number.
CardAssemblyDate	Displays the date the module was assembled.
HWVersion	Displays the hardware revision.
AdminStatus Note: Exception: only applies to 7830 Series.	Enables or disables the VIM.
OperStatus Note: Exception: only applies to 7830 Series.	Displays the operational status of the VIM.
SlotPower Note: Exception: only applies to 7830 Series.	Enables or disables power on the VIM slot.

View Power Supply Parameters

Perform this procedure to view information about the operating status of the power supplies.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Power Supply**.
3. Select the **Detail** tab.

Detail Field Descriptions

Use the data in the following table to use the **Detail** tab.

Name	Description
Type	Describes the type of power used. Note: To know what power types the switch supports, see the hardware installation documentation.
Description	Provides a description of the power supply.

Name	Description
SerialNumber	Specifies the power supply serial number.
HardwareRevision	Specifies the hardware revision number.
PartNumber	Specifies the power supply part number.
PowerSupplyOperStatus	Specifies the status of the power supply as one of the following: <ul style="list-style-type: none"> • unknown • empty • up • down
InputLineVoltage	Displays the input line voltage: <ul style="list-style-type: none"> • low110v—power supply connected to a 110 Volt source • high220v—power supply connected to a 220 Volt source • minus48v—power supply connected to a DC source • ac110vOr220v—power supply connected to a 110 Volt or 220 Volt source • dc—power supply connected to a DC power source
OutputWatts	Displays the output power of this power supply.
VoltageIn Note: Exception: Switches do not report this information for fixed power supply units.	Specifies the input voltage value in millivolts (mV) for the power supply unit.
Note: Exception: Switches do not report this information for fixed power supply units.	Specifies the output voltage value in millivolts (mV) for the power supply unit.
CurrentIn Note: Exception: Switches do not report this information for fixed power supply units.	Specifies the input current value in milliamps (mA) for the power supply unit.
CurrentOut Note: Exception: Switches do not report this information for fixed power supply units.	Specifies the output current value in milliamps (mA) for the power supply unit.

Name	Description
PowerIn Note: Exception: Switches do not report this information for fixed power supply units.	Specifies the input power value in milliwatts (mW) for the power supply unit.
PowerOut Note: Exception: Switches do not report this information for fixed power supply units.	Specifies the output power value in milliwatts (mW) for the power supply unit.
FanFlowType Note: Exception: only supported on 5520 Series	Displays the airflow direction of the power supply.

View System Temperature Information

View information about the temperature for each sensor on the device.

The system triggers an alarm when one of the zones exceeds the threshold temperature value.

Procedure

1. In the Device Physical View tab, select the chassis.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **System Temperature** tab.

System Temperature Field Descriptions

Use the data in the following table to use the **System Temperature** tab.

Name	Description
SensorIndex	Specifies the range of sensors on the device.
SensorDescription	Specifies the name of the sensor. This field identifies the inlet sensor for Energy Star certification on supporting platforms. For more information about Energy Star certification, see <i>Fabric Engine and VOSS Feature Matrix</i> .
Temperature (degrees celsius)	Specifies the sensor temperature measured in Celsius degrees.

Name	Description
WarningThreshold	Specifies the temperature value of the warning threshold for the sensor. When the temperature crosses the warning threshold a warning message is generated.
CriticalThreshold	Species the temperature value of the critical threshold for the sensor. When the temperature crosses the critical threshold, a critical message is generated or the system shuts down, depending on hardware capability.
Status	Specifies the current temperature status based on the warning and critical thresholds.



Address Resolution Protocol

[Address Resolution Protocol](#) on page 386

[Reverse Address Resolution Protocol](#) on page 387

[ARP configuration using the CLI](#) on page 388

[ARP configuration using Enterprise Device Manager](#) on page 397

Table 43: Address Resolution Protocol product support

Feature	Product	Release introduced
Address Resolution Protocol (ARP) including Proxy ARP and Static ARP	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Gratuitous ARP filtering	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Address Resolution Protocol

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host IP address, the network station uses Address Resolution Protocol (ARP) to determine the physical address for a network host by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts.

The network station uses ARP to determine the host physical address as follows:

- The network station broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.
- All network hosts receive the broadcast request.
- Only the specified host responds with its hardware address.
- The network station then maps the host IP address to its physical address and saves the results in an address-resolution cache for future use.
- The network station ARP table displays the associations of the known MAC address to IP address.

You can create ARP entries, and you can delete individual ARP entries.

Enable ARP traffic

The switch accepts and processes ARP traffic, spanning tree bridge packet data units (BPDU), and Topology Discovery Protocol packets on port-based VLANs with the default port action of drop. If a filter port action is drop for a packet, ARP packets are also dropped. As a result, ARP entries on that port are cleared and are not relearned when the ARP aging timer expires.

To prevent dropped ARP packets, configure the following options:

- A user-defined protocol-based VLAN for ARP EtherType (byprotocol usrDefined 0x0806).
- Ports as static members to this VLAN with the default port action of drop.
- The port default VLAN ID to the correct port-based VLAN where the ARPs are processed.

You do not need to make configuration changes for the BPDU and Topology Discovery Protocol packets.

Only one user-defined protocol-based VLAN for ARP is allowed for each Spanning Tree Group (STG). If the ports with the default port action of drop are in different STGs, you must create additional user-defined protocol-based VLANs.

Proxy ARP

A network station uses proxy ARP to respond to an ARP request from a locally attached host or end station for a remote destination. The network station sends an ARP response back to the local host with its own MAC address of the network station

interface for the subnet on which the ARP request was received. The reply is generated only if the device has an active route to the destination network.

The following figure shows an example of proxy ARP operation. In this example, the system displays host C with mask 24 to be locally attached to host B with mask 16, so host B sends an ARP request for host C. However, the switch is between the two hosts. To enable communication between the two hosts, the switch responds to the ARP request with the IP address of host C but with its own MAC address.

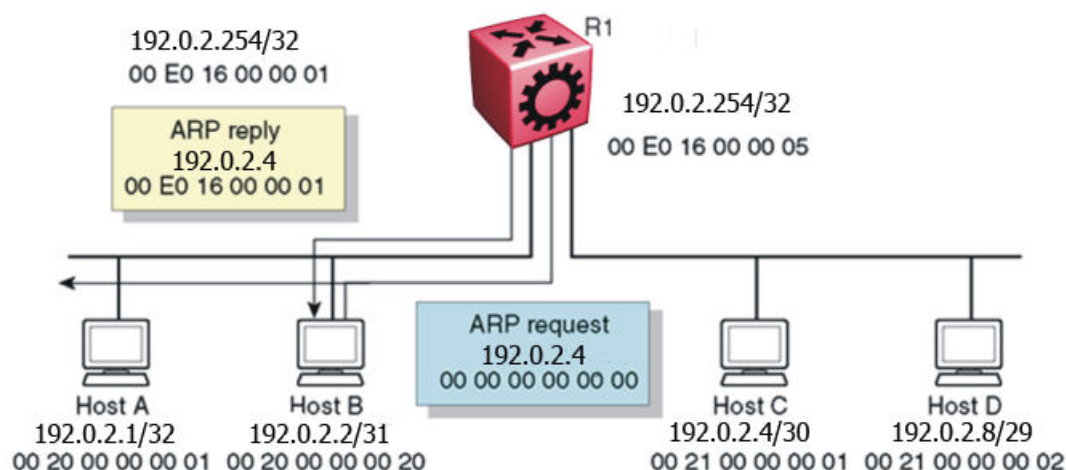


Figure 17: Proxy ARP operation

Loop detection

To prevent cases of ARP looping, configure the ARP loop detection flag to detect this situation. When a loop is detected, the port is shut down.

Flushing router tables

For administrative or troubleshooting purposes, sometimes you must flush the routing tables. Flush routing tables either by VLAN or by port. In a VLAN context, all entries associated with the VLAN are flushed. In a port context, all entries associated with the port are flushed.

Reverse Address Resolution Protocol

Certain devices use the Reverse Address Resolution Protocol (RARP) to obtain an IP address from a RARP server. MAC address information for the port is broadcast on all ports associated with an IP protocol-based or port-based VLAN. To enable a device to request an IP address from a RARP server outside its IP VLAN, you must create a RARP protocol-based VLAN.

RARP has the format of an ARP frame but its own Ethernet type (8035). You can remove RARP from the IP protocol-based VLAN definition and treat it as a separate protocol, thus creating a RARP protocol-based VLAN.

A typical network topology provides desktop switches in wiring closets with one or more trunk ports that extend to one or more data center switches where attached servers provide file, print, and other services. Use RARP functionality to define all ports in a network that require access to a RARP server as potential members of a RARP protocol-based VLAN. You must define all tagged ports and data center RARP servers as static or permanent members of the RARP VLAN. Therefore, a desktop host broadcasts an RARP request to all other members of the RARP VLAN. In normal operation, these members include only the requesting port, tagged ports, and data center RARP server ports. Because all other ports are potential members of this VLAN and RARP is only transmitted at startup, all other port VLAN memberships expire. With this feature, one or more centrally located RARP servers extend RARP services across traditional VLAN boundaries to reach desktops globally.

ARP configuration using the CLI

Network stations that use IP protocol require both a physical address and an IP address to transmit packets. In situations where the station knows only the network host IP address, the Address Resolution Protocol (ARP) lets you use the network station to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address.

A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

ARP response is enabled by default.

Enable ARP on a Port or a VLAN

Enable ARP on the device so that it answers local ARP requests.

About This Task

You can enable or disable ARP responses on the device. You can also enable ARP proxy, which lets a router answer a local ARP request for a remote destination.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable ARP on the device:

```
ip arp-response
```

Example

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface vlan 200  
Switch:1(config-if)#ip arp-response
```

Enable ARP Proxy

About This Task

Configure an ARP proxy to allow the platform to answer a local ARP request for a remote destination. ARP proxy is disabled by default.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable ARP proxy on the device:

```
ip arp-proxy enable
```

View ARP Information

The **show ip arp** command displays all of the configured and dynamically learned ARP entries in the ARP table.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display ARP information for a specified port or for all ports:

```
show ip arp interface gigabitethernet [slot/port[/sub-port]] [-slot/
port[/sub-port]][,...]
```
3. Display ARP information for a VLAN:

```
show ip arp interface vlan <1-4059>
```

Example

```
Switch:1>show ip arp interface
```

```
=====
                                Port Arp
=====
PORT_NUM  DOPROXY    DORESP
-----
1/1        false     true
1/2        false     true
1/3        false     true
1/4        false     true
1/5        false     true
1/6        false     true
1/7        false     true
1/8        false     true
1/9        false     true
1/10       false     true
1/11       false     true
1/12       false     true
1/13       false     true
1/14       false     true
1/15       false     true
1/16       false     true
1/17       false     true
```

```
--More-- (q = quit)
```

Variable definitions

Use the data in the following table to use the **show ip arp** command.

Variable	Value
<i>A.B.C.D</i>	Specifies the IP address of a network.
<i>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>interface</i>	Displays ARP interface configuration information.
<i>spbm-tunnel-as-mac</i>	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
<i>-s</i>	<p>Specifies a subnet.</p> <p>You must indicate the IP address followed by the subnet mask expressed as <A.B.C.D> <A.B.C.D>.</p>
<i>vlan <1-4059></i>	<p>Displays ARP entries for a particular VLAN ID.</p> <p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p>

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF name expressed as text from 1 to 16 characters in length. The total number of ARPs listed in the summary line of the show ip arp output represents the total number of ARPs on the chassis, including all VRFs.
<code>vrfids WORD<0-512></code>	Specifies a range of VRFIDs as text from 0 to 512 characters in length. The total number of ARPs listed in the summary line of the show ip arp output represents the total number of ARPs on the chassis, including all VRFs.

Use the data in the following table to help you understand the **show ip arp interface** command output.

Variable	Value
PORT_NUM	Indicates the port number.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.

Use the data in the following table to help you understand the **show ip arp interface vlan** command output.

Variable	Value
VLAN_ID	Indicates the VLAN ID.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.

Configure IP ARP Static Entries

About This Task

Configure ARP static entries to modify the ARP parameters on the device. The only way to change a static ARP is to delete the static ARP entry and create a new entry with new information.



Note

Static multicast ARP entries are not supported for NLB Unicast or NLB Multicast operations.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure ARP static entries on the device:

```
ip arp <A.B.C.D> 0x00:0x00:0x00:0x00:0x00:0x00 {slot/port[-slot/port]
[,...][slot/all][all]}
```

Example

Configure ARP static entries:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip arp 192.0.2.10 00-16-76-7D-80-C2 2/1
```

Variable Definitions

The following table defines parameters for the **ip arp** command.

Variable	Value
<code>request-threshold</code> <code><50-1000></code>	Configures the maximum number of outstanding ARP requests that a device can generate. The default value is 500.
<code>timeout</code> <code><1-32767></code>	Configures the length of time in seconds an entry remains in the ARP table before timeout. Note: The aging of ARP records is tied to the aging of MAC records. The ARP record for a given IP address is not removed unless the associated MAC record ages out and the router stops receiving a response to ARP requests for that IP address. In cases where the ARP aging time is set to less than the MAC aging time, the switch waits until the MAC ages out before deleting the ARP for an inactive host.
<code><A.B.C.D></code>	Adds ARP entries.

Clear ARP Entries

Use this procedure to clear dynamic ARP table entries associated with the interface or VLAN.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear ARP entries:

```
clear ip arp interface <gigabitethernet|vlan>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]} <1-4059>>
```

Example

Clear ARP entries:

```
Switch:1> enable
Switch:1# clear ip arp interface gigabitethernet 1/16
```

Variable Definitions

The following table defines parameters for the **clear ip arp interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>gigabitethernet vlan</i>	Specifies the interface type.
{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Show ARP Table Information

Show ARP information to view the configuration information in the ARP table.

About This Task

When you use the *interface* parameter with the **show ip arp** command, you can display ARP configuration information only for a specific switch.

The **show ip arp** command displays all configured and dynamically learned ARP entries in the ARP table.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the ARP table:

```
show ip arp [<A.B.C.D>] [-s <A.B.C.D>] [gigabitEthernet <slot/port[/sub-port>]] [interface <gigabitEthernet|vlan>] [nlb] [spbm-tunnel-as-mac] [vlan <1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example



Note

The **show ip arp** command output includes a column heading, MULTICAST-MAC-FLOODING, that is not supported on the switch. You cannot enable this feature.

```
Switch:1#show ip arp
```

```
=====
                        IP Arp - GlobalRouter
=====
IP_ADDRESS    MAC_ADDRESS    VLAN    PORT    TYPE    TTL(10 Sec)    TUNNEL    L3 ACCESS
-----
192.0.2.1     00:09:0f:09:00:08    20     1/3     DYNAMIC 2159                      GlobalRouter
192.0.2.12    b4:a9:5a:ff:f8:40    20     1/3     DYNAMIC 458                      GlobalRouter
192.0.2.25    e4:5d:52:3c:65:00    20     -       LOCAL   2160                      GlobalRouter
192.0.2.154   d4:ea:0e:c2:08:00    20     1/3     DYNAMIC 2131                      GlobalRouter
192.0.2.157   00:1c:17:b1:ec:80    20     1/3     DYNAMIC 2131                      GlobalRouter
192.0.2.161   fc:a8:41:fb:40:00    20     1/3     DYNAMIC 2131                      GlobalRouter
192.0.2.253   e0:db:55:d4:e5:7c    20     1/3     DYNAMIC 2041                      GlobalRouter
192.0.2.255   ff:ff:ff:ff:ff:ff    20     -       LOCAL   2160                      GlobalRouter
=====
```

```
=====
                        IP Arp Extn - GlobalRouter
=====
MULTICAST-MAC-FLOODING    AGING(Minutes)    ARP-THRESHOLD
-----
disable                    360                500
```

```
c: customer vid    u: untagged-traffic
```

```
8 out of 8 ARP entries displayed
```

```
ARPs on TX-NNI: Current = 0, re-ARP count = 0
```

Variable definitions

Use the data in the following table to help you use the **show ip arp** command.

Variable	Value
<code>-s</code>	Specifies the subnet for the table.
<code>gigabitEthernet</code>	Displays the entries for a particular brouter port.
<code>interface</code>	<p>Displays ARP interface configuration information.</p> <p>Use the following parameters to display ARP table information specifically for:</p> <ul style="list-style-type: none"> <code>gigabitethernet {slot/port[/sub-port]}[-slot/port[/sub-port]] [,...] [slot/all] [all]</code> displays IP ARP gigabitethernet interface information <code>VLAN <1-4059></code> displays IP ARP VLAN interface information <p>Example: <code>show ip arp interface vlan 1</code></p>
<code>nlb</code>	Displays the Network Load Balancing (NLB) ARP entries on the switch.
<code>spbm-tunnel-as-mac</code>	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
<code>vlan</code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <code>vrf-scaling</code> and <code>spbm-config-mode</code> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>Use these parameters to display ARP table information specifically for:</p> <ul style="list-style-type: none"> <code>vrf WORD<1-16></code>—the VLAN VRF name in a range from 1 to 16 characters <code>vrfids WORD<0-512></code>—the VLAN VRF ID in a range from 0 to 512 <p>Example: <code>show ip arp vlan 1 vrf 1</code></p>
<code>vrf WORD <1-16></code>	<p>Specifies the name of the VRF.</p> <p>The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis including all VRFs.</p>

Variable	Value
<code>vrfids WORD <0-512></code>	Specifies the VRF ID. The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis including all VRFs.
<code><A.B.C.D></code>	Specifies the network IP address for the table.

Configure Gratuitous ARP

Use the following procedure to configure Gratuitous Address Resolution Protocol (ARP). When Gratuitous ARP is enabled the switch allows all Gratuitous ARP request packets. The default is enabled.

If you disable Gratuitous ARP, the switch only allows Gratuitous ARP packets associated with Routed Split Multi-Link Trunking (RSMLT) or Virtual Router Redundancy Protocol (VRRP), and the switch discards all other Gratuitous ARP request packets.

About This Task

ARP translates network layer (layer 3) IP addresses into link layer (layer 2) MAC addresses. A host sends a Gratuitous ARP request packet to inform other hosts of the existence of an interface on the network, so other local hosts can update their ARP tables. If the IP or MAC address changes, or in the event of a failover, a host sends a Gratuitous ARP request packet to inform other hosts to update their ARP tables.

VRRP and RSMLT use gratuitous ARP to update the MAC address tables on switches.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable Gratuitous ARP:
`ip gratuitous-arp`
3. Save the changed configuration.
`save config [backup WORD<1-99>] [file WORD<1-99>] [verbose]`

ARP configuration using Enterprise Device Manager

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host IP address, the network station can use Address Resolution Protocol (ARP) to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

Enable or Disable ARP on a Port

After you assign the IP address, you can configure ARP. By default, ARP Response is enabled and Proxy ARP is disabled.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand: **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **ARP** tab.
5. In **DoProxy**, select **enable** to enable the Proxy ARP function.
6. In **DoResp**, select **enable** to configure the system to respond to an ARP. The default is enable.
7. Select **Apply**.

The ARP function is available only when the port or VLAN is routed; that is, it is assigned an IP address.

ARP Field Descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
DoProxy	Configures the system to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.

Enable or Disable ARP on a VLAN

Use the following procedure to enable ARP on VLAN level.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **ARP** tab.

7. In **DoProxy**, select **enable** to enable the Proxy ARP function.
8. In **DoResp**, select **enable** to configure the system to respond to an ARP. The default is enable.
9. Select **Apply**.

The ARP dialog box is available only if the port or VLAN is routed; that is, it is assigned an IP address.

ARP Field Descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
DoProxy	Configures the system to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.

View ARP Entries

You can view and manage known MAC address to IP address associations. In addition, you can create or delete individual ARP entries. For information about how to create a static ARP entry, see [Create Static ARP Entries](#) on page 400.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **ARP** tab.

ARP Field Descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
NetAddress	Specifies the IP address corresponding to the media-dependent physical address.
IfIndex	Identifies the router interface for this ARP entry: <ul style="list-style-type: none">• Brouter interfaces are identified by the slot/port number of the brouter port.• VLAN interfaces are identified by the vlan name.

Name	Description
PhysAddress	Specifies the media-dependent physical address (that is, the Ethernet address).
Type	Specifies the type of ARP entry: <ul style="list-style-type: none"> • local—a locally configured ARP entry • static—a statically configured ARP entry • dynamic—a learned ARP entry
TimeToLive	Indicates the time to live where the ARP is configured.
DestIfIndex	Indicates the slot/port on which the ARP entry was learned. For brouter interfaces this is the same value as IfIndex, but for VLAN interfaces, it designates the particular port in the VLAN on which the ARP was learned.
DestVlanId	VLAN ID where the ARP is configured.
BMac	Identifies the backbone MAC address if the entry is learned from an SPBM network.
DestCvid	Identifies the customer VLAN ID for a Switched UNI port.

Create Static ARP Entries

Use the following procedure to create a static ARP entry.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task



Note

Static multicast ARP entries are not supported for NLB Unicast or NLB Multicast operations.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **ARP** tab.
4. Select **Insert**.
5. In **NetAddress**, type the IP address.
6. Select **Port** or **Port in VLAN**.
7. In the dialog box, select the interface.
8. Select **OK**.
9. In **PhysAddress**, type the MAC address.

10. Select **Insert**.

Configure ARP Proxy

With an ARP proxy, the switch can respond to an ARP request from a locally attached host or end station for a remote destination. Proxy ARP does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet on which the ARP request was received. The reply is generated only if the system has an active route to the destination network.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **ARP** tab.
7. For **DoProxy**, select **enable**.
8. Select **Apply**.



Alternative Routes

[Route Preference](#) on page 403

[Preferences for Static Routes](#) on page 403

[Preferences for Dynamic Routes](#) on page 404

[Alternative Route Configuration using CLI](#) on page 404

[Alternative Route Configuration using EDM](#) on page 406

[IPv6 Alternative Routes Configuration Example](#) on page 407

Table 44: Alternative routes product support

Feature	Product	Release introduced
Alternative routes for IPv4	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Alternative routes for IPv6	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

To avoid traffic interruption, you can globally enable the alternative routes feature so the router can use the next-best route, also known as an alternative route, if the best route becomes unavailable.

Routers learn routes to a destination through routing protocols. Routers maintain a routing table of the learned alternative routes sorted in order by route preference, route costs, and route sources. The first route on the list is the best route and the route that the router prefers to use.

The alternative route concept also applies between routing protocols. For example, if an OSPFv3 route becomes unavailable and an alternative RIPng route is available, the system activates the RIPng route without waiting for the update interval to expire.

Route Preference

On the switch, all standard routing protocols have default preference values that determine the routing priority of the protocol. The router uses default preferences to select the best route when a clash exists in preference between the protocols.

You can modify the global preference for a protocol to give the protocol a higher or lower priority than other protocols. If you change the global preference for a static route and all best routes remain best routes, only the local route tables change. However, if the protocol preference change causes best routes to no longer be best routes, the change affects neighboring route tables.



Important

Changing route preferences is a process-intensive operation that can affect system performance and network reach while you perform route preference procedures. As a best practice, if you want to change preferences for static routes or routing protocols, do so when you configure routes or during a maintenance window.

If a router learns a route with the same network mask and cost values from multiple sources, the router uses the route preferences to select the best route to add to the forwarding database.



Note

To modify the preference for a route, you do *not* need to disable a route before you edit the configuration.

Preferences for Static Routes

When you configure a static route on the switch, you can specify a global preference for the route. You can also specify an individual route preference that overrides the global static route preference. The preference value can be between 0 and 255, with 0 reserved for local routes and 255 representing an unreachable route.

Preferences for Dynamic Routes

You can modify the preference value for dynamic routes through route filtering and IP policies, and this value overrides the global preference for the protocol.

The following table shows the default preferences for routing protocols and route types. Use this table to help you modify the global preference value.

Table 45: Routing protocol default preferences

Protocol	Default preference
Local	0
Static	5
SPBM_L1	7
OSPF intra-area	20
OSPF inter-area	25
Exterior BGP	45
RIP/RIPng	100
OSPF external type 1	120
OSPF external type 2	125
IBGP	175
Staticv6	5
OSPFv3 intra-area	20
OSPFv3 inter-area	25
OSPFv3 external type 1	120
OSPFv3 external type 2	125

Alternative Route Configuration using CLI

This section provides procedures to configure alternative routes using the command line interface (CLI).

Enable IPv4 Alternative Routes

About This Task

The default value is enabled. If you disable the alternative-route parameter, all existing alternative routes are removed. After you enable the parameter, all alternative routes are readed.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Activate the alternative route feature globally:

```
ip alternative-route
```

Enable IPv6 Alternative Routes

About This Task

Use this procedure to enable IPv6 alternative routes and view the configuration on the switch.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable IPv6 alternative routes:

```
ipv6 alternative-route
```



Note

IPv6 alternative routes are enabled by default.

3. Verify the configuration of the IPv6 alternative route:

```
show ipv6 global [vrf WORD<1-16> | vrfids WORD<0-512>]
```

```
show ipv6 route alternative [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Example:

```
Switch:1>enable
Switch:1#configure terminal
```

```
Switch:1(config)#router vrf globalRouter
Switch:1(router-vrf)#ipv6 alternative-route
```

```
Switch:1#show ipv6 global
```

```
=====
                        IPv6 Global Information - GlobalRouter
=====
forwarding                : enable
default-hop-cnt            : 64
number-of-interfaces       : 0
icmp-error-interval        : 1000
icmp-error-quota           : 50
icmp-unreach-msg           : disable
icmp-addr-unreach-msg      : enable
icmp-port-unreach-msg      : enable
icmp-echo-multicast-request : enable
icmp-drop-fragments        : disable
static-route-admin-status  : enable
alternative-route          : enable
ecmp                       : disable
ecmp-max-path              : 1
source-route               : disable
host-autoconfig            : disable
```

```
Switch:1#show ipv6 route alternative
```

```
=====
                        IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                NH VRF/ISID  VID/BID/TID  PROTO  COST  AGE  TYPE
-----
2910:0:0:1:0:0:0:0/64         39-4900-24XE            V-3          OSPF      2      0      B
20
2912:0:0:1:0:0:0:0/64         41-VSP8200              V-1001       LOCAL     1      0      B      0
2912:0:0:1:0:0:0:0/64         42-VSP-8404              T-10         BGP       1      0      A
45
3000:0:0:1:0:0:0:0/64         VSP-8404                 V-3          LOCAL     1      0      B      0
4001:0:0:1:0:0:0:0/64         45-VSP7400-32            T-10         LOCAL     1      0      B      0
5910:0:0:1:0:0:0:0/64         46-VSP8404C              T-10         BGP       1      0      B
45
5910:0:0:1:0:0:0:0/64         49-VSP4K                  V-3          OSPF      2      0      A
120
5910:0:0:2:0:0:0:0/64         219-5420-48P             T-10         BGP       1      0      B
45
5910:0:0:2:0:0:0:0/64         223-5420F-24P            V-3          OSPF      2      0      A
120
-----
13 out of 13 Total Num of Route Entries displayed.
-----
TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route
```

Alternative Route Configuration using EDM

This section provides procedures to configure alternative routes using Enterprise Device Manager (EDM).

Enable IPv4 Alternative Routes

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about

how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Enable alternative routes so that you can subsequently enable it on interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.
4. Select **AlternativeEnable**.

If the **AlternativeEnable** parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are re-added.

5. Select **Apply**.

Enable IPv6 Alternative Routes

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

To avoid traffic interruption, enable alternative routes on the switch, to replace the best route with the next-best route if the best route becomes unavailable. By default, this feature is enabled.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Globals** tab.
4. To enable IPv6 alternative routes, select **AlternativeRouteEnable**.
5. Select **Apply**.

IPv6 Alternative Routes Configuration Example

To avoid traffic interruption, you can enable alternative routes globally to replace the best route with the next-best route, if the best route becomes unavailable.

The concept of alternative route applies between routing protocols. For example, if an OSPFv3 route becomes unavailable and an alternative RIPng route is available, the system activates the RIPng route immediately without waiting for an update interval to expire.

By default, the alternative routes feature is globally enabled on the switch.

The following example demonstrates this behavior.

In this example, you configure OSPFv3 and RIPng routes on two switches, Switch-1 and Switch-2, as shown in the following table.

Table 46: Switch configuration

Configuration	Switch-1	Switch-2
VLAN ID	2	2 and 3
IPv6 address	2000:0:0:0:0:0:1/64	VLAN 2: 2000:0:0:0:0:0:2/64 VLAN 3: 3000:0:0:0:0:0:2/64
Protocols	OSPFv3 and RIPng	OSPFv3 and RIPng
VLAN port members	1/5	VLAN 2: 1/5 VLAN 3: 1/6

Switch-1 Configuration

On Switch-1, configure VLAN 2 and the IPv6 interface address 2000:0:0:0:0:0:1/64.

```
Switch1:1:1>enable
Switch1:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch1:1(config)#vlan create 2 type port-mstprstp 0
Switch1:1(config)#vlan members 2 1/5
Switch1:1(config)#interface vlan 2
Switch1:1(config-if)#ipv6 interface address 2000:0:0:0:0:0:1/64
Switch1:1(config-if)#ipv6 interface enable
Switch1:1(config-if)#exit
Switch1:1(config)#show vlan basic

=====
Vlan Basic
=====
VLAN      NAME      TYPE      MSTP
ID  NAME      TYPE      INST_ID  PROTOCOLID  SUBNETADDR  SUBNETMASK  VRFID  ORIGIN
-----
1      Default  byPort    0         none        N/A         N/A         0      CONFIG
2      VLAN-2   byPort    0         none        N/A         N/A         0      CONFIG

All 2 out of 2 Total Num of Vlans displayed

Switch1:1>show vlan members

=====
Vlan Port
=====
VLAN  PORT      ACTIVE      STATIC      NOT_ALLOW
ID   MEMBER      MEMBER      MEMBER      MEMBER
-----
1     1/1-1/4,1/6-  1/1-1/4,1/6-
     1/24         1/24
2     1/5          1/5
```

All 2 out of 2 Total Num of Port Entries displayed

Switch1:1(config)#show ipv6 interface vlan 2

```
=====
                        Vlan Ipv6 Interface
=====
IFINDX VLAN PHYSICAL      ADMIN  OPER  TYPE  MTU  HOP REACHABLE RETRANSMIT MCAST  IPSEC   RPC   RPCMODE
INDX      ADDRESS          STATE  STATE                LMT TIME      TIME      STATUS
-----
2050    2    b0:ad:aa:4e:59:00 enable  up    ETHER 1500 64   30000    1000    disable disable disable existonly
=====

                        Vlan Ipv6 Address
=====
IPV6 ADDRESS                                VLAN-ID    TYPE    ORIGIN    STATUS    NAME
-----
2000:0:0:0:0:0:1/64                        V-2        UNICAST MANUAL  PREFERRED
fe80:0:0:0:b2ad:aaff:fe4e:5900/64          V-2        UNICAST LINKLAYER PREFERRED

1 out of 1 Total Num of Interface Entries displayed.
2 out of 2 Total Num of Address Entries displayed.
```

Configure port 1/5.

```
Switch1:1(config)#interface gigabitEthernet 1/5
Switch1:1(config-if)#encapsulation dot1q
Switch1:1(config-if)#no shutdown
Switch1:1(config-if)#exit
```

Configure IPv6 globally.

```
Switch1:1(config)#ipv6 forwarding
Switch1:1(config)#show ipv6 forwarding
      Ipv6 forwarding - GlobalRouter : enable
      ecmp                : disable
      ecmp-max-path       : 1
```

Configure OSPFv3 on the VLAN.

```
Switch1:1(config)#interface vlan 2
Switch1:1(config-if)#ipv6 ospf area 0.0.0.0
Switch1:1(config-if)#ipv6 ospf enable
Switch1:1(config-if)#show ipv6 ospf interface vlan 2
      admin-status       : enable
      area                : 0.0.0.0
      dead-interval       : 40
      hello-interval      : 10
      metric              : 1
      poll-interval       : 120
      priority            : 1
      retransmit-interval : 5
      transit-delay       : 1
      type                : BROADCAST
```

Configure OSPFv3.

```
Switch1:1(config-if)#exit
Switch1:1(config)#router ospf ipv6-enable
Switch1:1(config)#show ipv6 ospf
```

```
=====
```

OSPFv3 Global Information - GlobalRouter

```

=====
router-id                : 170.78.88.0
admin-state              : ENABLED
version                  : 3
area-bdr-rtr-state      : FALSE
as-bdr-rtr-state        : FALSE
helper-mode              : ENABLED
as-scope-lsa-count       : 0
lsa-checksum             : 0
originate-new-lsas       : 22
rx-new-lsas              : 11
ext-lsa-count            : 0
=====

```

Switch1:1(config)#show ipv6 ospf neighbor

OSPF Neighbor - GlobalRouter

```

=====
IFINDX(VID/BRT/CLIP)  NBRROUTERID      NBRIPADDR                      STATE      TTL
=====
2050      (2)          170.78.84.0      fe80:0:0:0:b2ad:aaff:fe4e:5500      Full       31
=====

```

1 out of 1 Total Num of Neighbor Entries displayed.

OSPF Virtual Neighbor - GlobalRouter

```

=====
NBRAREAID      NBRROUTERID      VIRTINTFID  NBRIPV6ADDR                      STATE
=====

```

0 out of 0 Total Num of Virtual Neighbor Entries displayed.

OSPF NBMA Neighbor - GlobalRouter

```

=====
INTERFACE  NBRROUTERID      NBRIPADDR                      STATE
=====

```

0 out of 0 Total Num of NBMA Neighbor Entries displayed.

H = Helping a Restarting neighbor

Switch1:1(config-if)#exit

Configure RIPng on the VLAN.

```

Switch1:1(config)#interface vlan 2
Switch1:1(config-if)#ipv6 rip
Switch1:1(config-if)#ipv6 rip enable
Switch1:1(config-if)#show ipv6 rip interface

```

Total RIPng interfaces: 1

RIPng Interface - GlobalRouter


```
=====
IFINDEX      COST      POISON      SEND      ADMIN      OPER
              STATUS      DEFAULT      STATUS      STATUS
-----
2050 (2      ) 1        disable     disable     enable      enable

1 out of 1 Total Num of RIPng interfaces displayed
```

Configure RIPng globally.

```
Switch1:1(config)#router rip ipv6-enable
Switch1:1(config)#router rip
Switch1:1(config)#show ipv6 rip

=====
RIPng Global - GlobalRouter
=====

Rip : Enabled
HoldDown Time : 120
Timeout Interval : 180
Update Time : 30
Default Info Metric : 1
Default Info State : Disabled
Default Import Metric : 1
```

Switch-2 Configuration

On Switch-2, configure VLAN 2 and VLAN 3 with the IPv6 interfaces 2000:0:0:0:0:0:0:2/64 and 3000:0:0:0:0:0:0:2/64 respectively.

```
Switch2:1>enable
Switch2:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch2:1(config)#vlan create 2 type port-mstprstp 0
Switch2:1(config)#vlan members 2 1/5 portmember
Switch2:1(config)#interface vlan 2

Switch2:1(config-if)#ipv6 interface address 2000:0:0:0:0:0:0:2/64
Switch2:1(config-if)#ipv6 interface enable
Switch2:1(config-if)#ipv6 forwarding
Switch2:1(config-if)#exit

Switch2:1(config)#vlan create 3 type port-mstprstp 0
Switch2:1(config)#vlan members 3 1/6 portmember
Switch2:1(config)#interface vlan 3

Switch2:1(config-if)#ipv6 interface address 3000:0:0:0:0:0:0:2/64
Switch2:1(config-if)#ipv6 interface enable
Switch2:1(config-if)#ipv6 forwarding
Switch2:1(config-if)#exit

Switch2:1(config)#show vlan basic

=====
Vlan Basic
=====
VLAN
ID    NAME    TYPE    MSTP
INST_ID PROTOCOLID SUBNETADDR SUBNETMASK VRFID ORIGIN
```

```

-----
1      Default  byPort      0      none      N/A      N/A      0      CONFIG
2      VLAN-2   byPort      0      none      N/A      N/A      0      CONFIG
3      VLAN-3   byPort      0      none      N/A      N/A      0      CONFIG
-----

```

All 3 out of 3 Total Num of Vlans displayed

Switch2:1(config)#show vlan members

```

=====
                                Vlan Port
=====
VLAN  PORT                ACTIVE          STATIC          NOT_ALLOW
ID    MEMBER              MEMBER          MEMBER          MEMBER
-----
1     1/1-1/4,1/7-1/24    1/1-1/16,1/7-1/24
2     1/5                  1/5
3     1/6                  1/6
-----

```

All 3 out of 3 Total Num of Port Entries displayed

Switch2:1(config)#show ipv6 interface vlan

```

=====
                                Vlan Ipv6 Interface
=====
IFINDX  VLAN  PHYSICAL      ADMIN  OPER  TYPE  MTU  HOP  REACHABLE  RETRANSMIT  MCAST  IPSEC  RPC  RPCMODE
INDX    ADDRESS  STATE    STATE      LMT  TIME  TIME  TIME  TIME  STATUS
-----
2050    2     b0:ad:aa:4e:55:00 enable  up    ETHER 1500 64  30000    1000    disable disable disable existonly
2051    3     b0:ad:aa:4e:55:01 enable  up    ETHER 1500 64  30000    1000    disable disable disable existonly
-----

```

```

=====
                                Vlan Ipv6 Address
=====
IPV6 ADDRESS                                VLAN-ID  TYPE  ORIGIN  STATUS  NAME
-----
2000:0:0:0:0:0:2/64                        V-2     UNICAST MANUAL  PREFERRED
fe80:0:0:0:b2ad:aaff:fe4e:5500/64          V-2     UNICAST LINKLAYER PREFERRED
3000:0:0:0:0:0:2/64                        V-3     UNICAST MANUAL  PREFERRED
fe80:0:0:0:b2ad:aaff:fe4e:5501/64          V-3     UNICAST LINKLAYER PREFERRED
-----

```

2 out of 2 Total Num of Interface Entries displayed.

4 out of 4 Total Num of Address Entries displayed.

Configure ports 1/5 and 1/6.

```

Switch2:1(config)#interface GigabitEthernet 1/5
Switch2:1(config)#encapsulation dot1q
Switch2:1(config)#no shutdown

```

```

Switch2:1(config)#interface GigabitEthernet 1/6
Switch2:1(config)#encapsulation dot1q
Switch2:1(config)#no shutdown

```

Configure IPv6 globally.

```

Switch1:1(config)#ipv6 forwarding
Switch1:1(config)#show ipv6 forwarding

```

```

Ipv6 forwarding - GlobalRouter : enable
ecmp                : disable
ecmp-max-path       : 1

```

Configure OSPFv3 on the VLANs.

```

Switch2:1(config)#interface vlan 2
Switch2:1(config-if)#ipv6 ospf area 0.0.0.0
Switch2:1(config-if)#ipv6 ospf enable

```

```

Switch2:1(config)#interface vlan 3
Switch2:1(config-if)#ipv6 ospf area 0.0.0.0
Switch2:1(config-if)#ipv6 ospf enable
Switch2:1(config-if)#show ipv6 ospf

```

=====

OSPFv3 Global Information - GlobalRouter

=====

```

router-id                : 170.78.84.0
admin-state              : ENABLED
version                  : 3
area-bdr-rtr-state      : FALSE
as-bdr-rtr-state        : FALSE
helper-mode              : ENABLED
as-scope-lsa-count       : 0
lsa-checksum             : 0
originate-new-lsas      : 56
rx-new-lsas              : 62
ext-lsa-count            : 0

```

```
Switch2:1(config-if)#show ipv6 ospf interface
```

```
Total ospf areas: 1
```

```
Total ospf interfaces: 2
```

=====

OSPF Interface - GlobalRouter

=====

IFINDEX(VID/BRT/CLIP)	AREAID	ADM	IFSTATE	METRIC	PRI	DR/BDR	IFTYPE
2050 (2)	0.0.0.0	ena	BDR	1	1	170.78.88.0	BROADCAST
2051 (3)	0.0.0.0	ena	DR	1	1	170.78.84.0	BROADCAST
						0.0.0.0	

```
2 out of 2 Total Num of ospf interfaces displayed
```

```
Total ospf virtual interfaces: 0
```

=====

OSPF Virtual Interface - GlobalRouter

=====

AREAID	NBRIPADDR	STATE

```

0 out of 0 Total Num of ospf virtual interfaces displayed
Switch2:1(config-if)#show ipv6 ospf neighbor

=====
                        OSPF Neighbor - GlobalRouter
=====
IFINDEX(VID/BRT/CLIP)  NBRROUTERID      NBRIPADDR              STATE      TTL
-----
2050      (2)          170.78.88.0          fe80:0:0:0:b2ad:aaff:fe4e:5900 Full       30

1 out of 1 Total Num of Neighbor Entries displayed.

=====
                        OSPF Virtual Neighbor - GlobalRouter
=====
NBRAREAID      NBRROUTERID      VIRTINTFID NBRIPV6ADDR      STATE
-----
0 out of 0 Total Num of Virtual Neighbor Entries displayed.

=====
                        OSPF NBMA Neighbor - GlobalRouter
=====
INTERFACE NBRROUTERID      NBRIPADDR              STATE
-----

0 out of 0 Total Num of NBMA Neighbor Entries displayed.

H = Helping a Restarting neighbor

```

Configure OSPFv3 globally.

```

Switch2:1(config-if)#exit
Switch2:1(config)#router ospf ipv6-enable
Switch1:1(config)#show ipv6 ospf

=====
                        OSPFv3 Global Information - GlobalRouter
=====
router-id          : 170.78.88.0
admin-state        : ENABLED
version            : 3
area-bdr-rtr-state : FALSE
as-bdr-rtr-state   : FALSE
helper-mode        : ENABLED
as-scope-lsa-count : 0
lsa-checksum       : 0
originate-new-lsas : 22
rx-new-lsas        : 11
ext-lsa-count      : 0

```

Configure RIPng on the VLANs.

```
Switch2:1(config)#interface vlan 2
Switch2:1(config-if)#ipv6 rip
Switch2:1(config-if)#ipv6 rip enable
Switch2:1(config-if)#exit

Switch2:1(config)#interface vlan 3
Switch2:1(config-if)#ipv6 rip
Switch2:1(config-if)#ipv6 rip enable
Switch2:1(config-if)#exit
Switch2:1(config)#

Switch2:1(config)#show ipv6 rip interface

Total RIPng interfaces: 2

=====
RIPng Interface - GlobalRouter
=====
```

IFINDEX	COST	POISON STATUS	SEND DEFAULT	ADMIN STATUS	OPER STATUS
2050 (2)	1	disable	disable	enable	enable
2051 (3)	1	disable	disable	enable	enable

```
2 out of 2 Total Num of RIPng interfaces displayed
```

Configure RIPng globally.

```
Switch2:1(config)#router rip ipv6-enable
Switch2:1(config)#router rip

Switch2:1(config)#show ipv6 rip

=====
RIPng Global - GlobalRouter
=====
```

```

Rip : Enabled
HoldDown Time : 120
Timeout Interval : 180
Update Time : 30
Default Info Metric : 1
Default Info State : Disabled
Default Import Metric : 1

```

Viewing Route and Alternative Route Configuration on the Switches

On Switch-1 and Switch-2, the route 3000::2/64 is learned using the protocols RIPng and OSPFv3. The OSPFv3 route is learned as the best route because of its route preference value of 20. The RIPng route is added as alternative route because it has the route preference 100, which is greater than the OSPFv3 route preference of 20. On Switch-2, the route 3000::2/64 is a local route.

View route and alternative route configuration on Switch-1.

```
Switch1:1(config)#show ipv6 route alternative

=====
IPv6 Routing Table Information - GlobalRouter
=====
```

```

=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          OSPF   2     0    B    20
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          RIP    2     0    A   100
-----

3 out of 3 Total Num of Route Entries displayed.
-----

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

Switch1:1(config)#show ipv6 route

=====
                          IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          OSPF   2     0    B    20
-----

2 out of 2 Total Num of Route Entries displayed.
-----

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

```

View route and alternative route configuration on Switch-2.

```

Switch2:1(config)#show ipv6 route alternative

=====
                          IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:2/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          LOCAL  2     0    B    20
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          LOCAL  2     0    A   100
-----

3 out of 3 Total Num of Route Entries displayed.
-----

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

Switch2:1#show ipv6 route

=====
                          IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:2/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-3          LOCAL  1     0    B    20
-----

2 out of 2 Total Num of Route Entries displayed.
-----

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

```

Changing the Route Preference on Switch-1

On the switch, default preferences are assigned to all standard routing protocols. You can modify the global preference for a protocol to give it a higher or lower priority than other protocols. When you change the preference for a static route, if all best routes remain best routes, only the local route tables change. However, if changing the protocol preference causes best routes to no longer be best routes, neighboring route tables can be affected.

In the following example scenario, you configure a different routing preference for the RIPng protocol on Switch-1 and observe the learning of best and alternative routes. The existing route preference for RIPng is 100.

```
Switch1:1#show ipv6 route alternative
```

```
=====
```

```
IPv6 Routing Table Information - GlobalRouter
```

```
=====
```

Destination Address/PrefixLen	NEXT HOP	VID/BID/TID	PROTO	COST	AGE	TYPE	PREF
2000:0:0:0:0:0:1/64	39-4900-24XE	V-2	LOCAL	1	0	B	0
3000:0:0:0:0:0:2/64	SwitchXYZ	V-2	OSPF	2	0	B	20
3000:0:0:0:0:0:2/64	SwitchXYZ	V-2	RIP	2	0	A	100

```
-----
```

```
3 out of 3 Total Num of Route Entries displayed.
```

```
-----
```

```
TYPE Legend:
```

```
A=Alternative Route, B=Best Route, E=Ecmp Route
```

```
Switch1:1(config)#show ipv6 route
```

```
=====
```

```
IPv6 Routing Table Information - GlobalRouter
```

```
=====
```

Destination Address/PrefixLen	NEXT HOP	VID/BID/TID	PROTO	COST	AGE	TYPE	PREF
2000:0:0:0:0:0:1/64	39-4900-24XE	V-2	LOCAL	1	0	B	0
3000:0:0:0:0:0:2/64	SwitchXYZ	V-2	OSPF	2	0	B	20

```
-----
```

```
2 out of 2 Total Num of Route Entries displayed.
```

```
-----
```

```
TYPE Legend:
```

```
A=Alternative Route, B=Best Route, E=Ecmp Route
```

Configure a different route preference for the RIPng protocol, for example, 19.

```
Switch1:1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch1:1(config)#ipv6 route preference protocol ripng 19
```

```
Switch1:1(config)#exit
```

Verify the route preference configuration.

```
Switch1:1#show ipv6 route preference
```

```
=====
```

```
IPv6 Route Preference - GlobalRouter
```

```
=====
```

PROTOCOL	DEFAULT	CONFIG
RIP	100	19

LOCAL	0	0
STATIC	5	5
SPBM_L1	7	7
OSPFv3_INTRA	20	20
OSPFv3_INTER	25	25
EBGP	45	45
RIPNG	100	19
OSPFv3_E1	120	120
OSPFv3_E2	125	125
IBGP	175	175

View the updated route preference (for RIPng) on Switch-1. The RIPng route is now learned as the best route because its route preference value (19) is lower than that of OSPFv3 (20), as shown here.

```
Switch1:1(config)#show ipv6 route
```

```
=====
                        IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          RIP    2     0    B    19
=====

2 out of 2 Total Num of Route Entries displayed.

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

Switch1:1#show ipv6 route alternative
```

```
=====
                        IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          RIP    2     0    B    19
3000:0:0:0:0:0:2/64          SwitchXYZ                      V-2          OSPF   2     0    A    20
=====

3 out of 3 Total Num of Route Entries displayed.

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route
```

Disabling Alternative Route Learning on Switch-1

The following example demonstrates disabling alternative route learning on Switch-1.

View the alternative routes on Switch-1.

```
Switch1:1(config)#show ipv6 route alternative
```

```
=====
                        IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen  NEXT HOP                      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE                  V-2          LOCAL  1     0    B     0
=====
```



```

3000:0:0:0:0:0:2/64      SwitchXYZ      V-2      OSPF      2      0      B      20
3000:0:0:0:0:0:2/64      SwitchXYZ      V-2      RIP       2      0      A      100
-----

3 out of 3 Total Num of Route Entries displayed.
-----

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

```

Disable IPv6 alternative routes on Switch-1.

```

Switch1:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config)#no ipv6 alternative-route
Switch1:1(config)#exit

```

Verify that alternative route learning is disabled.

```

Switch1:1#show ipv6 global
      forwarding           : enable
      default-hop-cnt      : 64
      number-of-interfaces : 1
      icmp-error-interval  : 1000
      icmp-error-quota     : 50
      icmp-unreach-msg     : disable
      icmp-echo-multicast-request : enable
      static-route-admin-status : enable
      alternative-route    : disable
      ecmp                 : disable
      ecmp-max-path        : 1
      source-route         : disable

Switch1:1(config)#show ipv6 route

=====
                        IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen NEXT HOP          VID/BID/TID  PROTO COST  AGE TYPE PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE      V-2          LOCAL   1    0   B    0
3000:0:0:0:0:0:2/64          SwitchXYZ          V-2          OSPF    2    0   B    20
-----

2 out of 2 Total Num of Route Entries displayed.
-----

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

```

Verify that the alternative route (RIPng) is not learned.

```

Switch1:1(config)#show ipv6 route alternative

=====
                        IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen NEXT HOP          VID/BID/TID  PROTO  COST  AGE TYPE PREF
-----
2000:0:0:0:0:0:1/64          39-4900-24XE      V-2          LOCAL   1    0   B    0
3000:0:0:0:0:0:2/64          SwitchXYZ          V-2          OSPF    2    0   B    20
-----

2 out of 2 Total Num of Route Entries displayed.
-----

```

TYPE Legend:

A=Alternative Route, B=Best Route, E=Ecmp Route



Analytics

[Application Telemetry](#) on page 421

[IPFIX](#) on page 428

[sFlow](#) on page 433

[Common Elements Between sFlow and Application Telemetry](#) on page 438

[Application Telemetry Configuration Using CLI](#) on page 438

[IPFIX Configuration Using CLI](#) on page 442

[sFlow Configuration Using CLI](#) on page 449

[Application Telemetry Configuration Using EDM](#) on page 459

[IPFIX Configuration Using EDM](#) on page 461

[sFlow Configuration Using EDM](#) on page 465

This section provides details about the following features:

- Application Telemetry
- Internet Protocol Flow Information eXport (IPFIX)
- sFlow

Application Telemetry

Table 47: Application Telemetry product support

Feature	Product	Release introduced
Application Telemetry	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 47: Application Telemetry product support (continued)

Feature	Product	Release introduced
Application Telemetry Host Monitoring	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0.5
Application Telemetry — using Segmented Management Instance VLAN interface as source IP	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8
Application Telemetry — using Segmented Management Instance CLIP interface (GRT and VRF) as source IP	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8

Extreme Networks offers two Analytics solutions that monitor traffic on your network:

- sFlow

- Application Telemetry

**Important**

You can use either sFlow, or sFlow with Application Telemetry or both at the same time as they can coexist on a switch. Note that to enable Application Telemetry, you must enable sFlow first.

In both solutions, the switch collects flow information and sends it to a central server that processes the information and provides statistical data in the form of reports. Then you can use ExtremeCloud IQ Site Engine to analyze the reports to give you a full understanding of the applications on your network and learn who is using those applications. ExtremeCloud IQ Site Engine also provides information such as DoS tracking, security monitoring, and statistics for protocols, ports, and applications.

This section describes how Application Telemetry works and how to configure it. Because there is some commonality between the two features, this section also describes some sFlow features.

For further information about sFlow, see [sFlow](#) on page 433.

For more information about ExtremeCloud IQ Site Engine, see the documentation on the Extreme Networks Documentation portal with special attention to the *Application Analytics User Guide*.

How Application Telemetry Works

Both sFlow and Application Telemetry mirror packets to a server for deep packet inspection, but they collect streams in different ways:

- sFlow samples 1 out of n packets to create flow streams. This methodology achieves scalability and applies to high speed networks, but it provides limited application visibility.
- Application Telemetry does not sample some packets like sFlow; it monitors all traffic and uses policy rules to filter packets for analysis. This pattern matching methodology enables Application Telemetry to monitor all application-level traffic flows at wire speed on all interfaces simultaneously.

The policy rules that Application Telemetry uses are ACL and ACE filters that are pre-configured in a policy configuration file called `sflow.pol`. This policy file is not user configurable. These rules enable the switch to recognize several signatures that represent a combination of the following:

- IP protocol type (TCP/UDP)
- TCP flags
- Layer 4 port numbers
- data patterns (defined as offset/data/mask triplets)

Pattern matching enables Application Telemetry to target very specific, well-defined packets in each flow and not full streams of traffic. Thus, the switch mirrors only a

relatively few packets to the Analytics Engine. It is the Analytics Engine that performs deep packet inspection to create reports of statistical data.



Important

When you enable Application Telemetry, the switch loads the filter rules based on the following logic:

- Application Telemetry uses the `apptelemetry.pol` or the `sflow.pol` file because the filter rules can exist in either file. The `sflow.pol` file is the default file and is included with the image that is loaded on the switch. This file contains the default filter rules. The `apptelemetry.pol` file is the user-defined file, which can be updated by the ExtremeCloud IQ Site Engine. To use this file, configure Application Telemetry using the ExtremeCloud IQ Site Engine. When you run the Application Telemetry LiveUpdate script from ExtremeCloud IQ Site Engine, the updated `apptelemetry.pol` file is placed in `/intflash/`.
- When you enable Application Telemetry, the feature uses the files in the following order:
 - If the user-defined file (`apptelemetry.pol`) exists, then the switch loads the rules from this file.
 - If the `apptelemetry.pol` file does not exist or if there is a problem reading this file, then the switch uses the default `sflow.pol` file.

Application Telemetry Operational Considerations

The following section describes operational considerations for deploying Application Telemetry, including general considerations, followed by a summary of platform-specific considerations.

General Considerations

The following list describes general Application Telemetry operational considerations:

- When you enable Application Telemetry, it is globally enabled on all ports. You cannot disable the feature on a per-port basis.
- Application Telemetry supports IPv4 and IPv6 packets, although host monitoring is available for IPv4 hosts only.
- Application Telemetry filter rules are not user configurable. However, an updated `app-telemetry.pol` file can be installed through the ExtremeCloud IQ Site Engine.
- If a user-created filter rule (ACL) conflicts with an Application Telemetry defined filter, the user-created rule always takes precedence.
- There are two configurable sFlow collectors (Collector 1 and Collector 2). However, Application Telemetry uses Collector 1 only and you must configure it before enabling Application Telemetry.
- Application Telemetry supports Segmented Management Instance interfaces VLAN and CLIP. Segmented Management Instance interface OOB is not supported.

- Application Telemetry supports user created VRFs when you configure a Segmented Management Instance CLIP interface as the source IP.
- You can use the onboarding VLAN to manage the switch and enable Application Telemetry without any additional management or routing interface configuration.

Other Considerations

This section provides a summary of other operational considerations.

Table 48: Supported flow types

Attribute	Supported
Flows that ingress standard VLAN ports	Yes
Flows that ingress UNI ports	Yes
Flows that ingress NNI ports and egress UNI ports (Layer 2 VSN)	Yes
Flows that ingress NNI ports and egress UNI ports (Layer 3 VSN)	Yes
Flows that ingress NNI ports and terminate locally	Yes
Flow that ingress NNI ports and egress NNI ports	No
Flows on DvR Controllers or DvR Leafs	Yes Note: To know which switches support DvR Controllers or DvR Leafs, see <i>Fabric Engine and VOSS Feature Matrix</i> .

Table 49: Coexistence with security filters

Attribute	4220 Series 5320 Series 5420 Series	5520 Series 5720 Series 7520 Series 7720 Series
IPv6 security filters or IPv6 source guard	Not supported (consistency checks in place) Note: 4220 Series does not support these IPv6 features.	Allowed

Application Telemetry Configuration Overview

After the optional step of uploading the `apptelemetry.pol` file to flash memory using ExtremeCloud IQ Site Engine, activate Application Telemetry by configuring the following:

1. Configure the IP address of the egress interface for the GRE tunnel with the **sFlow agent-ip** command.
2. Enable sFlow with the **sflow enable** command.
3. Configure the IP address of the Analytics Engine with the **sFlow collector 1** command.
4. Enable Application Telemetry with the **app-telemetry enable** command.

The following figure shows the Application Telemetry agent on various routers and switches with packets being sent to the Analytics Engine.




Figure 18: Application Telemetry Overview

The following list identifies the components in the preceding figure:

1. Analytics Engine
2. GRE tunnels
3. Application Telemetry agents

Host Monitoring

You can use Application Telemetry to get better visibility for a selected host by performing a timed packet capture for both incoming and outgoing traffic specific to that host. Initiate the packet capture (PCAP) from ExtremeCloud IQ Site Engine and specify a source or destination IP address to match. ExtremeCloud IQ Site Engine pushes an additional rule to the Application Telemetry agent on the switch, which captures packets that match this rule and uses the existing ERSPAN GRE session to mirror these packets to Analytics Engine for analysis.

To use this feature, all configuration occurs in ExtremeCloud IQ Site Engine. The following prerequisites for configuration must be met:

- Application Telemetry is active.
- The Analytics Engine records application flows.
- You can see the flows in ExtremeCloud IQ Site Engine.

In ExtremeCloud IQ Site Engine, select a flow and configure packet capture. You can specify the host, either the originating or destination host for the flow, and a monitoring interval. For more information about how to configure packet capture in ExtremeCloud IQ Site Engine, see the ExtremeCloud IQ Site Engine documentation.

The following list identifies restrictions specific to host monitoring:

- You cannot configure monitoring of the same host twice.
- Host monitoring shares resources with the filter ACL application. The maximum number of hosts that can be monitored depends on the number of ACEs you configure. If no resources are available, the Resource Manager generates an error for both applications.
- You cannot configure monitoring of the sFlow agent IP address or collector IP address.

Although you use ExtremeCloud IQ Site Engine to configure the packet capture, the switch logs a message when this feature is activated or deactivated. Configuration of host monitoring is not saved; the monitoring is time-based.

**Note**

Host monitoring is supported beginning with ExtremeCloud IQ Site Engine version 8.2.4.

IPFIX

Table 50: Internet Protocol Flow Information eXport (IPFIX) product support

Feature	Product	Release introduced
Internet Protocol Flow Information eXport (IPFIX)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.0
IPFIX — using Segmented Management Instance VLAN as source IP	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.8 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.8

Table 50: Internet Protocol Flow Information eXport (IPFIX) product support (continued)

Feature	Product	Release introduced
IPFIX — using Segmented Management Instance CLIP (GRT and VRF) as source IP	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.8 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.8
Ability to configure IPFIX per UNI port or on all NNIs	4220 Series	Not Supported
	5320 Series	Fabric Engine 9.4 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.4
	5720 Series	Fabric Engine 9.4
	7520 Series	Fabric Engine 9.4
	7720 Series	Fabric Engine 9.4
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 9.4

Internet Protocol Flow Information eXport (IPFIX) is an Internet Engineering Task Force (IETF) standard of export for Internet Protocol flow information.

IPFIX monitors flows that pass an observation point. The switch organizes flows into a flow group, which is contained in an observation domain.

An IPFIX flow is a set of packets that pass an observation point in the network during a certain time interval. Packets that belong to a particular flow have a common set of properties. The switch defines each property using values from the following:

- Source IP address
- Destination IP address
- IP protocol
- L4 source port
- L4 destination port

A packet belongs to a flow if it completely satisfies all defined properties of the flow.

The switch logically organizes flows into a flow group, which corresponds to a single observation point. A flow can belong to only 1 flow group. A flow group is a collection of packet flows that meet match criteria. Examples of flow groups are packets ingressing a specific physical port, or packets with a destination IP address belonging to a specific subnet.

A flow group is contained in an observation domain. The switch assigns the flow group to an observation domain. The observation domain has a unique observation domain ID that you can configure. You can configure only 1 observation domain.

The IPFIX solution consists of the following processes:

- Filtering Rules process: The Filtering Rules process gathers information about flows through different ports, or the observation point. Flow information includes the following:
 - The IPv4 source address.
 - The IPv4 destination address.
 - The L4 source port.
 - The L4 destination port.
 - The transport protocol.
 - The total number of incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
 - The total number of octets in incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
 - The absolute timestamp of the first packet of this flow.
 - The absolute timestamp of the last packet of this flow.

The Filtering Rules process runs on the switch.

- Exporting process: The Filtering Rules process sends information to the Exporting process. The Exporting process uses the UDP transport protocol for network communication with the Collecting process.

The Exporting process runs on the switch.

- Collecting process: You can view flows and export flow information periodically to a collector. A collector can store a large number of flow records from several devices in the network. The IPFIX standard specifies the protocol for exporting the flows to a

collector, including the formatting of flow records and the underlying UDP transport protocol.

Use the collected information for network planning, troubleshooting a live network, and monitoring security threats.

The best practice is to use the ExtremeAnalytics™ solution as the collector. The ExtremeAnalytics solution provides an enhanced method of collecting IPFIX flow information.

The external collector for the IPFIX solution must support our IPFIX template, which contains the following element IDs defined by Internet Assigned Numbers Authority (IANA) IPFIX assignments.

Table 51: IPFIX element IDs

Element ID	Name	Description
0	unknown	Reserved
4	protocolIdentifier	The value of the protocol number in the IP packet header.
7	sourceTransportPort	The source port identifier in the transport header.
8	sourceIPv4Address	The IPv4 source address in the IP packet header.
11	destinationTransport Port	The destination port identifier in the transport header.
12	destinationIPv4Address	The IPv4 destination address in the IP packet header.
85	octetTotalCount	The total number of octets in incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
86	packetTotalCount	The total number of incoming packets for this flow at the observation point since the metering process (re-)initialization for this observation point.
145	templateId	The local template unique to the observation domain.
156	flowStartNanoseconds	The absolute timestamp of the last packet of this flow.
157	flowEndNanoseconds	The absolute timestamp of the last packet of this flow.
192	ipTTL	The value of the time-to-live (TTL) field in the IPv4 packet header.
234	ingressVRFID	The VRF name that receives packets for this flow.
243	dot1qVlanId	The VLAN ID in the Tag Control information of an Ethernet frame.

IPFIX is a push protocol. The Filtering Rules and Exporting processes periodically send IPFIX messages to configured receivers without interaction from the Collecting process.

IPFIX Operational Considerations

IPFIX collects IPv4 flow information on the switch and conforms with the following:

- IPFIX supports only 1 collector.
- IPFIX learns only IPv4 flows.
- IPFIX sends and receives only TCP/UDP flows.
- IPFIX uses only UDP to export packets.
- You can configure only the template exporting timer.
- The Out-of-Band (OOB) port does not support IPFIX.
- IPFIX exports TCP/UDP IPv4 flows on IS-IS interfaces that are members of a VLAN. IPFIX does not capture Mac-In-Mac encapsulated flows on IS-IS interfaces.

IPFIX processes IPv4 UDP or TCP Mac-in-Mac packet flows that are terminated by the switch. IPFIX does not process Mac-in-Mac packet flows that are only traversing the switch (Layer 2 switching).

- The switch supports only ingress sampling. The switch does not support egress sampling.
- IPFIX is supported on Segmented Management Instance interfaces VLAN and CLIP. Segmented Management Instance interface OOB is not supported.
- When you use Segmented Management Instance VLAN interface IP address as the source IP, IPFIX can function in Layer 2 only environments.
- You must use Segmented Management Instance CLIP to access VRFs and Layer 3 VSNs.
- IPFIX supports user created VRFs when you configure a Segmented Management Instance CLIP interface as the source IP.
- You can use the onboarding VLAN to manage the switch and enable IPFIX without any additional management or routing interface configuration.
- When an interface transitions between UNI and NNI roles, its IPFIX operational state is dynamically adjusted according to the IPFIX configuration and interface role:
 - When an interface changes from UNI to NNI, the IPFIX operational state of the interface becomes controlled by the IPFIX NNI configuration.
 - When an interface changes from NNI to UNI, the IPFIX operational state of the interface becomes controlled by the IPFIX UNI port configuration.
- If you disable IPFIX on a specific UNI port or NNI, records learned on that UNI port or NNI are not flushed. Flows age out according to aging timers, even if the flow continues to be received.

sFlow

Table 52: sFlow product support

Feature	Product	Release introduced
sFlow	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
sFlow — using Segmented Management Instance OOB interface as source IP	4220 Series	Fabric Engine 9.2
	5320 Series	Not Supported
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8
sFlow — using Segmented Management Instance VLAN interface as source IP	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8

Table 52: sFlow product support (continued)

Feature	Product	Release introduced
sFlow — using Segmented Management Instance CLIP interface (GRT and VRFs) as source IP	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8

sFlow monitors traffic in a data network. Use sFlow to monitor routers and switches in the network, and capture traffic statistics about those devices. sFlow uses sampling to provide scalability for network-wide monitoring, and therefore applies to high speed networks. The switch sends the sampled data as a User Datagram Protocol (UDP) packet to the specified host and port.

**Note**

sFlow and Application Telemetry send mirrored packets from a common source to a common destination. sFlow sends samples directly to the destination, while Application Telemetry sends mirrored packets through a GRE tunnel, to the same destination. For more information, see [Common Elements Between sFlow and Application Telemetry](#) on page 438.

sFlow consists of the following:

- sFlow agent—Performs two types of sampling:
 - Flow samples: Flow sampling randomly samples an average of 1 out of n packets for each operation.
 - Counter samples: Counter sampling periodically polls and exports counters for a configured interface. This type of sampling uses a counter to determine if the packet is sampled. Each packet that an interface receives, and that a filter does not drop, reduces the counter by one. After the counter reaches zero, the sFlow agent takes a sample.

**Note**

Only generic interface counters and Ethernet interface counters are supported.

- sFlow datagrams—Supports both flow samples and counter samples. Datagrams can be sent from the front panel port or an out-of-band (OOB) port. Each datagram provides information about the sFlow version, the originating IP address of the

device, a sequence number, the number of samples it contains, and one or more flow and/or counter samples.

- sFlow collector—Located on a central server and runs software that analyzes and reports on network traffic. Two sFlow collectors can be configured to be reachable over a management network or Shortest Path Bridging (SPB).

sFlow Operational Considerations

The following section lists operational considerations for deploying sFlow.

- Application-specific integrated circuit (ASIC) or Software Development Kit (SDK) limitation—To avoid wobbling, the counter interval for sFlow is 20 seconds. Minor wobbling can still occur even after configuring the counter interval due to the interaction between the sFlow agent counter export schedule and the frequency with which the switch ASIC SDK copies and caches counters from the ASIC.
- sFlow supports a maximum of two collectors.
- UDP datagram size and the collector buffer are restricted to 1400 bytes. sFlow sends datagrams to the collector when the buffer reaches the 1400-byte capacity or after a timeout of one second is triggered. The collector buffer size cannot be modified.
- The switch supports IPv4 collector IP addresses.
- VLAN counters/statistics are not supported.
- sFlow can be enabled only on the front panel ports.
- You cannot configure the sampling limit. The sampling limit applies system-wide rather than on a per port basis. Sampling rates differ depending on the hardware platform so any sampled packets beyond the limit are dropped. For more information about feature support, see *Fabric Engine Release Notes*.
- The switch supports only ingress sampling. The switch does not support egress sampling.
- The switch does not support enabling sFlow on a link aggregation group (LAG) interface. However, you can enable sFlow on the member interfaces of a LAG.
- When you specify the sFlow source IP as the Segmented Management Instance OOB instance, the Management Instance OOB routing table determines reachability to the collector.
- When you specify the sFlow source IP as the Segmented Management Instance VLAN instance, the Management Instance VLAN routing table determines reachability to the collector.
- When you specify the sFlow source IP as the Segmented Management Instance CLIP instance, the VRF associated with the Management Instance CLIP checks reachability to the collector, including GRT, VRFs or Layer 3 VSNs.
- When you specify the sFlow source IP as a routing VLAN or loopback IP, only the GRT checks reachability to the collector. User-defined VRFs or Layer 3 VSNs are not supported for this option.
- For Segmented Management Instance interfaces, both collectors must be reachable through the Management Instance interface bound to the sFlow source IP. Configurations where one sFlow collector is a Management Instance OOB interface and the other sFlow collector is on an in-band interface are not supported.

- You must use Segmented Management Instance CLIP to access VRFs and Layer 3 VSNs.
- sFlow supports user created VRFs when Segmented Management Instance CLIP interface IP address is used as the sFlow source IP.
- You can use the onboarding VLAN to manage the switch and enable sFlow without any additional management or routing interface configuration.


sFlow Configuration Considerations

The following section describes configuration considerations for sFlow.

If the sFlow collector has two network interface controller (NIC) cards, add a route to the agent IP address for the NIC card on which the sFlow datagrams are received to avoid dropped sFlow datagrams due to reverse path checks.

After you configure the sFlow agent on the network device that you want to monitor, the system collects flow samples or counter samples, and exports these traffic statistics as sFlow datagrams to the sFlow collector on a server or appliance. For example, after the buffers reach capacity or a timeout is triggered, an sFlow datagram, which is a UDP packet, sends the measurement information to the sFlow collector buffers. The UDP payload contains the sFlow datagram.

The following figure shows the sFlow agent on various routers and switches with sFlow datagrams being sent to the sFlow collector.



The following list identifies components in the preceding figure:

1. sFlow collector
2. sFlow datagrams
3. sFlow agents

As a general rule, drop action occurs after sampling completes. However, in situations related to Layer 1 errors such as, MTU exceeded packets, the drop action occurs before sampling begins. For errors such as frame too long, packets are dropped due to the size of the frame being greater than the interface MTU. In this situation, the packets are dropped before sampling begins so only counter polling occurs. To enable trace, use **line-card 1 trace level 232 <0-4>**.



Important

The defined sampling rate, an average of 1 out of n packets/operations does not provide a 100% accurate result, but it does provide a result with quantifiable accuracy.

Common Elements Between sFlow and Application Telemetry

sFlow and Application Telemetry send mirrored packets from a common source to a common destination. sFlow sends samples directly to the destination, while Application Telemetry sends mirrored packets through a GRE tunnel, to the same destination.

The tunnel source is the switch that you want to monitor:

- sFlow sends sampled flows.
- Application Telemetry sends packets that match its policy rules.

Both sFlow and Application Telemetry use an agent to package either the sFlow streams or the Application Telemetry packets. To configure the agent, they both use the **sflow agent-ip** command.



Note

The switch sends only one mirrored copy, even if the packet matches two or more policies. For information on which mirrored copies take precedence, see [Application Telemetry Operational Considerations](#) on page 424.

The tunnel destination for the mirrored traffic is a server where software performs a deep packet inspection of the mirrored traffic.

- sFlow sends flow and counter samples as datagrams to the sFlow Collector.
- Application Telemetry sends packets that match the policy rules over a GRE tunnel to the Analytics Engine.

To configure the tunnel destination, they both use the **sflow collector <1-2>** command.



Important

You can configure two Collectors, but Application Telemetry uses Collector 1 only. You must configure Collector 1 before you enable Application Telemetry.

Application Telemetry Configuration Using CLI

Use Application Telemetry to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure this feature using CLI.

Configure the Agent IP Address

Use this procedure to configure the source of the Application Telemetry packets.

You can configure Application Telemetry on Segmented Management Instance interfaces OOB, VLAN, and CLIP. You can also use routing VLANs or loopback IPs as the source IP address. Configure the agent IP address to match the Management Instance interface, routing VLAN or loopback IP address.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable the agent IPv4 address:

```
sflow agent-ip {A.B.C.D}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#sflow agent-ip 192.0.2.27
```

Variable Definitions

Use the data in the following table to use the **sflow agent-ip** command.

Variable	Definition
{A.B.C.D. }	Specifies the agent-ip address (IPv4). Note: For Segmented Management Instance interfaces, you must configure the IP address to the IP address of the Segmented Management Instance interface on which datagrams egress.

Configure an Analytics Engine and Enabling Application Telemetry Globally

Use this procedure to enable Application Telemetry and configure the device used as either an sFlow Collector or an Application Telemetry Analytics Engine. This device is where the agent sends sFlow datagrams and Application Telemetry packets for analysis.

sFlow supports up to two collectors for each interface slot in the chassis. However, Application Telemetry supports Collector 1 only.



Note

- You can configure two Collectors, but Application Telemetry uses Collector 1 only. You must configure Collector 1 before you enable Application Telemetry.
- Before you change or remove Collector 1, you must disable Application Telemetry.
- By default, Application Telemetry is globally disabled.

Before You Begin

- You must configure the sFlow agent IP address.
- You must enable sFlow before you can enable Application Telemetry.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the Analytics Engine information using Collector 1:

```
sflow collector 1 address {A.B.C.D} [owner WORD<1-20>] [port <1-65535>]
```
3. Verify the Analytics Engine configuration:

```
show sflow collector 1
```
4. Enable Application Telemetry:

```
app-telemetry enable
```
5. Verify the global configuration:

```
show app-telemetry status
```



Note

The output of this command shows whether Application Telemetry is enabled or not and if the collector is reachable.

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#sflow collector 1 address 192.0.2.26 owner sflow1 port 6343 timeout 497
Switch:1(config)#show sflow collector 1
```

```
=====
                        sFlow Collector Configuration Info
=====
```

Id	Owner	Collector-IP	Port	Timeout(secs)	Reachable via
1	sflow1	192.0.2.26	6343	497	192.0.2.15

```
-----
All 1 out of 1 Total Num of sflow collector entries displayed
Switch:1(config)#app-telemetry enable
Switch:1(config)#show app-telemetry status
Application Telemetry is enabled
Collector is reachable via 192.0.2.26
```

Variable Definitions

Use the data in the following table to use the **sflow collector** command.

Variable	Value
<1-2>	Specifies the ID of the collector where you want to send packets for analysis. Application Telemetry uses Collector 1 only.
owner WORD<1-20>	Specifies the name of the collector.

Variable	Value
<i>Collector-IP</i> {A.B.C.D.}	Specifies the IP address of the collector.
<i>port</i> <1-65535>	Specifies the destination port. The default port is 6343. Note: Application Telemetry does not use this parameter.
<i>timeout</i> <1-65535>	Specifies the time remaining (in seconds) before the collector is released. The default timeout is 0, which means the timeout is not used and the switch sends data forever. Note: Application Telemetry does not use this parameter.

View Application Telemetry Counters

Use the following procedure to view the Application Telemetry status counters. The switch assigns an ID to each counter and displays information about each filter rule by name. The information includes how many packets were transmitted to the Analytics Engine that matched the specified pattern in the rule and the total number of bytes in the packets.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View Application Telemetry counters:

```
show app-telemetry counter [name <WORD<1-32> | id <1-2000>]
```

Example

```
Switch:1>show app-telemetry counter
```

```
=====
Application Telemetry Counters
=====
EntryId      Name      Packets    Bytes
-----
1            ssh       1258       72145
2            sslclient 457        27000
=====
```

```
All 2 out of 2 Total Num of Application Telemetry counters entries displayed
```

Clear Application Telemetry Counters

Use this procedure to clear the Application Telemetry status counters. You can clear all of the counters or specify just the counters you want to clear by name or ID.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Clear Application Telemetry counters:

```
clear app-telemetry counter [name <WORD<1-32> | id <1-2000>]
```
3. Verify that the counters were cleared:

```
show app-telemetry counter [name <WORD<1-32> | id <1-2000>]
```

Example

Clear the counters.

```
Switch:1>enable
Switch:1#clear app-telemetry counter
Switch:1>show app-telemetry counter
```

Application Telemetry Counters			
EntryId	Name	Packets	Bytes
1	ssh	0	0
2	sslclient	0	0

All 2 out of 2 Total Num of Application Telemetry counters entries displayed

IPFIX Configuration Using CLI

This section provides procedures to configure IPFIX using Command Line Interface (CLI).

Enable IPFIX Globally

About This Task

Use the following procedure to enable IPFIX globally. IPFIX provides the ability to monitor IPv4 traffic flows.

If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.

The default global state is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable IPFIX:

```
ip ipfix enable
```


*Examples***Enable IPFIX globally:**

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#ip ipfix enable
```

Disable IPFIX globally:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#no ip ipfix enable

By disabling IPFIX globally, all the processes and traffic
sent to collector(s) will be stopped.Do you agree (y/n) ? y
```

NEW Disable IPFIX on All NNI**About This Task**

Perform this procedure to disable IPFIX on all NNI. By default, IPFIX is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Disable IPFIX on all NNI:

```
no ip ipfix nni enable
```

NEW Disable IPFIX on a UNI Port**About This Task**

Perform this procedure to disable IPFIX on a UNI port. By default, IPFIX is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Disable IPFIX on a UNI port:

```
no ip ipfix port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]
[slot/all][all]} enable
```

Display IPFIX Global Status**About This Task**

Use the following procedure to display global status information for IPFIX.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IPFIX global status:

```
show ip ipfix
```

Example

The output can differ by hardware platform.



Note

For 7830 Series, the Flow Limit displays as 262144.

```
Switch:1#show ip ipfix
```

```
=====
                                IPFIX Global
=====
      Global-State : enable
  Observation-Domain ID : 1
        Flow Limit : 20000
        Flow Count : 0
      Aging Interval : 40
  Enabled UNI ports : 1/1,1/2, 1/10
        NNI state  : disabled
```

Configure the IPFIX Aging Interval

About This Task

Use the following procedure to configure an aging interval for IPFIX. The aging interval determines how long a traffic flow that is no longer being received, is retained as a flow.

Procedure

1. Enter Global Configuration mode:


```
enable
```

```
configure terminal
```
2. Depending on the switch, use the appropriate command to configure a value for the aging interval:

Command

```
ip ipfix aging-interval <1-60>
```

```
ip ipfix aging-interval <100ms | 1.6s | 6.4s |
25.6s | 1m42.4s | 6m49.6s | 27m18.4s | 1h49m13.6s |
7h16m54.5s | 1d5h7m37.6s>
```

```
ip ipfix aging-interval <100ms | 1s | 10s | 60s | 10m
| 30m | 1h |10h |1d>
```

Description

7520 Series and 7720 Series

7830 Series

Other supported Universal hardware switches

Example

```
Switch:1>enable
Switch:1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ipfix aging-interval 1d
```

Variable Definitions

The following table defines parameters for the **ip ipfix aging-interval** command.

Variable	Value
<code><1-60></code> Note: Exception: only applies to 7520 Series and 7720 Series	Specifies, in seconds, the flow record aging interval. The aging interval determines how long a traffic flow that is no longer being received is retained as a flow. The default is 40 seconds.
<code><100ms 1s 10s 60s 10m 30m 1h 10h 1d></code> Note: Exception: does not apply to 7830 Series, 7520 Series, and 7720 Series	Specifies, in seconds, minutes, hours, or days, the flow record aging interval. The aging interval determines how long a traffic flow that is no longer being received is retained as a flow. The default is 10 seconds.
<code><100ms 1.6s 6.4s 25.6s 1m42.4s 6m49.6s 27m18.4s 1h49m13.6s 7h16m54.5s 1d5h7m37.6s></code> Note: Exception: only applies to 7830 Series.	Specifies, in seconds, minutes, hours, or days, the flow record aging interval. The aging interval determines how long a traffic flow that is no longer being received is retained as a flow. The default is 1 minute 42.4 seconds.

Configure the IPFIX Collector

Before You Begin

IPFIX source port behavior on the 7830 Series works as follows. When you configure the IPFIX source port on a 7830 Series switch, the system automatically reserves four consecutive UDP ports, starting from the base port you specify. For example, if you configure port 2055, the switch internally reserves the three subsequent ports—2056, 2057, and 2058. Only the base port is configurable; the hardware automatically allocates the remaining ports.

You must ensure that the entire range $[n, n+3]$ remains available and is not restricted by external firewalls or overlapping service configurations. You cannot manually configure ports $[n+1]$ through $[n+3]$.

If any port in the block is unavailable, the switch can fail to export IPFIX templates or send incomplete flow data.

About This Task

Use the following procedure to configure a collector for IPFIX. Use the ExtremeAnalytics™ solution as the collector.

You can configure IPFIX on Segmented Management Instance interfaces VLAN and CLIP. You can also use routing VLANs or loopback IPs as the exporter IP address. Configure the exporter IP address to match the Management Instance interface IP address.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure values for the collector ID, the IP address of the collector, and the IP address of the exporter. Optionally, you can configure values for the source port sending flow information and the destination port receiving flow information:

```
ip ipfix collector <1-1> {A.B.C.D} exporter-ip {A.B.C.D} [dest-port <1-65535>] [src-port <1-65535>]
```



Note

You cannot configure collector or exporter IP addresses in the following formats:

- 255.255.255.255
- 127. x.x.x
- 0.x.x.x
- 224.0.0.0 to 239.255.255.255

If you configure a collector or exporter IP address in any of these formats, the following error message is displayed:

```
Error: Invalid IP address
```

3. (Optional) Configure a value for the export interval:

```
ip ipfix collector 1 export-interval <1-120>
```

4. (Optional) Configure a value for the initial burst of template packets:

```
ip ipfix collector 1 initial-burst <1-10>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ipfix collector 1 192.0.2.15 exporter-ip 192.0.2.16
dest-port 2 src-port 4
Switch:1(config)#ip ipfix collector 1 export-interval 40
Switch:1(config)#ip ipfix collector 1 initial-burst 4
```

Variable Definitions

Use the data in the following table to use the **ip ipfix collector** command.

Variable	Value
<1-1>	Specifies the IPFIX collector ID.
{A.B.C.D}	Specifies the IP address of the collector.
exporter-ip {A.B.C.D}	Specifies the IP address of the exporter. Note: For Segmented Management Instance interfaces, the exporter-IP address is the IP address of the Segmented Management Instance interface on which datagrams egress.
dest-port <1-65535>	Specifies the destination port receiving flow information.
src-port <1-65535>	Specifies the source port sending flow information.
export-interval	Specifies, in seconds, the frequency of template packet exports to the collector. The default value is 60 seconds.
initial-burst	Specifies the number of template packets sent when the collector becomes reachable. The default value is 5.

Display IPFIX Collector Information

About This Task

Use the following procedure to display information about the IPFIX collector. The IPFIX collector can store flow information from multiple network devices.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about the IPFIX collector:

```
show ip ipfix collector [<1-1>]
```

Example

IPFIX Collector-Info								
ID	Collector	Exporter	Source	Destination	Collector	Reachable	Exporting	Initial
	IP Address	IP Address	Port	Port	State	Via	Interval	Burst
1	20.20.20.2	20.20.20.1	2055	2055	Enabled	-	60	5

Variable Definitions

Use the data in the following table to use the **show ip ipfix collector** command.

Variable	Value
<1-1>	Specifies the IPFIX collector ID.

Configure an IPFIX Observation Domain

About This Task

An observation domain consists of a collection of flow groups. Use this procedure to assign a unique ID to an observation domain. The default value is 0.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Configure the observation domain ID:

```
ip ipfix observation-domain <0-4294967295>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ipfix observation-domain 1
```

Variable Definitions

Use the data in the following table to use the **ip ipfix observation-domain** command.

Variable	Value
observation-domain <0-4294967295>	Specifies a value for the observation domain used to send IPFIX messages. The default is 0 (no observation domain). If the value is 0, data that is sent is not applied to a single observation domain.

Display IPFIX Flows

About This Task

Use the following procedure to display information about IPFIX flows. You can display information for all IPFIX flows, or you can specify a single IPFIX flow and display additional information about that flow.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display information about all IPFIX flows:

```
show ip ipfix flows
```

3. Display information about an IPFIX flow learned on an ingress VRF:

```
show ip ipfix flows source-addr {A.B.C.D} dest-addr {A.B.C.D} source-  
port <1-65535> dest-port <1-65535> protocol {udp|tcp} in-vrf <0-511>  
{in-vlan <1-4095> | untagged} ip-ttl <0-255>
```

Example

Display all IPFIX flows:



Note

Although you can run the **show ip ipfix flows** command on the 7830 Series, it always returns empty output. To display the number of flows learned, use the **show ip ipfix** command.

```
Switch:1>show ip ipfix flows
```

IPFIX Flows					
Source IP	Destination IP	Source Port	Destination Port	Protocol	In Port
192.0.2.1	198.51.100.1	63	63	UDP	1/17
203.0.113.2	203.0.113.1	63	63	UDP	RX-NNI

Variable Definitions

The following table defines parameters for the **show ip ipfix flows** command.

Variable	Value
<i>dest-addr {A.B.C.D}</i>	Specifies an IP address for the flow destination.
<i>dest-port <1-65535></i>	Specifies a value for the destination port.
<i>in-vlan <1-4095></i>	Specifies an ingress VLAN ID for the flow.
<i>in-vrf <0-511></i>	Specifies an ingress VRF ID for the flow.
<i>ip-ttl <0-255></i>	Specifies the IP time-to-live (TTL) for the flow.
<i>protocol {udp tcp}</i>	Specifies the transport protocol.
<i>source-addr {A.B.C.D}</i>	Specifies an IP address for the flow source.
<i>source-port <1-65535></i>	Specifies a value for the source port.
<i>untagged</i>	Specifies ingress untagged flows.

sFlow Configuration Using CLI

Use sFlow to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure sFlow using CLI.

Configure the Agent IP Address and Enable sFlow Globally

About This Task

Configure the sFlow agent IPv4 address, and then enable sFlow before the system can monitor and capture traffic statistics to send to an sFlow collector. By default, sFlow is globally disabled.

You can configure sFlow on Segmented Management Instance interfaces OOB, VLAN, and CLIP. You can also use routing VLANs or loopback IPs as the source IP address. Configure the agent IP address to match the Management Instance interface, routing VLAN or loopback IP address.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
 2. Configure the sFlow source IP address:
`sflow agent-ip {A.B.C.D}`
 3. Configure the sFlow collector:
`sflow collector <1-2> address {A.B.C.D}`
 4. Configure the sFlow destination port number:
`sflow collector <1-2> port <1-65535>`
 5. Configure the sFlow collector timeout:
`sflow collector <1-2> timeout <1-65535>`
 6. Enter GigabitEthernet Interface Configuration mode to bind the port the collector:
`interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]`
`[, ...] [slot/all] [all]}`
-
- Note**
- If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
7. (Optional) Configure the interface sampling rate:
`sflow sampling-rate <1024-100000>`
 8. (Optional) Configure the interface counter intervals:
`sflow counter-interval <1-3600>`
 9. (Optional) Configure the interface max header size:
`sflow max-header-size <64-256>`
 10. Enable sFlow:
`sflow enable`

Examples

Displays sFlow configuration:

```
Switch:1(config)#show sflow
```


sFlow Global Configuration	
Global State	: Enabled
Agent IP	: 192.0.2.27

Displays sflow collector configuration:

sFlow Collector Information					
Id	Owner	Collector-IP	Port	Timeout(secs)	Reachable via
1	sflow1	192.0.2.26	6343	497	192.0.2.15
2	sflow2	192.0.2.27	6343	531	192.0.2.16

Displays sFlow port configuration:

sFlow Port Configuration Info				
Port	Packet-Sample-Rate	Max-Header-Size	Counter-interval (in secs)	Collector-list
1/1	0	255	525	1,2
1/2	0	255	525	1,2

Displays sFlow statistics:

sFlow Statistics Info	
Collector-id	sFlow-Datagrams
1	1001
2	0

Variable Definitions

The following table defines parameters for the **sflow** commands related to configuring sFlow as the source IP.

Variable	Value
<i>agent-ip</i>	Specifies the agent-ip address (IPv4).
<i>collector <1-2></i>	Specifies the id to export sFlow datagrams to the collector id.
<i>address {A.B.C.D.}</i>	Specifies the collector IP address.
<i>owner WORD<1-20></i>	Specifies the sFlow collector name.
<i>port <1-65535></i>	Specifies the destination UDP port. The default port is 6343.

Variable	Value
<i>timeout</i> <1-65535>	Specifies the time remaining (in seconds) before the collector is released. The default timeout is 0, which means the timeout is not used and the collector sends data forever.
<i>sampling-rate</i> <1024-1000000>	Configures the packet sampling rate on a port. The default value is 0 (disabled).
<i>counter-interval</i> <1-3600>	Specifies the polling interval on a port to determine how often the sFlow agent polls and exports counters for a configured interface. The default value is 0 (disabled).
<i>max-header-size</i> <64-256>	Specifies the maximum number of bytes to be copied from the sampled packet. The default is 128 bytes.

Configure an sFlow Collector

Configure an sFlow collector to determine the device to which the sFlow agent sends sFlow datagrams. You can configure up to two collectors for each interface slot in the chassis.

Before You Begin

- You must globally enable sFlow.

About This Task

The sFlow datagrams that the agent sends to the collector are not encrypted. Use a VLAN to create a secure measurement network to route sFlow datagrams.

To further protect the sFlow collector, configure it to accept only sFlow datagrams, or to check sequence numbers and verify source addresses.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Configure the collector information:

```
sflow collector <1-2> address {A.B.C.D} [Owner WORD <1-20>] [port <1-65535>] [timeout <1-65535>]
```
- Verify the collector configuration:

```
show sflow collector <1-2>
```

Example

Configure collector ID, and then verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
```

```

Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#sflow collector 1 address 192.0.2.26 owner sflow1 port 6343 timeout 497
Switch:1(config)#sflow collector 2 address 192.0.2.27 owner sflow2 port 6343 timeout 531
Switch:1(config)#show sflow collector

```

sFlow Collector Configuration Info

Id	Owner	Collector-IP	Port	Timeout(secs)	Reachable via
1	sflow1	192.0.2.26	6343	497	192.0.2.15
2	sflow2	192.0.2.27	6343	531	192.0.2.16

All 2 out of 2 Total Num of sflow collector entries displayed

What to Do Next

After you configure the sFlow collector, configure the packet sampling rate to enable sFlow on a port or ports.

Variable Definitions

Use the data in the following table to use the **sflow collector** command.

Variable	Value
<i>collector</i> <1-2>	Specifies the id to export sFlow datagrams to the collector id.
<i>address</i> {A.B.C.D.}	Specifies the collector IP address.
<i>owner</i> WORD<1-20>	Specifies the sFlow collector name.
<i>port</i> <1-65535>	Specifies the destination UDP port. The default port is 6343.
<i>timeout</i> <1-65535>	Specifies the time remaining (in seconds) before the collector is released. The default timeout is 0, which means the timeout is not used and the collector sends data forever.

Configure the Packet Sampling Rate

Configure the packet sampling rate at port level to determine how many packets the system counts before it takes a sample. Configuring the sampling rate enables sFlow on the port.

Before You Begin

- You must globally enable sFlow.

About This Task

If you configure a conservative sampling rate to prevent overloading the sFlow agent, the result will reflect high values that do not reflect typical traffic levels.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the collector id:

```
sflow collector <1-2>
```

3. Configure the sampling rate:

```
sflow sampling-rate <1024-1000000>
```

4. Verify the configuration:

```
show sflow interface {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

Example

Configure sampling rates for ports 2/1, 2/2, 2/3, and 2/4.

```
Switch:1(config-if)#interface gigabitethernet 2/1,2/2
Switch:1(config-if)#sflow collector 1
Switch:1(config-if)#sflow sampling-rate 10000
Switch:1(config-if)#interface gigabitethernet 2/3
Switch:1(config-if)#sflow collector 2
Switch:1(config-if)#sflow sampling-rate 8192
Switch:1(config-if)#interface gigabitethernet 2/4
Switch:1(config-if)#sflow collector 2
Switch:1(config-if)#sflow sampling-rate 12001
Switch:1(config-if)#show sflow interface enabled
=====
sFlow Port Configuration Info
=====
Port    Packet-Sample-Rate  Max-Header-Size  Counter-interval  Collector-list
                        (in secs)
-----
2/1     10000               128              0                  1
2/2     10000               128              0                  1
2/3     8192                128              0                  2
2/4     12001               128              0                  2
-----
All 4 out of 4 Total Num of sflow port entries displayed
```

Variable Definitions

Use the data in the following table to use the **sflow sampling-rate** and **show sflow interface** commands.

Variable	Value
<1024-1000000>	Configures the packet sampling rate on a port. The default value is 0 (disabled).
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure sFlow Maximum Header Size

Configure the maximum header size on a single port or multiple ports.

Before You Begin

- You must globally enable sFlow.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all] [all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the maximum-header size:

```
sflow max-header-size <64-256>
```

3. Verify the configuration:

```
show sflow interface {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all] [all]}
```

Example

For ports 1/1 to 1/10, configure the maximum header size, and then verify the configuration.

```
Switch:1(config-if)#interface gigabitethernet 1/1-1/10
Switch:1(config-if)#sflow max-header-size 255
Switch:1(config-if)#show sflow interface 1/1-1/10
```

```
=====
sFlow Port Configuration Info
```

Port	Packet-Sample-Rate	Max-Header-Size	Counter-interval (in secs)	Collector-list
1/1	0	255	525	1,2
1/2	0	255	525	1,2
1/3	0	255	525	1,2
1/4	0	255	525	1,2
1/5	0	255	525	1,2
1/6	0	255	525	1,2
1/7	0	255	525	1,2
1/8	0	255	525	1,2
1/9	0	255	525	1,2
1/10	0	255	525	1,2

Variable Definitions

Use the data in the following table to use the **max-header-size** command.

Variable	Value
<64-256>	Identifies the maximum number of bytes to be copied from the sampled packet. Default 128 bytes.

Configure the Counter Sampling Interval

Configure the counter sampling interval values at port level to determine how often the sFlow agent polls and exports counters for a configured interface.

Before You Begin

- You must globally enable sFlow.

Procedure

- Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- Configure the counter sampling interval:

```
sflow counter-interval <1-3600>
```

- Verify the configuration:

```
show sflow interface {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

Example

Verify all slots use the default polling-interval configuration.

```
Switch:1(config-if)#sflow counter-interval 525
Switch:1(config-if)#show sflow interface 1/1-1/10
```

```
=====
                        sFlow Port Configuration Info
=====
```

Port	Packet-Sample-Rate	Max-Header-Size	Counter-interval (in secs)	Collector-list
1/1	0	128	525	1,2
1/2	0	128	525	1,2
1/3	0	128	525	1,2
1/4	0	128	525	1,2
1/5	0	128	525	1,2
1/6	0	128	525	1,2
1/7	0	128	525	1,2
1/8	0	128	525	1,2
1/9	0	128	525	1,2
1/10	0	128	525	1,2

```
-----
All 10 out of 10 Total Num of sflow port entries displayed
```

Variable Definitions

Use the data in the following table to use the **sflow counter-interval** and **show sflow interface** commands.

Variable	Value
<1-3600>	Specifies the polling interval for a slot. Default value is 0 (disabled).
{slot/port[/sub-port] [-slot/port[/sub-port]][,...] [slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View sFlow Statistics

Display statistics for sFlow datagrams.

Before You Begin

- You must globally enable sFlow.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View sFlow statistics:

```
show sflow statistics [collector <1-2>]
```

Example

```
Switch:1>show sflow statistics

=====
sFlow Statistics Info
=====
Collector-id      sFlow-Datagrams
-----
1                1001
2                0
-----

All 2 out of 2 Total Num of sflow statistics entries displayed
```

Clear sFlow Statistics

Use this procedure to clear the statistics for each collector.

Before You Begin

- You must globally enable sFlow.

Procedure

1. Enter Privileged EXEC mode:


```
enable
```
2. Clear sFlow statistics:


```
clear sflow statistics [collector <1-2>]
```
3. Verify the collector information:


```
show sflow statistics [collector <1-2>]
```

Example

Clear the statistics for collector ID 1.

```
Switch:1>enable
Switch:1#clear sflow statistics collector 1
Switch:1#show sflow statistics collector 1

=====
sFlow Statistics Info
=====
Collector-id      sFlow-Datagrams
-----
1                0
-----

All 1 out of 1 Total Num of sflow statistics entries displayed
```


Application Telemetry Configuration Using EDM

Use Application Telemetry to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure this feature using EDM.

sFlow and Application Telemetry send mirrored packets from a common source to a common destination. sFlow sends samples directly to the destination, while Application Telemetry sends mirrored packets through a GRE tunnel, to the same destination.

Both sFlow and Application Telemetry use an agent to package either the sFlow streams or the Application Telemetry packets. To configure the agent, they both use the **Serviceability > Sflow > Globals** and **Serviceability > Sflow > Collector** tabs. For more information, see [sFlow Configuration Using EDM](#) on page 465.

Enable Application Telemetry Globally

Use this procedure to globally enable Application Telemetry so it can send packets to an Analytics Engine. By default, Application Telemetry is globally disabled.

Before You Begin

You must complete the following:

- Configure an agent IP address.
- Enable sFlow.
- Configure Collector 1.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Application Telemetry**.
3. Select the **Globals** tab.
4. Select **AdminEnable**.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminEnable	Shows whether Application Telemetry is enabled. By default, the check box is not enabled.
ClearCounterStats	Clears the Application Telemetry status counters.

View Application Telemetry Counters

Use the following procedure to view the Application Telemetry status counters. The switch assigns an ID to each counter and displays information about each filter rule by name. The information includes how many packets were transmitted to the Analytics Engine that matched the specified pattern in the rule and the total number of bytes in the packets.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Application Telemetry**.
3. Select the **Counter** tab.

Counter Field Descriptions

Use the data in the following table to use the **Counter** tab.

Name	Description
CounterId	Shows the Application Telemetry rule ID.
CounterName	Shows the rule name.
CounterPkts	Shows the number of packets transmitted to the Analytics Engine that matched the specified pattern in the rule.
CounterBytes	Shows the total number of bytes in the packets.

Clear Application Telemetry Counters

Use this procedure to clear the Application Telemetry status counters. You can clear all of the counters or specify just the counters you want to clear by name or ID.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Application Telemetry**.
3. Perform one of the following actions:
 - To clear all the counters, select the **Globals** tab, and then select **ClearCounterStats**.
 - To clear specific counters, select the **Counter** tab, select the counter ID you want to clear, and then select **ClearStats**.
4. Select **Apply**.

View Application Telemetry Status

About This Task

Use this procedure to view the status of the Application Telemetry collector.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Application Telemetry**.
3. Select the **Status** tab.

Status Field Descriptions

Use the data in the following table to use the **Status** tab.

Name	Description
Collector IP Address	Shows the address of the Application Telemetry collector.
IsReachable	Shows whether the Application Telemetry collector is reachable.
NextHop	If the collector is reachable, shows the name or address of the next hop through which the collector is reachable.

IPFIX Configuration Using EDM

This section provides procedures to configure IPFIX in Enterprise Device Manager (EDM).

Enable IPFIX Globally

About This Task

IPFIX allows you to monitor IPv4 traffic flows. Use the following procedure to enable IPFIX globally.

If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.

The default global state is disabled.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IPFIX**.
3. Select the **Globals** tab.
4. To enable IPFIX, select **enable** in the **ConfState** option box.
5. (Optional) To configure an observation domain ID, type a value in the **ObservationDomainID** field.
6. Depending on the switch, select a value for the aging interval:

Command	Description
AgingInterval	7520 Series and 7720 Series

Command	Description
AgingIntervalV3	7830 Series
AgingIntervalV2	Other supported Universal hardware switches

7. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
ConfState	Specifies whether IPFIX is enabled or disabled. The default is disabled.
ObservationDomainID	Specifies a value for the observation domain used to send IPFIX messages. The default is 0 (no observation domain). If the value is 0, data that is sent is not applied to a single observation domain.
AgingInterval Note: Exception: only applies to 7520 Series and 7720 Series	Specifies (in seconds) the flow record aging interval. The aging interval determines how long a traffic flow that is no longer being received is retained as a flow. The default is 40 seconds.
AgingIntervalV2 Note: Exception: does not apply to 7830 Series, 7520 Series, and 7720 Series	Specifies (in seconds, minutes, hours, or days) the flow record aging interval. The aging interval determines how long a traffic flow that is no longer being received is retained as a flow. The default is 10 seconds.
AgingIntervalV3 Note: Exception: only applies to 7830 Series	Specifies (in seconds, minutes, hours, or days) the flow record aging interval. The aging interval determines how long a traffic flow that is no longer being received is retained as a flow. The default is 1 minute 42.4 seconds.
NniConfState	Enables IPFIX on all NNI. The default is enable. If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.

NEW Disable IPFIX on All NNI

About This Task

Perform this procedure to disable IPFIX on all NNI. By default, IPFIX is enabled.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IPFIX**.
3. Select the **Globals** tab.
4. In **NniConfState**, select **disable**.
5. Select **Apply**.

NEW Disable IPFIX on a UNI Port

About This Task

Perform this procedure to disable IPFIX on a UNI port. By default, IPFIX is enabled.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IPFIX**.
3. Select the **Interface** tab.
4. In **IpfixEnable** for the port, select **false**.
5. Select **Apply**.

NEW Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Index	Displays the port.
IpfixEnable	Enables or disables IPFIX on the port. The default is true, which means enabled. If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.

Configure the IPFIX Collector

Before You Begin

IPFIX source port behavior on the 7830 Series works as follows. When you configure the IPFIX source port on a 7830 Series switch, the system automatically reserves four consecutive UDP ports, starting from the base port you specify. For example, if you configure port 2055, the switch internally reserves the three subsequent ports—2056,

2057, and 2058. Only the base port is configurable; the hardware automatically allocates the remaining ports.

You must ensure that the entire range $[n, n+3]$ remains available and is not restricted by external firewalls or overlapping service configurations. You cannot manually configure ports $[n+1]$ through $[n+3]$.

If any port in the block is unavailable, the switch can fail to export IPFIX templates or send incomplete flow data.

About This Task

Use the following procedure to configure a collector for IPFIX. Use the ExtremeAnalytics™ solution as the collector.

You can configure IPFIX on Segmented Management Instance interfaces VLAN and CLIP. You can also use routing VLANs or loopback IPs as the exporter IP address. Configure the exporter IP address to match the Management Instance interface IP address.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IPFIX**.
3. Select the **Collector** tab.
4. Select **Insert**.
5. In **Num**, specify a value for the collector ID number.
6. In **AddressType**, specify a value for the IP address type of the collector.
7. In **Address**, specify a value for the IP address of the collector.
8. (Optional) In **SrcPort**, specify a value for the source port sending flow information.
9. (Optional) In **DestPort**, specify a value for the destination port receiving flow information.
10. In **ExporterIpType**, specify a value for the IP address type of the exporter for the collector.
11. In **ExporterIp**, specify the IP address of the exporter for the collector.
12. (Optional) In **ExportInterval**, specify the frequency of template packets exports to the collector.
13. (Optional) In **InitialBurst**, specify the number of template packets to send when the collector becomes reachable.
14. Select **Insert**.

Collector Field Descriptions

Use the data in the following table to use the **Collector** tab.

Name	Description
Num	Specifies the ID number of the collector
AddressType	Specifies the IP address type of the collector.

Name	Description
Address	Specifies the IP address of the collector.
Protocol	Specifies the protocol used to export data from the exporter to the collector.
SrcPort	Specifies the source port sending flow information. The default value is 2055.
DestPort	Specifies the destination port receiving flow information. The default value is 2055.
ExporterIpType	Specifies the IP address type of the exporter for the collector.
ExporterIp	Specifies the IP address of the exporter for the collector. Note: For Segmented Management Instance interfaces, the exporter-IP address is the IP address of the Segmented Management Instance interface on which datagrams egress.
State	Specifies the state of the collector. The default value is enabled.
IsReachable	Specifies whether the collector is reachable. The default value is false (not reachable)
ViaNextHopName	Specifies the next-hop through which the collector is reachable.
ExportInterval	Specifies, in seconds, the frequency of template packets exports to the collector. The default value is 60 seconds.
InitialBurst	Specifies the number of template packets sent when the collector becomes reachable. The default value is 5.

sFlow Configuration Using EDM

Use sFlow to capture traffic statistics to monitor traffic in a data network. This section provides procedures to view and configure sFlow using EDM.

Enable sFlow and Configure the Agent IP Address

Use this procedure to enable sFlow and configure the sFlow agent IP address so the system can send packets to an sFlow collector.

Before You Begin

You must enable sFlow before you enable Application Telemetry.

About This Task

Application Telemetry and sFlow both use the **sFlow Globals** tab.

You can configure sFlow on Segmented Management Instance interfaces OOB, VLAN, and CLIP. You can also use routing VLANs or loopback IPs as the source IP address. Configure the agent IP address to match the Management Instance interface, routing VLAN or loopback IP address.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Sflow**.
3. Select the **Globals** tab.
4. Select **AdminEnable** to enable sFlow.
5. In the **AgentAddress** field, enter the agent IPv4 address.
6. Select **Apply**.

What to Do Next

After you configure the agent IP address and globally enable sFlow, proceed to configuring the sFlow collector.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminEnable	Shows whether sFlow is enabled. By default, the check box is not enabled.
AgentAddressType	Specifies the collector IP address type. Only IPv4 collector addresses are supported.
AgentAddress	<p>Specifies the agent IP address of an interface.</p> <p>Note: For Segmented Management Instance interfaces, you must configure AgentAddress to the IP address of the Segmented Management Instance interface on which datagrams egress.</p>

Configure an sFlow Collector

About This Task

Use this procedure to configure the device used as either an sFlow Collector or an Application Telemetry Analytics Engine. This device is where the agent sends sFlow datagrams and Application Telemetry packets for analysis.

sFlow supports up to two collectors for each interface slot in the chassis. However, Application Telemetry supports Collector 1 only.

**Note**

- You can configure two Collectors, but Application Telemetry uses Collector 1 only. You must configure Collector 1 before you enable Application Telemetry.
- Before you change or remove Collector 1, you must disable Application Telemetry.
- By default, Application Telemetry is globally disabled.

**Tip**

You can configure the **Collector** tab to select only the columns you are interested in seeing. By default, the system does not display **AddressType** option. To make the **AddressType** column visible, select the down arrow on one of the menu headings, navigate to **Columns**, and select the **AddressType** check box.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Sflow**.
3. Select the **Collector** tab.
4. Configure the fields for Collector 1.
5. Select **Apply**.

What to Do Next

After you configure the sFlow collector, configure the packet sampling rate to enable sFlow on a port or ports.

Collector Field Descriptions

Use the data in the following table to use the **Collector** tab.

Name	Description
Index	Shows collector 1 and collector 2. The switch exports sFlow and Application Telemetry traffic to the collector.
Owner	Specifies the sFlow collector name. The string length is 1 to 20 characters.
Timeout	Specifies the time remaining (in seconds) before the collector is released and stops sampling. The default timeout is 0, which means the timeout is not used and the switch sends data forever.

Name	Description
Address	Specifies the collector IP address. If the default address is set to 0.0.0.0, no traffic is sent.
Port	Specifies the destination port. The default port is 6343.
IsReachable	Shows whether the sFlow collector is reachable.
NextHop	If the collector is reachable, shows the name or address of the next hop through which the collector is reachable.

Configure the Packet Samples and Counter Samples

Before You Begin

- You must globally enable sFlow.

About This Task

Configure the packet sampling rate to determine how many packets the system counts before it takes a sample and configure the counter sampling interval to determine how often the sFlow agent polls and exports counters for a configured interface. You can also configure the maximum header size on a single port or multiple ports.

Configuring the sampling rate enables sFlow on the port.



Tip

You can configure the **Interfaces** tab to select only the columns you are interested in seeing. By default, the system does not display **Instances** option. To make the **Instances** column visible, select the down arrow on one of the menu headings, navigate to **Columns**, and select the **Instances** check box.

Procedure

- In the navigation pane, expand **Configuration > Serviceability**.
- Select **Sflow**.
- Select the **Interfaces** tab.
- In the **DataSource** column, navigate to the slot and port where you want to configure sFlow, and configure the following:
 - PacketSamplingRate**—Double-click the field, and type a sampling rate value.
 - MaximumHeaderSize**—Double-click the field, and type a maximum header size value.
 - Interval**—Double-click the field, and type the counter sampling interval value in seconds.
- Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
DataSource	Shows the slot and port for which traffic statistics are collected. The field name is different on different hardware platforms.
Instance	Shows the number of sFlow samplers associated with a specific datasource. Note: You must select this field for it to display on the Interfaces tab.
Collectors	Shows the collectors that have been configured for the sFlow agent to send sFlow datagrams. Two collectors are supported. The field name is different on different hardware platforms.
PacketSamplingRate	Specifies the packet sampling rate to determine how many packets the system counts before it take a sample. The default is 0.
MaximumHeaderSize	Specifies the maximum header size on a single port or multiple ports. The default is 128 bytes.
Interval	Specifies the counter sampling interval to determine how often the sFlow agent polls and exports counters for a configured interface. The default is 0.

Display sFlow Statistics

Use the following procedure to display (true) sFlow statistics. Statistics for sFlow are cleared (false), by default.

About This Task



Tip

You can configure the **Stats** tab to select only the columns you are interested in seeing. The system displays all the options, by default. To hide a column, select the down arrow on one of the menu headings, navigate to **Columns**, and select the check box for the column you want to hide.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **sFlow**.
3. Select the **Stats** tab.
4. In **ClearStats**, double-click the field, and select **true** or **false**.
5. Select **Apply**.

Stats Field Descriptions

Use the data in the following table to use the **Stats** tab.

Name	Description
Index	Shows sFlow collector ID 1 and 2
DatagramCount	Shows the number of datagrams that have been sent to the collector.
ClearStats	Shows whether the sFlow statistics are displayed (true) or cleared (false). The default is false.



Anycast IP Gateway

[Anycast IP Gateway Operations](#) on page 472

[I-SID Attributes TLV 187](#) on page 473

[Anycast IP Gateway Benefits](#) on page 473

[Anycast IP Gateway Restrictions](#) on page 473

[Anycast IP Gateway Configuration Using CLI](#) on page 474

[Anycast IP Gateway Configuration Using EDM](#) on page 480

Table 53: Anycast IP Gateway product support

Feature	Product	Release introduced
Anycast IP Gateway	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.1
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1

Anycast IP Gateway provides robust active-active router instances with shortest path forwarding. This feature is beneficial for deployments that share I-SIDs among many sites.

Configure Anycast IP Gateway to use many switches as the default gateway for attached subnets and to distribute routing. All configured routers can route packets they receive with the destination address of the gateway MAC. Unlike with VRRP, you do not need a backup router to forward packets to a primary router.

Anycast IP Gateway distributes routing by announcing the virtual router interface (the gateway IP and gateway MAC) through the IS-IS LSDB, and then BEBs use the SPB shortest path calculation to forward traffic to the closest router. Layer 2 VSNs at the edge of the SPB network process the advertised information, and then bridge the traffic without requiring dynamic MAC learning.

Anycast IP Gateway Operations

Anycast IP Gateway Router

Anycast IP Gateway routers use TLV 187 to advertise the configured routing interfaces. You must associate the routing interfaces with Layer 2 VSNs, and configure each Anycast IP Gateway-enabled interface with a gateway IP and a gateway MAC.

The gateway MAC uses the VRRP MAC address format with the virtual router ID (VRID) in the last byte, 00:00:5e:00:01:xx. The gateway MAC must be the same on different routers for the same associated Layer 2 VSN. Different interfaces can share the gateway MAC. As a best practice, use the same VRID on different interfaces whenever possible.

You can configure a different gateway IP from the VLAN IP, or as a One IP interface, which uses the same IP as the VLAN. The gateway IP must be in different subnets for different interfaces, and must be the same on different routers for the same associated Layer 2 VSNs.

The BEB floods broadcast ARP queries for the Anycast IP Gateway address to all routers. Anycast IP Gateway routers can also issue ARP queries in the subnets of the routing interfaces. For ARP queries initiated by Anycast IP Gateway routers, the sender hardware address (MAC address) and sender protocol address (IP address) differ by scenario. The following table provides details for each scenario.

Table 54: ARP payload by scenario

Scenario	Ethernet Source Address	Sender Hardware Address	Sender Protocol Address
ARP reply	VLAN chassis MAC	GW MAC	GW IP
ARP query non One IP	VLAN chassis MAC	VLAN chassis MAC	VLAN IP
ARP query One IP without CLIP	VLAN chassis MAC	VLAN chassis MAC	0.0.0.0
ARP query One IP with CLIP	VLAN chassis MAC	VLAN chassis MAC	CLIP IP

Anycast IP Gateway BEB

BEBs process the I-SID Attributes TLVs that contain Anycast IP Gateway information advertised by routers for each locally-created Layer 2 VSN, and install a static MAC entry for the gateway MAC with the next hop of the closest router based on SPB cost. BEBs do not dynamically learn the gateway MAC. If multiple gateway routers are reachable with the same SPB cost, the BEB makes the selection based on the IS-IS system ID:

- Even I-SIDs use the router with the higher system ID value.
- Odd I-SIDs use the router with the lower system ID value.

If a vIST node advertises the routing interface TLV, the next-hop BMAC for the static MAC entry is the virtual BMAC of the cluster. If only one of the vIST peers is configured

as a router, then the other peer always selects its peer as the next-hop, regardless of the SPB cost.

Anycast IP Gateway With Multi-area SPB

Multi-area SPB boundary nodes can redistribute I-SID Attributes TLVs if you enable the Multi-area SPB Layer 2 VSN redistribution policy. The TLVs contain a metric field to advertise the SPB cost from a physical node to the Multi-area SPB virtual node to compare against the cost of routers in the local area.

The designated boundary node only redistributes the routing information for its closest Anycast IP Gateway router to the other SPB area for each Layer 2 VSN.

I-SID Attributes TLV 187

Anycast IP Gateway routers use the existing IS-IS LSDB and TLV 187 to advertise routing interface information to BEBs. This TLV has two sub-TLVs:

- Sub-TLV 1 defines the list of gateway MACs that the Anycast IP Gateway router advertises. This list includes the gateway MACs associated to VRIDs that you configure on the router.
- Sub-TLV 2 defines the Layer 2 VSN routing interfaces within Layer 3 VSNs that the Anycast IP Gateway router advertises. The MAC index field specifies which gateway MAC in sub-TLV 1 is the gateway MAC for a specific Layer 2 VSN.

Anycast IP Gateway Benefits

The following list outlines the benefits of Anycast IP Gateway:

- Anycast IP Gateway uses the router with the shortest path.
- Anycast IP Gateway runs over the existing SPB network.
- Anycast IP Gateway load shares traffic between routers.
- Anycast IP Gateway routers use the existing IS-IS LSDB to advertise routing interface information to BEBs.
- BEBs forward traffic through the Layer 2 VSN to the closest router by using the SPB cost.
- BEBs react to SPB cost changes, and automatically select the closest router as the next hop.
- Anycast IP Gateway offers an alternative to VRRP, RSMLT, and DvR in campus deployments.

Anycast IP Gateway Restrictions

The following list identifies feature restrictions:

- Anycast IP Gateway supports only IPv4.
- You can enable Anycast IP Gateway only on VLANs.
- You cannot enable Anycast IP Gateway on a vIST VLAN.

- You cannot configure VRRP or DvR on a Layer 2 VSN I-SID that has an Anycast IP Gateway interface. The BEB does not install the static entry if you configure DvR or VRRP on the VLAN associated with the Anycast IP Gateway Layer 2 VSN I-SID.
- Anycast IP Gateway does not support Secondary IP Interfaces.
- All BEB Fabric nodes must run a software image that supports Anycast IP Gateway and TLV 187. For more information, see *Fabric Engine and VOSS Feature Matrix*.
- To avoid duplicate MAC addresses with third-party devices that use a conflicting virtual router ID (VRID) in the subnet, you can assign up to 16 unique Anycast IP Gateway MAC addresses. If you do not encounter this situation, use common VRIDs as a best practice to avoid reaching the 16 Anycast IP Gateway MAC limit.

Anycast IP Gateway Configuration Using CLI

The topics in this section provide CLI procedures to configure and verify Anycast IP Gateway.

Configure the Anycast IP Gateway Router

Before You Begin

- Create the VLAN.
- Configure an IP address for the VLAN.
- Create the Layer 2 I-SID.
- Map the customer VLAN to the I-SID.
- Optionally, if you associate the VLAN with a VRF (Layer 3 VSN), enable IP-VPN on that VRF.

About This Task

Anycast IP Gateway routers advertise the configured routing interfaces, including the gateway MAC addresses and the gateway IP addresses. Associate the routing interfaces with Layer 2 VSNs, and configure each Anycast IP Gateway-enabled interface with a gateway IP and a gateway MAC.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. (Optional) Add the interface to a Layer 3 VSN:

```
vrf WORD<1-16>
```



Note

If you do not specify a VRF, the interface is added in the GRT.

3. Configure the gateway IP and a gateway MAC:

```
ip anycast-gateway {A.B.C.D} [vrid <1-255>]
```


4. Enable Anycast IP Gateway on the VLAN:

```
ip anycast-gateway enable
```

5. Display the routable MAC entry for the gateway MAC:

```
show vlan mac-address-entry <1-4059>
```

6. Display the local ARP entry for the gateway IP and the gateway MAC:

```
show ip arp vlan <1-4059>
```

7. Display the static gateway MAC entry:

```
show i-sid mac-address-entry <1-4059>
```

Example

Configure Anycast IP Gateway on VLAN 10, assuming all necessary prerequisites are complete:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 10
Switch:1(config-if)#ip anycast-gateway 192.0.2.254
Switch:1(config-if)#ip anycast-gateway enable
Switch:1(config-if)#exit
Switch:1(config)#show vlan mac-address-entry 10
=====
Vlan Fdb
=====
VLAN      MAC
ID  STATUS ADDRESS      INTERFACE  SMLT  REMOTE  TUNNEL
-----
10   self   00:00:5e:00:01:27 Port-mgmt   false -
10   self   d8:84:66:f6:dd:01 Port-cpp    false -
c: customer vid  u: untagged-traffic
2 out of 2 entries in all fdb(s) displayed.

Switch:1(config)#show i-sid mac-address-entry 10
=====
I-SID Fdb Table
=====
I-SID STATUS MAC-ADDRESS      INTERFACE TYPE  DEST-MAC      BVLAN DEST-SYSNAME AREA-ROLE AREA-NAME
-----
10   self   00:00:5e:00:01:27 Port-mgmt LOCAL  00:00:00:00:00:00 0      HOME      area-49.0010
10   self   d8:84:66:f6:dd:01 Port-cpp  LOCAL  00:00:00:00:00:00 0      HOME      area-49.0010
c: customer vid  u: untagged-traffic
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch r: discover by remote VIST switch
All 2 out of 2 Total Num of i-sid FDB Entries displayed

Switch:1(config)#show ip arp vlan 10
=====
IP Arp - GlobalRouter
=====
IP_ADDRESS      MAC_ADDRESS      VLAN  PORT  TYPE  TTL(10 Sec) TUNNEL  L3 ACCESS
-----
192.0.2.15       d8:84:66:f6:dd:01 10    -     LOCAL 2160          GlobalRouter
192.0.2.254      00:00:5e:00:01:27 10    -     LOCAL 2160          GlobalRouter
192.0.2.255      ff:ff:ff:ff:ff:ff 10    -     LOCAL 2160          GlobalRouter
c: customer vid  u: untagged-traffic
3 out of 3 ARP entries displayed
```

Variable Definitions

The following table defines parameters for the **ip anycast-gateway** command.

Variable	Value
<i>A.B.C.D</i>	Specifies the gateway IPv4 address. The gateway IP must be in different subnets for different interfaces, and must be the same on different routers for the same associated Layer 2 VSNs. You can configure a different gateway IP from the VLAN IP, or as a One IP interface to use the same IP as the VLAN. For more information, see Create a One IP Anycast IP Gateway Interface on page 476.
<i>enable</i>	Enables Anycast IP Gateway on the VLAN. The default is disabled.
<i>vrid <1-255></i>	Specifies the virtual router ID, which defines the gateway MAC as a VRRP MAC. The gateway MAC must be the same on different routers for the same associated Layer 2 VSN. If you do not specify the VRID as part of the command, the default value is VRID 39 (hexadecimal 27), which translates to the MAC of 00:00:5e:00:01:27.

Create a One IP Anycast IP Gateway Interface

Before You Begin

- Create the VLAN.
- Create the Layer 2 I-SID.
- Map the customer VLAN to the I-SID.
- Optionally, if you associate the VLAN with a VRF (Layer 3 VSN), enable IP-VPN on that VRF and configure a CLIP interface in the VRF.

About This Task

Configure the interface as a One IP interface, which means the gateway IP is the same as the VLAN IP. The gateway IP must be in different subnets for different interfaces, and must be the same on different routers for the same associated Layer 2 VSNs.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. (Optional) Add the interface to a Layer 3 VSN:

```
vrf WORD<1-16>
```

**Note**

If you do not specify a VRF, the interface is added in the GRT.

3. Configure the gateway IP and a gateway MAC:

```
ip anycast-gateway one-ip {A.B.C.D/X} [vrid <1-255>] [<0-511>]
```

4. Enable Anycast IP Gateway on the VLAN:

```
ip anycast-gateway enable
```

5. Display the routable MAC entry for the gateway MAC:

```
show vlan mac-address-entry <1-4059>
```

6. Verify that the local ARP entry is installed with the VLAN MAC address:

```
show ip arp vlan <1-4059>
```

Example

Configure a One IP interface on VLAN 11, assuming all necessary prerequisites are complete:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 11
Switch:1(config-if)#ip anycast-gateway one-ip 198.0.2.254/24
Switch:1(config-if)#ip anycast-gateway enable
Switch:1(config-if)#exit
Switch:1(config)#show vlan mac-address-entry 11
=====
Vlan Fdb
=====
VLAN      MAC
ID  STATUS ADDRESS      INTERFACE  SMLT  REMOTE  TUNNEL
-----
11   self   00:00:5e:00:01:27 Port-mgmt   false  -
11   self   d8:84:66:f6:dd:01 Port-cpp    false  -
c: customer vid  u: untagged-traffic
2 out of 2 entries in all fdb(s) displayed.

Switch:1(config)#show i-sid mac-address-entry 11
=====
I-SID Fdb Table
=====
I-SID STATUS MAC-ADDRESS      INTERFACE TYPE  DEST-MAC      BVLAN DEST-SYSNAME AREA-ROLE AREA-NAME
-----
11   self   00:00:5e:00:01:27 Port-mgmt LOCAL  00:00:00:00:00:00 0      HOME      area-49.0010
11   self   d8:84:66:f6:dd:01 Port-cpp  LOCAL  00:00:00:00:00:00 0      HOME      area-49.0010
c: customer vid  u: untagged-traffic
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
All 2 out of 2 Total Num of i-sid FDB Entries displayed

Switch:1(config)#show ip arp vlan 11
=====
IP Arp - VRF testvrf
=====
IP_ADDRESS  MAC_ADDRESS      VLAN  PORT  TYPE  TTL(10 Sec)  TUNNEL
-----
198.0.2.254 d8:84:66:f6:dd:01 11    -     LOCAL 2160
198.0.2.255 ff:ff:ff:ff:ff:ff 11    -     LOCAL 2160
c: customer vid  u: untagged-traffic
2 out of 2 ARP entries displayed
```

Variable Definitions

The following table defines parameters for the **ip anycast-gateway one-ip** command.

Variable	Value
<0-511>	Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated.
A.B.C.D/X	Specifies the gateway IPv4 address and subnet mask. The gateway IP must be in different subnets for different interfaces, and must be the same on different routers for the same associated Layer 2 VSNs.
enable	Enables Anycast IP Gateway on the VLAN. The default is disabled.
vrid <1-255>	Specifies the gateway MAC as a VRRP MAC. The gateway MAC must be the same on different routers for the same associated Layer 2 VSN. If you do not specify the VRID as part of the command, the default value is VRID 39 (hexadecimal 27), which translates to the MAC of 00:00:5e:00:01:27.

Display Anycast IP Gateway Interfaces

About This Task

Display the entries for local routing interfaces.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the interfaces:

```
show ip anycast-gateway interfaces [i-sid <1-15999999>] [vlan
<1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Display the interface entries:

```
Switch:1>show ip anycast-gateway interfaces
```

```
=====
                        Anycast interfaces - GlobalRouter
=====
```

L2VSN	VLAN	GW	IPv4	GW MAC	VRID	ADMIN STATUS	OPER STATE	ONE-IP	VRF ID	VRF Name
10	10	192.0.2.254	00:00:5e:00:01:27	39	enable	up	false	0	GlobalRouter	
11	11	198.0.2.254	00:00:5e:00:01:27	39	enable	up	true	0	GlobalRouter	

2 out of 2 total number of interfaces displayed

Display Anycast IP Gateway I-SIDs

About This Task

Display routing information for remote Anycast IP Gateway interfaces on BEBs.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the interfaces:

```
show ip anycast-gateway i-sid [<1-15999999>] [adv-sys-id
<xxxx.xxxx.xxxx>] [all-routers] [home] [remote]
```

Examples

Display routing information:

```
Switch:1>show ip anycast-gateway i-sid
=====
Anycast I-SID
=====
L2VSN  L3VSN  GW IPv4      GW MAC      PATH  ADD
COST  METRIC  NEXT-HOP-BMAC  AREA  ADV-SYS-ID  ADV-RTR
-----
10      0      192.0.2.254  00:00:5e:00:01:27  50    0      00:01:00:01:44:16  HOME  0001.0001.4416  R2
11      0      198.0.2.254  00:00:5e:00:01:27  50    0      00:01:00:01:44:15  HOME  0001.0001.4415  R1
-----
2 out of 2 total number of entries displayed
```

Display information for all advertising routers:

```
Switch:1>show ip anycast-gateway i-sid all-routers
=====
Anycast I-SID
=====
L2VSN  L3VSN  GW IPv4      GW MAC      PATH  ADD
COST  METRIC  NEXT-HOP-BMAC  AREA  ADV-SYS-ID  ACTIVE  ADV-RTR
-----
10      0      192.0.2.254  00:00:5e:00:01:27  50    0      00:01:00:01:44:15  HOME  0001.0001.4415  false  R1
10      0      192.0.2.254  00:00:5e:00:01:27  50    0      00:01:00:01:44:16  HOME  0001.0001.4416  true   R2
11      0      198.0.2.254  00:00:5e:00:01:27  50    0      00:01:00:01:44:15  HOME  0001.0001.4415  true   R1
11      0      198.0.2.254  00:00:5e:00:01:27  50    0      00:01:00:01:44:16  HOME  0001.0001.4416  false  R2
-----
4 out of 4 total number of entries displayed
```

View TLV Information for Anycast IP Gateway

About This Task

View the information that the Anycast IP Gateway router advertises.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information for TLV 187:

```
show isis lsdb tlv 187 detail
```

Example

Display information for TLV 187.

```
Switch:1>show isis lsdb tlv 187 detail
=====
ISIS LSDB (DETAIL)
=====
Level-1 LspID: 0001.0001.4415.00-02      SeqNum: 0x00000002      Lifetime: 1190
      Chksum: 0xf3b9  PDU Length: 62
      Host_name: R1
      Attributes:      IS-Type 1
TLV:187 SUB-TLV 1 Anycast GW MAC:
      Index 1: 00:00:5e:00:01:27

      SUB-TLV 2 I-SID EXTENDED ATTRIBUTES:
        L3VSN I-SID: 0
          L2VSN I-SID: 10
            Anycast
            MAC index: 1
            IPv4 address: 192.0.2.254

          L2VSN I-SID: 11
            Anycast
            MAC index: 1
            IPv4 address: 198.0.2.254
```

Anycast IP Gateway Configuration Using EDM

The topics in this section provide EDM procedures to configure and verify Anycast IP Gateway.

Create an Anycast IP Gateway Interface

Before You Begin

- Create the VLAN.
- Configure an IP address for the VLAN.
- Create the Layer 2 VSN I-SID. If the I-SID you manually type in this task does not exist, creation of the Anycast IP Gateway interface fails. If you are uncertain which I-SIDs already exist, or to avoid typing errors, you can create an interface using [Create an Anycast IP Gateway Interface On a Specific I-SID](#) on page 482.
- Map the customer VLAN to the I-SID.
- Optionally, associate the VLAN with a VRF (Layer 3 VSN), and enable IP-VPN on that VRF.
- For a Onelp interface, configure a CLIP interface in the corresponding VRF.

About This Task

Anycast IP Gateway routers advertise the configured routing interfaces, including the gateway MAC addresses and the gateway IP addresses. Associate the routing interfaces with Layer 2 VSNs, and configure each Anycast IP Gateway-enabled interface with a gateway IP and a gateway MAC.

In addition to adding new interfaces through the **Insert Interfaces** dialog, you can use the table-based **Interfaces** tab to modify existing configuration of all Anycast IP Gateway interfaces.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Anycast Gateway**.
3. Select **Interfaces**.
4. Select **Insert**.
5. In **L2vsnlid**, type the I-SID to associate with the interface.
6. In **GwAddr**, type the gateway IP address.
7. (Optional) In **Vrld**, type the virtual router ID for the gateway MAC.
8. (Optional) Select **Onelp** to make the gateway IP the same as the VLAN IP.
9. (Optional) For a One IP interface, specify the prefix length for the address in **PrefixLen**.
10. Select **Enable**.
11. Select **Insert**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
L2vsnlid	Specifies the Layer 2 VSN I-SID associated with the Anycast IP Gateway routing interface.
GwAddrType	Displays the address type. Currently, only IPv4 is supported. This parameter only displays on the Insert Interfaces dialog.
GwAddr	Specifies the gateway IPv4 address. The gateway IP must be in different subnets for different interfaces, and must be the same on different routers for the same associated Layer 2 VSNs.
PrefixLen	Specifies the prefix length for a One IP address.
Vrld	Specifies the virtual router ID, which defines the gateway MAC as a VRRP MAC. The gateway MAC must be the same on different routers for the same associated Layer 2 VSN. If you do not specify the VRID as part of the command, the default value is VRID 39 (hexadecimal 27), which translates to the MAC of 00:00:5e:00:01:27.
GwMac	Displays the Anycast IP Gateway MAC address.
Vrflid	Displays the VRF ID associated with the Anycast IP Gateway routing interface.
VlanId	Displays the VLAN ID associated with the Anycast IP Gateway routing interface.
Onelp	Configures the routing interface as a One IP interface. You can configure a different gateway IP from the VLAN IP, or as a One IP interface to use the same IP as the VLAN.

Name	Description
Enable	Configures if the Anycast IP Gateway routing interface is enabled. The default is disabled (false).
Operating State	Displays the operating state of the Anycast IP Gateway routing interface: up or down.

Create an Anycast IP Gateway Interface On a Specific I-SID

Before You Begin

- Create the VLAN.
- Configure an IP address for the VLAN.
- Create the Layer 2 VSN I-SID.
- Map the customer VLAN to the I-SID.
- Optionally, associate the VLAN with a VRF (Layer 3 VSN), and enable IP-VPN on that VRF.
- For a Onelp interface, configure a CLIP interface in the corresponding VRF.

About This Task

Anycast IP Gateway routers advertise the configured routing interfaces, including the gateway MAC addresses and the gateway IP addresses. Associate the routing interfaces with Layer 2 VSNs, and configure each Anycast IP Gateway-enabled interface with a gateway IP and a gateway MAC.

In addition to adding new interfaces through the **Insert Anycast GW** dialog, you can use the table-based **Anycast GW** tab to modify existing configuration on a specific I-SID.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. On the **Service** tab, select an existing Layer 2 VSN I-SID.
4. Select **IP**.
5. On the **Anycast GW** tab, select **Insert**.
6. In **GwAddr**, type the gateway IP address.
7. (Optional) In **Vrld**, type the virtual router ID for the gateway MAC.
8. (Optional) Select **Onelp** to make the gateway IP the same as the VLAN IP.
9. (Optional) For a One IP interface, specify the prefix length for the address in **PrefixLen**.
10. Select **Enable**.
11. Select **Insert**.

Anycast GW *Field Descriptions*

Use the data in the following table to use the **Anycast GW** tab.

Name	Description
L2vsnlSid	Specifies the Layer 2 VSN I-SID associated with the Anycast IP Gateway routing interface.
GwAddrType	Displays the address type. Currently, only IPv4 is supported. This parameter only displays on the Insert Anycast GW dialog.
GwAddr	Specifies the gateway IPv4 address. The gateway IP must be in different subnets for different interfaces, and must be the same on different routers for the same associated Layer 2 VSNs.
PrefixLen	Specifies the prefix length for a One IP address.
VrId	Specifies the virtual router ID, which defines the gateway MAC as a VRRP MAC. The gateway MAC must be the same on different routers for the same associated Layer 2 VSN. If you do not specify the VRID as part of the command, the default value is VRID 39 (hexadecimal 27), which translates to the MAC of 00:00:5e:00:01:27.
GwMac	Displays the Anycast IP Gateway MAC address.
VrfId	Displays the VRF ID associated with the Anycast IP Gateway routing interface.
VlanId	Displays the VLAN ID associated with the Anycast IP Gateway routing interface.
Onelp	Configures the routing interface as a One IP interface. You can configure a different gateway IP from the VLAN IP, or as a One IP interface to use the same IP as the VLAN.
Enable	Configures if the Anycast IP Gateway routing interface is enabled. The default is disabled (false).
Operating State	Displays the operating state of the Anycast IP Gateway routing interface: up or down.

Display Anycast IP Gateway I-SIDs for the Home Area

About This Task

Display routing information for Anycast IP Gateway interfaces on BEBs.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Anycast Gateway**.
3. Select **ISID**.

ISID Field Descriptions

Use the data in the following table to use the **ISID** tab.

Name	Description
L2vsn Isid	Displays the Layer 2 VSN I-SID.
Adv Sys Id	Displays the IS-IS system ID of the advertising node.
Address Type	Displays the address type. Currently, only IPv4 is supported.
Address	Displays the IPv4 gateway address.
L3vsn Isid	Displays the Layer 3 VSN I-SID.
Mac	Displays the gateway MAC.
Path Cost	Displays the SPB path cost to the next hop.
Add Metric	Displays the SPB cost advertised by the Multi-area SPB designated boundary node.
Next-Hop BMAC	Displays the BMAC of the next hop.
Rtr HostName	Displays the host name of the advertising node.
Is Gw Node	Displays if the advertising node is the installed Anycast IP Gateway node.

Display Anycast IP Gateway I-SIDs for the Remote Area

**Note**

This procedure only applies to a switch with boundary-node abilities.

About This Task

Display routing information for Anycast IP Gateway interfaces on BEBs.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Anycast Gateway Remote**.
3. Select **ISID**.

ISID Field Descriptions

Use the data in the following table to use the **ISID** tab.

Name	Description
L2vsn Isid	Displays the Layer 2 VSN I-SID.
Adv Sys Id	Displays the IS-IS system ID of the advertising node.
Address Type	Displays the address type. Currently, only IPv4 is supported.
Address	Displays the IPv4 gateway address.
L3vsn Isid	Displays the Layer 3 VSN I-SID.

Name	Description
Mac	Displays the gateway MAC.
Path Cost	Displays the SPB path cost to the next hop.
Add Metric	Displays the SPB cost advertised by the Multi-area SPB designated boundary node.
Next-Hop BMAC	Displays the BMAC of the next hop.
Rtr HostName	Displays the host name of the advertising node.
Is Gw Node	Displays if the advertising node is the installed Anycast IP Gateway node.



Bidirectional Forwarding Detection

[BFD Fundamentals](#) on page 487

[BFD Configuration using CLI](#) on page 491

[BFD Configuration using EDM](#) on page 506

Table 55: Bidirectional Forwarding Detection (BFD) product support

Feature	Product	Release introduced
BFD (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.1

Table 55: Bidirectional Forwarding Detection (BFD) product support (continued)

Feature	Product	Release introduced
BFD (IPv6)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 demonstration feature
	5420 Series	VOSS 8.4 demonstration feature
	5520 Series	VOSS 8.2.5 demonstration feature
	5720 Series	Fabric Engine 8.7 demonstration feature
	7520 Series	Fabric Engine 8.10 demonstration feature
	7720 Series	Fabric Engine 8.10 demonstration feature
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1 demonstration feature
	VSP 7400 Series	VOSS 8.1 demonstration feature
BFD over Fabric Extend Tunnels (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.2
	VSP 7400 Series	VOSS 8.2

Use Bidirectional Forwarding Detection (BFD) to provide a failure detection mechanism between two systems.

The following sections provide information and procedures for BFD.

BFD Fundamentals

The following sections provide fundamentals information about Bidirectional Forwarding Detection (BFD).

BFD Overview

Bidirectional Forwarding Detection (BFD) is a simple Hello protocol used between two peers. In BFD, peer systems periodically transmit BFD packets to each other. If one of the systems does not receive a BFD packet after a certain period of time, the system assumes that the link or other system is not operating.

A path is considered operational when bidirectional communication is established between systems. However, this does not preclude the use of unidirectional links.

BFD provides low-overhead, short-duration failure detection between two systems. BFD also provides a single mechanism for connectivity detection over any media, at any protocol layer.

Because BFD sends rapid failure-detection notifications to the routing protocols that run on the local system, which initiates routing table recalculations, BFD helps reduce network convergence time.

BFD supports IPv4/IPv6 single hop detection for static routes, OSPFv2, OSPFv3, iBGP, iBGPv6. Forwarding path failure detection for Fabric Extend tunnels is supported over an IPv4 network only.



Note

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.



Note

iBGPv6 is not supported in VRF.

BFD Operation

The switch uses one BFD session for all protocols with the same destination. For example, if a network runs OSPFv2 and BGP across the same link with the same peer, only one BFD session is established, and BFD shares session information with both routing protocols.

You can enable BFD over data paths with specified OSPFv2 and OSPFv3 neighbors, BGP neighbors, static routing next-hop addresses, and Fabric Extend tunnels.

The switch supports BFD asynchronous mode, which sends BFD control packets between two systems to activate and maintain BFD neighbor sessions. To reach an agreement with its neighbor about how rapidly failure detection occurs, each system estimates how quickly it can send and receive BFD packets.

A session begins with the periodic, slow transmission of BFD control packets. When bidirectional communication is achieved, the BFD session comes up.

After the session is up, the transmission rate of Control packets can increase to achieve detection time requirements. If Control packets are not received within the calculated detection time, the session is declared down. After a session is down, Control packet transmission returns to the slow rate.

If a session is declared down, it cannot come back up until the remote end signals that it is down (three-way handshake). A session can be kept administratively down by configuring the state of AdminDown.

In asynchronous mode, detection time is equal to the value of DetectMult received from the remote system multiplied by the agreed transmit interval of the remote system (the greater of RequiredMinRxInterval and DesiredMinTxInterval.) DetectMult is approximately equal to the number of sequential packets that must be missed to declare a session down.

BFD States

A session normally proceeds through three states; two states are used to establish a session (Init and Up) and one state is used to tear down a session (Down). This allows a three-way handshake for both session establishment and session teardown, assuring that both systems are aware of all session state changes. There is a fourth state (AdminDown) that you can use to administratively put a session down indefinitely.

- Down state: Indicates the session is down or has just been created. The session will remain in Down state until the remote system sends a BFD control packet indicating anything other than Up state. If the control packet signals Down state, the session advances to Init state. If the control packet signals Init state, the session advances to Up state.
- Init state: In this state, the host system establishes communications with the remote system and sends a request to move the session to the Up state, but the remote system has not yet recognized the request. A session remains in Init state until it receives a BFD control packet signaling Init or Up state, or until the connectivity timer expires, indicating communication with the remote system is lost.
- Up state: Indicates the BFD session is established and connectivity is working. A session remains in Up state until connectivity fails or until the session is taken down administratively.
- AdminDown state: Indicates the BFD session is being held down administratively. This causes the remote system to enter Down state and remain there until the local system exits AdminDown state.

BFD Configuration

The following sections provide conceptual information about BFD configuration. For detailed procedural information about BFD configuration, see [BFD Configuration using CLI](#) on page 491 and [BFD Configuration using EDM](#) on page 506.

Enable BFD

To enable Bidirectional Forwarding Detection (BFD) between 2 peers:

1. Configure BFD globally.
2. Configure BFD on the required interfaces of both peer systems.
3. Enable BFD on the required routing protocols.
4. Specify the next-hop device with which the switch initiates the BFD session.

Delete a BFD Session

To delete a BFD session, disassociate all applications with the BFD session, then administratively bring down the BFD session.

**Note**

To successfully delete a BFD session, you must run the commands in the following order:

1. Disassociate all applications from the BFD session.
2. Disable BFD at the global level or interface level, which transitions the BFD session to the `AdminDown` state.

If you change the order of operations, the BFD session is not deleted.

BFD Considerations

The following considerations apply to Bidirectional Forwarding Detection (BFD):

- BFD is supported only in asynchronous mode. Demand mode and echo functionalities are not supported.
- You configure BFD parameters on a per session basis, not on a per next-hop basis.
- BFD creates multiple sessions even though a neighbor shares an IP address.
- The granularity of the fault detection interval in BFD is 100 ms, and the minimum multiplier is 2.

The minimum value for the transmit interval or the receive interval is 100 ms. If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.

You can configure a total of 16 BFD sessions. Of the 16 possible BFD sessions, you can configure a maximum of 4 BFD sessions with the minimum value for transmit interval or receive interval. You can configure the remaining BFD sessions with a transmit interval or a receive interval that is greater than or equal to the 200 ms default value.

- BFD is not supported over RSMLT links. This applies to BFD sessions over IPv4 interfaces and IPv6 interfaces.
- Inter-tunnel routing with 6in4 tunnels is not supported. This means that incoming IPv6 packets over a tunnel cannot be forwarded over another tunnel configured on the same switch.
- BFD for Interior Border Gateway Protocol (iBGP) and BGPv6 in VRF is not supported.
- BFD for eBGPv6 in VRF is not supported.
- Session dampening is not supported for BFD.
- The switch supports BFD multihop only at the eBGP application level. For other applications, the switch does not support BFD multihop, as defined by RFC 5883. However, there is no requirement for source and destination IP addresses to be in the same subnet.
- BFD over IPv6 Fabric Extend (FE) tunnels is not supported.

- The minimum value for the transmit interval or the receive interval is 1 second with a fault detection time of 3 seconds for BFD over IPv4 FE tunnels.
- BFD does not support a static route flag.
- BFD is not supported on a Virtual Router Redundancy Protocol (VRRP) interface.
- You can configure a total of 256 BFD and Virtual Link Aggregation Control Protocol (VLACP) sessions.
- BFD packets cannot be mirrored when BFD is configured on the switch.
- You cannot enable BFD on a logical interface for a virtual NNI link on boundary nodes.

BFD Configuration using CLI

Use the following procedures to configure Bidirectional Forwarding Detection (BFD) using CLI. BFD provides low-overhead, short-duration failure-detection between two systems.

Enable BFD Globally

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD globally.



Note

Enabling BFD globally does not establish a BFD session. To establish a BFD session, you must also configure BFD at the interface level and at the application level.

Procedure

1. Enter BFD Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bfd
```

2. Enable BFD:

```
router bfd enable
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router bfd
Switch:1(router-bfd)#router bfd enable
```

Configure BFD on an IPv4 Interface

About This Task

Use the following procedure to enable and to configure Bidirectional Forwarding Detection (BFD) on an IPv4 interface. All interface configuration is performed at the VLAN, GigabitEthernet, or Loopback level.



Note

Enabling BFD on an interface does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the application level.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

followed by one of the following:

- `interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}`
- `interface loopback <1-256>`
- `interface vlan <1-4059>`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable BFD on an interface:

```
ip bfd enable
```

3. (Optional) Configure the transmit interval:

```
ip bfd interval <100-65335>
```

4. (Optional) Configure the minimum receive interval:

```
ip bfd min-rx <100-65335>
```

5. (Optional) Configure the multiplier:

```
ip bfd multiplier <1-20>
```

6. (Optional) In GigabitEthernet Interface Configuration mode, you can configure a value for port:

```
ip bfd port {slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}
```

7. (Optional) In VLAN Interface Configuration mode, you can configure a value for VLAN:

```
ip bfd vlan <1-4094>
```

8. (Optional) In Loopback Interface Configuration mode, you can configure a value for loopback:

```
ip bfd loopback <1-256>
```

Variable Definitions

The following table defines parameters for the **ip bfd** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>enable</code>	Enable BFD on a port, VLAN, or loopback.
<code>interval <100-65335></code>	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
<code>min-rx <100-65535></code>	<p>Specifies the receive interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>
<code>multiplier <1-20></code>	<p>Specifies the multiplier used to calculate the amount of time BFD waits before declaring a receive timeout. The default is 3.</p> <p>Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.</p>

Variable	Value
port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
vlan <1-4094>	Specifies the VLAN ID.
loopback <1-256>	Specifies the Loopback ID.

Configure BFD on an IPv6 Interface

About This Task



Note

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Use the following procedure to enable and to configure BFD on an IPv6 interface. All interface configuration is performed at the VLAN or GigabitEthernet level.



Note

Enabling BFD on an interface does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the application level.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[, ...] [slot/all] [all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable BFD on an interface:

```
ipv6 bfd enable
```

3. (Optional) Configure the transmit interval:

```
ipv6 bfd interval <100-65335>
```

4. (Optional) Configure the minimum receive interval:

```
ipv6 bfd min-rx <100-65335>
```

5. (Optional) Configure the multiplier:

```
ipv6 bfd multiplier <1-20>
```

6. (Optional) In GigabitEthernet Interface Configuration mode, you can configure a value for port:

```
ipv6 bfd port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}
```

7. (Optional) In VLAN Interface Configuration mode, you can configure a value for VLAN:

```
ipv6 bfd vlan <1-4094>
```

Variable Definitions

The following table defines parameters for the **ip bfd** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
enable	Enable BFD on a port, VLAN, or loopback.
interval <100-65335>	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
min-rx <100-65535>	<p>Specifies the receive interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>

Variable	Value
multiplier <1-20>	Specifies the multiplier used to calculate the amount of time BFD waits before declaring a receive timeout. The default is 3. Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.
port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vlan <1-4094>	Specifies the VLAN ID.
loopback <1-256>	Specifies the Loopback ID.

Enable BFD at the BGP Application Level

About This Task

BFD supports internal Border Gateway Protocol (iBGP) and external Border Gateway Protocol (eBGP) on IPv4 interfaces. You configure BFD on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. BFD does not support BGPv6 for VRF on IPv6 interfaces.



Note

Enabling BFD at the BGP application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. Enter BGP Router Configuration mode:
enable

configure terminal

router bgp
2. Enable BFD for the BGP protocol:
neighbor WORD<0-1536> fall-over bfd

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config)#router bgp
Switch:1(router-bgp)#neighbor 192.0.2.15 fall-over bfd
```

Variable Definitions

The following table defines parameters for the **neighbor** command.

Variable	Value
WORD<0-1536>	Specifies the peer IP address or the peer group name.

Enable BFD at the OSPF Application Level

About This Task

BFD supports Open Shortest Path First (OSPF) for IPv4 interfaces and OSPFv3 for IPv6 interfaces.

Use the following procedure to enable BFD at the OSPF application level.



Note

Enabling BFD at the OSPF application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. (Optional) Enable BFD on an IPv4 interface under the OSPF protocol:

```
ip ospf bfd
```

3. Enable BFD on an IPv6 interface under the OSPF protocol:

```
ipv6 ospf bfd
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/3
Switch:1(config-if)#ip ospf bfd
```

Variable Definitions

The following table defines parameters for the **ip ospf bfd** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub- port]][,...][slot/ all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code><1-4059></code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p>

Configure BFD on an IPv4 Static Route

About This Task

Use the following procedure to configure BFD on an IPv4 static route.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Configure BFD on an IPv4 static route:


```
ip route bfd {A.B.C.D}
```

Variable Definitions

The following table defines parameters for the **ip route bfd** command.

Variable	Value
<code>{A.B.C.D}</code>	Specifies the BFD static route IPv4 address.

Configure BFD on an IPv6 Static Route

About This Task



Note

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Use the following procedure to configure BFD on an IPv6 static route.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure BFD on an IPv6 static route:

```
ipv6 route bfd WORD<0-128>
```
3. (Optional) Configure an IPv6 static route for a port:

```
ipv6 route bfd WORD<0-128> port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}
```
4. (Optional) Configure an IPv6 static route for a VLAN:

```
ipv6 route bfd WORD<0-128> vlan <1-4094>
```

Variable Definitions

The following table defines parameters for the **ipv6 route bfd** command.

Variable	Value
WORD<0-128>	Specifies the BFD static route IPv6 address.
port {slot/port[-slot/port] [, ...] [slot/all] [all]}	Specifies the port number for the BFD IPv6 static route.
vlan <1-4094>	Specifies the VLAN ID for the BFD IPv6 static route.

Clear BFD Session Statistics

About This Task

Use the following procedure to clear local and remote Bidirectional Forwarding Detection (BFD) session statistics for IPv4 or IPv6 interfaces.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. (Optional) Clear BFD session statistics for an IPv4 interface:

```
clear ip bfd stats
```

3. Clear BFD session statistics for an IPv6 interface:

```
clear ipv6 bfd stats
```

Variable Definitions

The following table defines parameters for the **clear ip bfd stats** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies a VRF instance by VRF name.
<i>vrfids</i> WORD<0-512>	Specifies a VRF or range of VRFs by ID.

Display BFD Global Configuration

About This Task

Use this procedure to display global configuration information for BFD.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display global BFD configuration information:

```
show ip bfd [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>show ip bfd
=====
                        BFD information - GlobalRouter
=====
                        BFD Version : 1
                        Admin Status : TRUE
                        Trap Enable  : FALSE
=====
Total session number : 1

UP: 1, DOWN: 0, AdminDown: 0, Init: 0
=====
```

Variable Definitions

The following table defines parameters for the **show ip bfd** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies a VRF instance by VRF name.
<i>vrfids</i> WORD<0-512>	Specifies a VRF or range of VRFs by ID.

Display BFD Configuration for an IPv4 Interface

About This Task

Use the following procedure to display BFD configuration on an interface.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display BFD on a Gigabit Ethernet interface:

```
show ip bfd interfaces GigabitEthernet [{slot/port[/sub-port]}[-  
slot/port[/sub-port]][,...][slot/all][all]] [vrf WORD<1-16>] [vrfids  
WORD<0-512>]
```

3. Display BFD on a VLAN interface:

```
show ip bfd interfaces vlan [<1-4059>] [vrf WORD<1-16>] [vrfids  
WORD<0-512>]
```

Examples

The following example displays VLAN interface configuration information for BFD.

```
Switch:1>show ip bfd interfaces vlan 11
```

Vlan Bfd					
VLAN	STATUS	MIN_RX	INTERVAL	MULTIPLIER	VRF-ID
11	enable	200	200	3	0

The following example displays Loopback interface configuration information for BFD:

```
Switch:1>show ip bfd interfaces loopback
```

Circuitless IP Interface Bfd					
INTF ID	STATUS	MIN_RX	INTERVAL	MULTIPLIER	VRF-ID
1	enable	200	200	3	0
2	enable	200	200	3	2

Variable Definitions

The following table defines parameters for the **show ip bfd interfaces** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD<1-16></code>	Specifies a VRF instance by VRF name.
<code>vrfids WORD<0-512></code>	Specifies a VRF or range of VRFs by ID.

Display BFD Configuration for an IPv6 Interface

About This Task



Note

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Use the following procedure to display BFD configuration on an IPv6 interface.

Procedure

- To enter User EXEC mode, log on to the switch.
- Display BFD on a Gigabit Ethernet interface:

```
show ipv6 bfd interfaces GigabitEthernet [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]  
[slot/all] [all]}
```
- Display BFD on a VLAN interface:

```
show ipv6 bfd interfaces vlan <1-4059>
```

Example

```
Switch:1>show ipv6 bfd interfaces gigabitethernet 1/3
```

```
=====
Port Bfd
=====
PORT      STATUS    MIN_RX    INTERVAL  MULTIPLIER  VRF-ID
```

```
-----
1/3      enable      200      200      3      0
```

Variable Definitions

The following table defines parameters for the **show ip bfd interfaces** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD<1-16></code>	Specifies a VRF instance by VRF name.
<code>vrfids WORD<0-512></code>	Specifies a VRF or range of VRFs by ID.

Display BFD IPv4 Neighbor Information

About This Task

Use this procedure to display BFD session information for IPv4 neighbors.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display BFD neighbor information:

```
show ip bfd neighbors
```
3. (Optional) Display BFD neighbor next-hop information:

```
show ip bfd neighbors next-hop {A.B.C.D}
```
4. (Optional) Display BFD neighbor information for a particular VRF:

```
show ip bfd neighbors vrf WORD<1-16>
```
5. (Optional) Display BFD neighbor information for a VRF ID or a range of VRF IDs:

```
show ip bfd neighbors vrfids WORD<0-512>
```

Example

```
Switch:1>show ip bfd neighbors
```

```
-----
BFD Session - GlobalRouter
-----
```

MY_DISC	YOUR_DISC	NEXT_HOP	STATE	MULTI	MIN_TX	MIN_RX	ACT_TX	DETECT_TIME	REMOTE_STATE	APP	RUN
1	0	192.0.2.11	Down	3	200	200	1000	600	Down	O	

1 out of 1 BFD session displayed											

APP and RUN Legend:											
B=BGP, O=OSPF, S=Static Route											

Variable Definitions

The following table defines parameters for the **show ip bfd neighbors** command.

Variable	Value
{A.B.C.D}	Specifies the next-hop IP address in the format a.b.c.d.
vrf WORD<1-16>	Specifies a VRF instance by name.
vrfids WORD<0-512>	Specifies a range of VRFs by ID number.

Display BFD IPv6 Neighbor Information

About This Task



Note

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Use this procedure to display information about BFD IPv6 neighbors.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display BFD neighbor information:

```
show ipv6 bfd neighbors
```
3. (Optional) Display BFD neighbor next-hop information:

```
show ipv6 bfd neighbors next-hop WORD<0-128>
```
4. (Optional) Display BFD neighbor information for a particular VRF:

```
show ipv6 bfd neighbors vrf WORD<1-16>
```
5. (Optional) Display BFD neighbor information for a range of VRFs:

```
show ipv6 bfd neighbors vrfids WORD<0-512>
```

Example

```
Switch:1>show ipv6 bfd neighbors
=====
BFD Session - GlobalRouter
=====
MY_DISC  YOUR_DISC  NEXT_HOP                STATE  MULTI  MIN_TX  MIN_RX  ACT_TX  DETECT_TIME  REMOTE_STATE  APP  RUN
1         0          2001:DB8:0:0:25AB:0:0:1  Down   3      200    200    1000    0          Down         O    
```

1 out of 1 BFD session displayed

APP and RUN Legend:

B=BGP_IPv6, O=OSPFv3, S=IPv6 Static Route

Variable Definitions

The following table defines parameters for the **show ipv6 bfd neighbors** command.

Variable	Value
WORD<0-128>	Specifies the next-hop IPv6 address in the format a:b:c:d:e:f:g:h.
vrf WORD<1-16>	Specifies a VRF instance by name.
vrfids WORD<0-512>	Specifies a range of VRFs by ID number.

Display BFD Statistics

About This Task

Use the following procedure to display BFD statistics for IPv4 or IPv6 interfaces.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display BFD IPv4 statistics:

```
show ip bfd stats [vrf] [vrfids]
```
3. Display BFD IPv6 statistics:

```
show ipv6 bfd stats [vrf] [vrfids]
```

Example

The following example displays BFD statistics for IPv4 interfaces.

```
Switch:1>show ip bfd stats
=====
BFD staticstics - GlobalRouter
=====
MY_DISC  YOUR_DISC  NEXT_HOP  PACKET_IN  PACKET_OUT  LAST_UP  LAST_DOWN
-----
1         0         192.0.2.10  4661750    4620630    Mon Sep  6 15:31:15 2021  Mon Sep  6 15:28:08 2021
=====
```

The following example displays BFD statistics for IPv6 interfaces.

```
Switch:1>show ipv6 bfd stats
=====
BFD staticstics - GlobalRouter
=====
MY_DISC  YOUR_DISC  NEXT_HOP  PACKET_IN  PACKET_OUT  LAST_UP  LAST_DOWN
-----
1         0         2001:DB8:0:0:0:0:0:ffff 4661750    4620630    Mon Sep  6 15:31:15 2021  Mon Sep  6 15:28:08 2021
=====
```

Variable Definitions

The following table defines parameters for the **show ip bfd stats** command.

Variable	Value
vrf	Specifies a VRF instance by VRF name.
vrfids	Specifies a VRF or range of VRFs by ID.

BFD Configuration using EDM

Use the following procedures to configure Bidirectional Forwarding Detection (BFD) using EDM. BFD provides low-overhead, short-duration failure-detection between two systems.

Enable BFD Globally

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD globally.



Note

Enabling BFD globally does not establish a BFD session. To establish a BFD session, you must enable BFD at the interface level and at the application level.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **BFD**.
3. Select the **Globals** tab.
4. In **AdminStatus**, select **enabled**.
5. (Optional) Select **TrapEnabled** to send BFD traps.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminStatus	Specifies whether BFD is enabled.
VersionNumber	Specifies the current version number of the BFD protocol.
TrapEnabled	Specifies whether BFD traps are sent.

Display BFD Sessions

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to display information about BFD sessions. You can optionally display BFD session information for IPv4 or IPv6 interfaces.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **BFD**.
3. Select the **Sessions** tab.
4. (Optional) Select **Filter**.
5. (Optional) Select **AddrType**.
6. (Optional) In **AddrType**, specify a value for address type.

Sessions Field Descriptions

Use the data in the following table to use the **Sessions** tab.

Name	Description
Discriminator	Specifies the local discriminator that uniquely identifies the BFD session.
RemoteDiscr	Specifies the discriminator of the remote system in the BFD session.
UdpPort	Specifies the UDP Port for the BFD session. The default value is the well-known value for the port.
State	Specifies the state of the BFD session. Possible values are Down, Up, Init, and AdminDown.

Name	Description
Addr	Specifies the IP address of the interface associated with the BFD session. A value of unknown (0) indicates the BFD session is not associated with a specific interface.
DesiredMinTxInterval	Specifies the preferred minimum interval for transmitting BFD control packets by the local system.
ReqMinTxInterval	Specifies the minimum interval for transmitting BFD control packets that the local system can support.
DestAddr	Specifies the destination IP address of the interface associated with the BFD session.
OldState	Specifies the old state of the BFD session.
App	Specifies the applications configured on the BFD session.
AppRun	Specifies the applications running on the BFD session.
AddrType	Specifies the IP address type of the interface associated with this BFD session. Possible values are ipv4 and ipv6.

Configure BFD for an IPv4 Interface on a Port

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable and configure BFD for an IPv4 interface on a port.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IP**.
3. Select the **BFD** tab.
4. Select **Enable**.
5. (Optional) In **MinRxInterval**, specify the minimum receive interval.
6. (Optional) In **TxInterval**, specify the transmit interval.
7. (Optional) In **Multiplier**, specify a value for the multiplier used to calculate a receive timeout.

BFD Field Descriptions

Use the data in the following table to use the **BFD** tab.

Name	Description
Enable	Enable BFD on the port.
MinRxInterval	<p>Specifies the minimum interval, in milliseconds, between received BFD control packets that the local system is capable of supporting. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>
TxInterval	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
Multiplier	<p>Specifies a value for the multiplier used to calculate a receive timeout. The default is 3.</p> <p>Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.</p>

Configure BFD for an IPv6 Interface on a Port

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable and configure BFD for an IPv6 interface on a port.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IPv6**.
3. Select the **IPv6 BFD Interface** tab.
4. (Optional) In the **MinRxInterval** column, double-click the field and type a value for **MinRxInterval**.
5. (Optional) In the **TxInterval** column, double-click the field and type a value for **TxInterval**.
6. (Optional) In the **Multiplier** column, double-click the field and type a value for **Multiplier**.

BFD Field Descriptions

Use the data in the following table to use the **BFD** tab.

Name	Description
Interface	Specifies the BFD interface.
MinRxInterval	<p>Specifies the minimum interval, in milliseconds, between received BFD control packets that the local system is capable of supporting. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>
TxInterval	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
Multiplier	<p>Specifies a value for the multiplier used to calculate a receive timeout. The default is 3.</p> <p>Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.</p>

Enable BFD on a CLIP interface

Before You Begin

- Create the CLIP Interface.

About This Task

Enable Bidirectional Forwarding Detection (BFD) over Fabric Extend tunnels on a CLIP interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select the required CLIP interface.
5. Select **BFD**.
6. Select **Enable**.

7. (Optional) In **MinRxInterval**, specify the minimum receive interval.
8. (Optional) In **TxInterval**, specify the transmit interval.
9. (Optional) In **Multiplier**, specify the multiplier used to calculate a receive timeout.
10. Select **Apply**.

BFD Field Descriptions

Use the data in the following table to use the **BFD** tab.

Name	Description
Enable	Enable BFD on the CLIP interface.
MinRxInterval	<p>Specifies the minimum interval, in milliseconds, between received BFD control packets that the local system is capable of supporting. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>
TxInterval	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
Multiplier	<p>Specifies a value for the multiplier used to calculate a receive timeout. The default is 3.</p> <p>Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.</p>

Configure BFD for an IPv4 Interface on a VLAN

About This Task

BFD provides a failure detection-mechanism between two systems. Use the following procedure to enable and configure BFD for an IPv4 interface on a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN on which you want to configure BFD.

5. Select **IP**.
6. Select **BFD**.
7. Select **Enable**.
8. (Optional) In **MinRxInterval**, specify the minimum receive interval.
9. (Optional) In **TxInterval**, specify the transmit interval.
10. (Optional) In **Multiplier**, specify a value for the multiplier used to calculate a receive timeout.

BFD Field Descriptions

Use the data in the following table to use the **BFD** tab.

Name	Description
Enable	Enable BFD on the VLAN.
MinRxInterval	<p>Specifies the minimum interval, in milliseconds, between received BFD control packets that the local system is capable of supporting. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>
TxInterval	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
Multiplier	<p>Specifies a value for the multiplier used to calculate a receive timeout. The default is 3.</p> <p>Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.</p>

Configure BFD for an IPv6 Interface on a VLAN

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable and configure BFD for an IPv6 interface on a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.

2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN on which you want to configure BFD.
5. Select **IPv6**.
6. Select **IPv6 BFD Interface**.
7. (Optional) In the **MinRxInterval** column, double-click the field and type a value for **MinRxInterval**.
8. (Optional) In the **TxInterval** column, double-click the field and type a value for **TxInterval**.
9. (Optional) In the **Multiplier** column, double-click the field and type a value for **Multiplier**.

IPv6 BFD Interface Field Descriptions

Use the data in the following table to use the **IPv6 BFD Interface** tab.

Name	Description
Interface	Specifies an index value that uniquely identifies the interface.
Enable	Enable BFD on the VLAN.
MinRxInterval	<p>Specifies the minimum interval, in milliseconds, between received BFD control packets that the local system is capable of supporting. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.</p>
TxInterval	<p>Specifies the transmit interval in milliseconds. The default is 200 ms.</p> <p>Note: The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.</p>
Multiplier	<p>Specifies a value for the multiplier used to calculate a receive timeout. The default is 3.</p> <p>Note: If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.</p>

Enable BFD for BGP Peers

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD for Border Gateway Protocol (BGP) peers.



Note

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Peers** tab.
4. Select **Insert**.
5. Select **BfdEnable**.
6. Select **Insert**.

Peers Field Descriptions

Use the data in the following table to use the **Peers** tab.

Name	Description
Instance	Specifies the BGP peer instance.
LocalAddrType	Specifies the local IP address type of the entered BGP peer.
LocalAddr	Specifies the local IP address of the entered BGP peer.
RemoteAddrType	Specifies the remote IP address type of the entered BGP peer.
RemoteAddr	Specifies the remote IP address of the entered BGP peer.
AdminStatus	Specifies the administrative status of the BGP peer.
GroupName	Specifies the peer group name to which the peer belongs (optional).

Name	Description
AccessType Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Configures a BGP neighbor as part of the trusted or unrestricted segment of a Segmented VRF. The default is the trusted segment. The switch adds existing BGP neighbors in the VRF context to the trusted segment automatically. EDM only displays this field if you launch the VRF context view. This field does not apply to the GlobalRouter.
PeerState	Specifies the BGP peer connection state.
RemoteAs	Configures a remote AS number for the peer or peer-group.
Enable	Controls whether the peer connection is enabled or disabled. The default is disabled.
EbgpMultiHop	Enables or disables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable.
RoutePolicyIn	Specifies the policy (by name) that applies to all network routes learned from this peer.
RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing route updates.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
UpdateSourceInterface	Specifies the source IP address to use when the switch sends eBGP packets to this peer or peer group.
ConnectRetryInterval	Specifies the time interval, in seconds, for the connect retry timer. The suggested value for this timer is 120 seconds.
HoldTimeConfigured	Specifies the time interval, in seconds, for the hold time for this BGP speaker with this peer. This value is in an open message sent to this peer by this BGP speaker. To determine the hold time with the peer, the switch compares this value with the HoldTime value in an open message received from the peer. The HoldTime must be at least three seconds. If the value is zero, the hold time does not establish with the peer. The suggested value for this timer is 180 seconds.
KeepAliveConfigured	Specifies the time interval, in seconds, for the KeepAlive timer configured for this BGP speaker with this peer. KeepAliveConfigured determines the keep alive message frequency relative to HoldTimeConfigured; KeepAlive indicates the actual time interval for the keep alive messages. The maximum value for this timer is one-third of HoldTimeConfigured. If KeepAliveConfigured is zero, no periodic keep alive messages are sent to the peer after the peers establish a BGP connection. Configure a value of 60 seconds.
MD5Authentication	Enables and disables MD5 authentication.

Name	Description
AdvertisementInterval	<p>Specifies the time interval, in seconds, that elapses between each transmission of an advertisement from a BGP neighbor. The default value is 30 seconds.</p> <p>The route advertisement interval feature is implemented using the time stamp that indicates when each route is advertised. The time stamp is marked to each route so that the route advertisement interval is compared to the time stamp and BGP is then able to make a decision about whether the route advertisement can be sent or should be delayed when a better route is received. This feature does not work for a withdraw route because the route entry is already removed when the processing route advertisement is sent and the time stamp marked in the route entry cannot be obtained.</p>
DefaultOriginate	When enabled, specifies that the current route originated from the BGP peer. This parameter enables or disables sending the default route information to the specified neighbor or peer. The default value is false.
DefaultOriginateIpv6	When enabled, specifies that the current IPv6 route originated from the BGP peer. This parameter enables or disables sending the default IPv6 route information to the specified neighbor or peer. The default value is false.
Weight	Specifies the peer or peer group weight, or the priority of updates the system can receive from this BGP peer. The default value is 100.
MaxPrefix	Configures a limit on the number of routes accepted from a neighbor. The default value is 12000 routes. A value of 0 means no limit exists.
NextHopSelf	Specifies that the next-hop attribute in an iBGP update is the address of the local router or the router that generates the iBGP update. The default is disable.
RouteReflectorClient	<p>Specifies that this peer is a route reflector client.</p> <p>Note: This parameter only applies to VRF 0.</p>
SoftReconfigurationIn	<p>When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable.</p> <p>Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).</p>

Name	Description
DebugMask	Displays the specified debug information for the BGP peer. The default value is none. <ul style="list-style-type: none"> • None disables all debug messages. • Event enables the display of debug event messages. • State enables display of debug state transition messages. • Update enables display of debug messages related to updates transmission and reception. • Error enables the display of debug error messages. • Trace enables the display of debug trace messages. • Init enables the display of debug initialization messages. • All enables all debug messages. • Packet enables the display of debug packet messages. • Warning enables the display of debug warning messages. • Filter enables the display of debug messages related to filtering.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer. The default value is disable.
Vpnv4Address	Specifies the vpnv4 routes.
IpvpnLiteCap	Enable or disable IP VPN-lite capability on the BGP neighbor peer.
Ipv6Cap	Enable or disable the IPv6 capability on the BGP neighbor peer. The default value is disable.
SooAddress	Specifies the site-of-origin (SoO) address of the BGP peer.
SooAsNumber	Specifies the site-of-origin (SoO) Autonomous System (AS) number of the BGP peer.
SooAssignedNum	Specifies the site-of-origin (SoO) assigned number of the BGP peer.
SooType	Specifies the site-of-origin (SoO) type of the BGP peer.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates in the database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer. The default is disable.
AllowAsIn	Specifies the number of AS-in allowed for the BGP peer.
Ipv6RoutePolicyIn	Specifies the policy (by name) that applies to all network IPv6 routes learned from this peer.

Name	Description
Ipv6RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing IPv6 route updates.
Ipv6MaxPrefix	Configures a limit on the number of IPv6 routes accepted from a neighbor. A value of 0 means no limit exists.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for this BGP peer.

Enable BFD for BGP Peer Groups

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD for Border Gateway Protocol (BGP) peer groups.



Note

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

- In the navigation pane, expand **Configuration > IP**.
- Select **BGP**.
- Select the **Peer Groups** tab.
- Select **Insert**.
- Select **BfdEnable**.
- Select **Insert**.

Peer Groups Field Descriptions

Use the data in the following table to use the **Peer Groups** tab.

Name	Description
Index	Specifies the index of this peer group.
GroupName	Specifies the peer group to which this neighbor belongs (optional).

Name	Description
AccessType Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Configures a BGP peer group as part of the trusted or unrestricted segment of a Segmented VRF. The default is the trusted segment. The switch adds existing BGP neighbors in the VRF context to the trusted segment automatically. EDM only displays this field if you launch the VRF context view. This field does not apply to the GlobalRouter.
Enable	Enables or disables the peer group.
RemoteAs	Configures a remote AS number for the peer-group.
DefaultOriginate	When enabled, the BGP speaker (the local router) sends the default route 0.0.0.0 to a group of neighbors for use as a default route. The default is disabled.
DefaultOriginateIpv6	When enabled, the BGP speaker (the local router) sends the default route to a group of neighbors for use as a default route. The default is disabled.
EbgpMultiHop	When enabled, the switch accepts and attempts BGP connections to external peers that reside on networks that do not directly connect. The default is disabled.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between BGP routing updates. The default value is 30 seconds.
KeepAlive	Specifies the time interval, in seconds, between sent BGP keep alive messages to remote peers. The default value is 60.
HoldTime	Configures the hold time for the group of peers in seconds. Use a value that is three times the value of the KeepAlive time. The default value is 180.
Weight	Assigns an absolute weight to a BGP network. The default value is 100.
MaxPrefix	Limits the number of routes accepted from this group of neighbors. A value of zero indicates no limit. The default value is 12,000 routes.
NextHopSelf	Specifies that the switch must set the NextHop attribute to the local router address before sending updates to remote peers.
RoutePolicyIn	Specifies the route policy that applies to all networks learned from this group of peers.
RoutePolicyOut	Specifies the route policy that applies to all outgoing updates to this group of peers.
RouteReflectorClient	Specifies that this peer group is a route reflector client. Note: This parameter only applies to VRF 0.

Name	Description
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is enable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
MD5Authentication	Enables and disables MD5 authentication. The default is disable.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer group. The default value is disable.
AfUpdateSourceInterfaceType	Specifies the interface type.
AfUpdateSourceInterface	Specifies the IP address used for circuitless IP (CLIP) for this peer group.
Vpnpv4Address	Enables BGP address families for IPv4 (BGP) and Layer 3 VPN (MP-BGP) support. Enable this parameter for VPN/VRF Lite routes.
IppvpnLiteCap	Specifies (when enabled) that IP VPN Lite capability can be enabled or disabled on the BGP neighbor peer. The default is disable.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates in the database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer group. The default is disable.
AllowedAsIn	Specifies the number of AS-in allowed for the BGP peer group.
IPv6Cap	Enable or disable the IPv6 capability on the BGP neighbor peer. The default value is disable.
Ipv6RoutePolicyIn	Specifies the policy (by name) that applies to all network IPv6 routes learned from this peer.
Ipv6RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing IPv6 route updates.
Ipv6MaxPrefix	Configures a limit on the number of IPv6 routes accepted from a neighbor. A value of 0 means no limit exists.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for the BGP peer group.

Enable BFD for BGPv6 Peers

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD for BGPv6 peers.

**Note**

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

**Note**

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see .

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **BGP+**.
3. Select the **Peers** tab.
4. Select **Insert**.
5. Select **BfdEnable**.

Peers Field Descriptions

Use the data in the following table to use the **Peers** tab.

Name	Description
RemoteAddr	Specifies the remote IPv6 address of the entered BGP+ peer.
GroupName	Specifies the peer group name to which the peer belongs (optional).
PeerState	Specifies the BGPv6 peer connection state.
RemoteAs	Configures a remote AS number for the peer or peer-group in the range 0 to 65535.
Enable	Controls whether the peer connection is enabled or disabled. The default is disabled.
EbgpMultiHop	Enables or disables a connection to a BGPv6 peer that is more than one hop away from the local router. The default value is disable.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
UpdateSourceInterface	Specifies the source IP address to use when the switch sends eBGP packets to this peer or peer group.
ConnectRetryInterval	Specifies the time interval, in seconds, for the connect retry timer. The suggested value for this timer is 120 seconds. The range is 1 to 65535.

Name	Description
HoldTimeConfigured	Specifies the time interval, in seconds, for the hold time for this BGP speaker with this peer. This value is in an open message sent to this peer by this BGP speaker. To determine the hold time with the peer, the switch compares this value with the HoldTime value in an open message received from the peer. The HoldTime must be at least three seconds. If the value is zero, the hold time does not establish with the peer. The suggested value for this timer is 180 seconds. The range is 0 to 65535.
KeepAliveConfigured	Specifies the time interval, in seconds, for the KeepAlive timer configured for this BGP speaker with this peer. KeepAliveConfigured determines the keep alive message frequency relative to HoldTimeConfigured; KeepAlive indicates the actual time interval for the keep alive messages. The maximum value for this timer is one-third of HoldTimeConfigured. If KeepAliveConfigured is zero, no periodic keep alive messages are sent to the peer after the peers establish a BGP connection. Configure a value of 60 seconds. The range is 0 to 21845.
MD5Authentication	Enables and disables MD5 authentication.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between each transmission of an advertisement from a BGPv6 neighbor. The default value is 30 seconds and the range is 5 to 120 seconds. The route advertisement interval feature is implemented using the time stamp that indicates when each route is advertised. The time stamp is marked to each route so that the route advertisement interval is compared to the time stamp and BGP is then able to make a decision about whether the route advertisement can be sent or should be delayed when a better route is received. This feature does not work for a withdraw route because the route entry is already removed when the processing route advertisement is sent and the time stamp marked in the route entry cannot be obtained.
DefaultOriginateIpv6	When enabled, specifies that the current IPv6 route originated from the BGP peer. This parameter enables or disables sending the default IPv6 route information to the specified neighbor or peer. The default value is false.
Weight	Specifies the peer or peer group weight, or the priority of updates the system can receive from this BGP peer. The default value is 100 and the range is 0 to 65535.
NextHopSelf	Specifies that the next-hop attribute in an iBGP update is the address of the local router or the router that generates the iBGP update. The default is disable.

Name	Description
RouteReflectorClient	Specifies that this peer is a route reflector client. Note: This parameter only applies to VRF 0.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
DebugMask	Displays the specified debug information for the BGP peer. The default value is none. <ul style="list-style-type: none"> • None disables all debug messages. • Event enables the display of debug event messages. • State enables display of debug state transition messages. • Update enables display of debug messages related to updates transmission and reception. • Error enables the display of debug error messages. • Trace enables the display of debug trace messages. • Init enables the display of debug initialization messages. • All enables all debug messages. • Packet enables the display of debug packet messages. • Warning enables the display of debug warning messages. • Filter enables the display of debug messages related to filtering.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer. The default value is disable.
Vpnv4Address	Enables BGP address families for IPv4 (BGP) and Layer 3 VPN (MP-BGP) support. Enable this parameter for VPN/VRF Lite routes.
IvpnLiteCap	Enable or disable IP VPN-lite capability on the BGP neighbor peer.
Ipv6Cap	Enable or disable the IPv6 capability on the BGP neighbor peer. The default value is disable.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates in the database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride Note:	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer. The default is disable.

Name	Description
AllowAsIn Note:	Specifies the number of AS-in allowed for the BGP peer. The range is 1–10.
Ipv6RoutePolicyIn	Specifies the policy (by name) that applies to all network IPv6 routes learned from this peer.
Ipv6RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing IPv6 route updates.
Ipv6MaxPrefix	Configures a limit on the number of IPv6 routes accepted from a neighbor. A value of 0 means no limit exists.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for this peer.

Enable BFD for OSPF on an IPv4 Port Interface

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD for the OSPF protocol on an IPv4 port interface.



Note

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **OSPF** tab.
5. Select **BfdEnable**.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
Enable	Enables or disables OSPF routing on the specified port. The default is false.
HelloInterval	Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet, and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
DesigRtrPriority	Specifies the priority of this port in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for the type of service (TOS) on this port. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. <ul style="list-style-type: none"> FFFF—No route exists for this TOS. IPCP links—Defaults to 0. 0—Use the interface speed as the metric value when the state of the interface is up.

Name	Description
AuthType	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha1—Specifies secure hash algorithm (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. You can only access and enable the SHA-1 authentication type after you enable enhanced secure mode. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
AuthKey	Specifies the key (up to 8 characters) when you specify simple password authentication in the port AuthType variable.
AreaId	<p>Specifies the OSPF area name in dotted-decimal format.</p> <p>The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).</p>
AdvertiseWhenDown	<p>Advertises the network on this port as up, even if the port is down. The default is false.</p> <p>After you configure a port with no link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even if the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.</p>
IfType	<p>Specifies the type of OSPF interface (broadcast, NBMA, passive, or p2p).</p> <p>Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.</p>
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must have the same poll interval.

Name	Description
IfMtuIgnore	Specifies whether the interface ignores the global maximum transmission unit (MTU) configuration. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
BfdEnable	Enable Bidirectional Forwarding Detection (BFD) for OSPF.

Enable BFD for OSPF on an IPv4 VLAN Interface

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable OSPF BFD on an IPv4 VLAN interface.



Note

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN on which you want to enable BFD for OSPF.
5. Select **IP**.
6. Select **OSPF**.
7. Select **BfdEnable**.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
Enable	Enables or disables OSPF routing on the specified VLAN. The default is false.
HelloInterval	<p>Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds.</p> <p>After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.</p>

Name	Description
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
DesigRtrPriority	Specifies the priority of this VLAN in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for this TOS on this VLAN. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. <ul style="list-style-type: none"> FFFF—No route exists for this TOS. IPCP links—Defaults to 0. 0—Use the interface speed as the metric value when the state of the interface is up.
AuthType	Specifies the type of authentication required for the interface. <ul style="list-style-type: none"> none—Specifies that no authentication required. simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. sha1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. You can only access and enable the SHA-1 authentication type after you enable enhanced secure mode. sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
AuthKey	Specifies the key (up to eight characters) when you specify simple password authentication in the VLAN AuthType variable.

Name	Description
AreaId	Specifies the OSPF area name in dotted-decimal format. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdvertiseWhenDown	Advertises the network even if the port is down. If true, OSPF advertises the network on this VLAN as up, even if the port is down. The default is false. After you configure a port without a link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown.
IfType	Specifies the type of OSPF interface (broadcast, NBMA, passive, or p2p). Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must use the same poll interval.
IfMtuIgnore	Specifies whether the VLAN ignores the MTU configuration. To allow the switch to accept OSPF DD packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for OSPF.

Enable BFD for OSPF on an IPv6 Port Interface

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable BFD for the OSPF protocol on an IPv6 port interface.



Note

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select the **IPv6 OSPF Interface** tab.
5. Select **Insert**.
6. Select **BfdEnable**.

IPv6 OSPFv3 Interface *Field Descriptions*

Use the data in the following table to use the **IPv6 OSPFv3 Interface** tab.

Name	Description
Index	Specifies the interface index for the IPv6 interface on which OSPFv3 is configured.
Areaid	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Type Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none"> • broadcast • NBMA • point-to-point • point-to-multipoint • passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this parameter to the same value for all routers attached to a common network. The default is 10.

Name	Description
RtrDeadInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this parameter to the same value for all routers attached to a common network. The default is 40.
PollInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies whether Link LSA suppression is enabled.

Enable BFD for OSPF on an IPv6 VLAN Interface

About This Task

BFD provides a failure-detection mechanism between two systems. Use the following procedure to enable OSPF BFD on an IPv6 VLAN interface.



Note

Enabling BFD at the application level does not establish a BFD session. To establish a BFD session, you must enable BFD globally and at the interface level.

Procedure

1. In the navigation pane, expand the **Configuration** > **VLAN** folders.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN on which you want to enable BFD for OSPF.
5. Select **IPv6**.
6. Select **IPv6 OSPF Interface**.
7. Select **Insert**.
8. Select **BfdEnable**.

IPv6 OSPF Interface Field Descriptions

Use the data in the following table to use the **IPv6 OSPF Interface** tab.

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.
Areaid	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Type	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none"> • broadcast • NBMA • point-to-point • point-to-multipoint • passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.

Name	Description
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this parameter to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this parameter to the same value for all routers attached to a common network. The default is 40.
PollInterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.

Name	Description
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for OSPF.

Configure BFD on an IPv4 Static Route

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BFD**.
3. Select the **Static Route** tab.
4. Select **Insert**.
5. In **NextHop**, type the IPv4 address of the next hop of the BFD session.
6. (Optional) In **VrfId**, type the ID of the VRF associated with the BFD session.
7. Select **Insert**.

Static Route Field Descriptions

Use the data in the following table to use the **Static Route** tab.

Name	Description
NextHop	Specifies the IPv4 address of the next hop of the BFD session.
VrfId	Specifies the ID of the VRF associated with the BFD session.
VrfName	Specifies the name of the VRF associated with the BFD session.

Configure BFD on an IPv6 Static Route

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 BFD**.
3. Select **Insert**.
4. In **Interface**, select either **Port** or **Vlan**, and then select an interface.

5. In **NextHop**, type the IPv6 address of the next hop of the BFD session.
6. (Optional) In **Vrflid**, type the ID of the VRF associated with the BFD session.
7. Select **Insert**.

Static Route *Field Descriptions*

Use the data in the following table to use the **Static Route** tab.

Name	Description
Interface	Specifies either a port or VLAN interface.
NextHop	Specifies the IPv4 address of the next hop of the BFD session.
Vrflid	Specifies the ID of the VRF associated with the BFD session.
VrfName	Specifies the name of the VRF associated with the BFD session.

Display BFD Performance Counters

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **BFD**.
3. Select the **Performance counters** tab.

Performance counters *Field Descriptions*

Use the data in the following table to use the **Performance counters** tab.

Name	Description
PktIn	Specifies the total number of BFD messages received for this BFD session.
PktOut	Specifies the total number of BFD messages sent for this BFD session.



BGP

- [BGP Fundamentals](#) on page 536
- [BGP Configuration Using CLI](#) on page 573
- [BGP Verification Using CLI](#) on page 602
- [BGP Configuration Using EDM](#) on page 618
- [BGP Configuration Examples](#) on page 654

The following sections provide conceptual information and procedures that you can use to configure Border Gateway Protocol (BGP) services.

Unless otherwise noted, examples and network illustrations in these sections apply to all supported platforms.

BGP Fundamentals

Table 56: Border Gateway Protocol product support

Feature	Product	Release introduced
Border Gateway Protocol for IPv4 (BGPv4)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none">• 5320-16P-2MXT-2X• 5320-24T-4X-XT• 5320-24T-24S-4XE-XT• 5320-48P-8XE• 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 56: Border Gateway Protocol product support (continued)

Feature	Product	Release introduced
BGP+ (BGPv4 for IPv6)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Not supported on 5320-16P-2MXT-2X
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
BGPv6	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
External BGP (eBGP)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 56: Border Gateway Protocol product support (continued)

Feature	Product	Release introduced
Internal BPG (iBGP)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Route metric for BGP route redistribution	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 56: Border Gateway Protocol product support (continued)

Feature	Product	Release introduced
iBGP over user-created VRFs	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.1

Border Gateway Protocol (BGP) is an inter-domain routing protocol that provides loop-free routing between autonomous systems (AS) or within an AS. This section describes the major BGP features.

Autonomous Systems

An Autonomous system (AS) is a group of routers and hosts run by a single technical administrator that has a single, clearly defined routing policy. Each AS uses a unique AS number assigned by the appropriate Internet Registry entity. LANs and WANs that interconnect by IP routers form a group of networks called an internetwork. For administrative purposes, internetworks divide into boundaries known as autonomous systems.

The following figure shows a sample internetwork segmented into three autonomous systems.

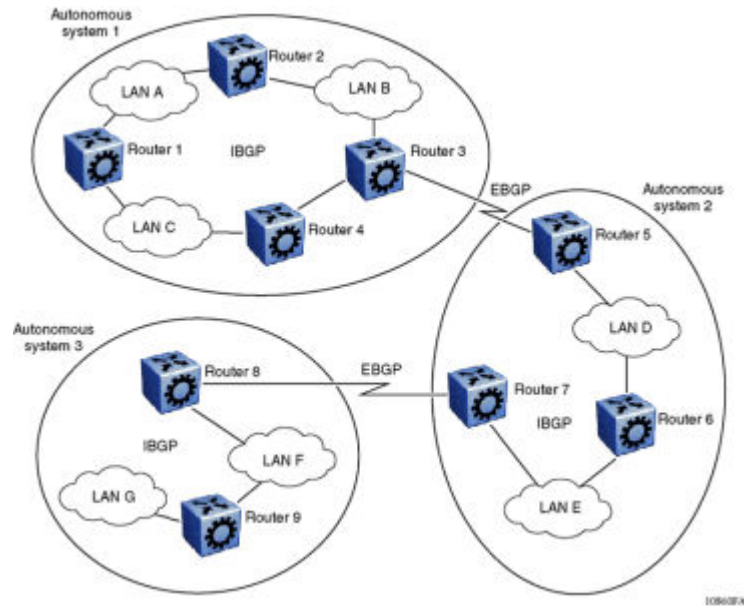


Figure 19: Internetwork segmented into three autonomous systems

BGP exchanges information between autonomous systems as well as between routers within the same AS. As shown in the preceding figure, routers that are members of the same AS and exchange BGP updates run internal BGP (iBGP), and routers that are members of different autonomous systems and exchange BGP updates run external BGP (eBGP).

Internal and External BGP Routing

The switch supports both iBGP intra-AS routing and eBGP external-AS routing. With iBGP, each router within an AS runs an interior gateway protocol (IGP), such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). The iBGP information, along with the IGP route to the originating BGP border router, determines the next hop to use to exchange information with an external AS. Each router uses iBGP exclusively to determine reachability to external autonomous systems. After a router receives an iBGP update destined for an external AS, it passes the update to IP for inclusion in the routing table only if a viable IGP route to the correct border gateway is available.

BGP speakers in different autonomous systems use eBGP to communicate routing information.

BGP Speaker

BGP routers employ an entity within the router, referred to as a BGP speaker, which transmits and receives BGP messages and acts upon them. BGP speakers establish a peer-to-peer session with other BGP speakers to communicate.

All BGP speakers within an AS must be fully meshed. The following figure shows a BGP network with fully-meshed BGP speakers.




Figure 20: BGP networks

Transit AS

An AS with more than one BGP speaker can use iBGP to provide a transit service for networks located outside the AS. An AS that provides this service is a transit AS. As shown in the preceding figure, [BGP networks](#), AS 40 is the transit AS. AS 40 provides information about the internal networks, as well as transit networks, to the remaining autonomous systems. The iBGP connections between routers D, E, and F provide consistent routing information to the autonomous systems.

Stub and Multihomed Autonomous Systems

As shown in the preceding figure, [BGP networks](#), an AS can include one or more BGP speakers that establish peer-to-peer sessions with BGP speakers in other autonomous systems to provide external route information for the networks within the AS.

A stub AS has a single BGP speaker that establishes a peer-to-peer session with one external BGP speaker. In this case, the BGP speaker provides external route information only for the networks within its own AS.

A multihomed AS has multiple BGP speakers.

Peers

BGP uses Transmission Control Protocol (TCP) as a transport protocol. When two routers open a TCP connection to each other for the purpose of exchanging routing information, they form a peer-to-peer relationship. In the preceding figure, [BGP](#)

[networks](#), Routers A and D are BGP peers, as are Routers B and E, C and E, F and G, and Routers D, E, and F.

Although Routers A and D run eBGP, Routers D, E, and F within AS 40 run iBGP. The eBGP peers directly connect to each other, while the iBGP peers do not. As long as an IGP operates and allows two neighbors to logically communicate, the iBGP peers do not require a direct connection.

**Note**

You cannot create the same iBGP peers on two different VRFs, or the same eBGP peers on two different chassis. Only one local autonomous system (AS) can exist for each chassis or VRF.

Because all BGP speakers within an AS must be fully meshed logically, the iBGP mesh can grow to large proportions and become difficult to manage. You can reduce the number of peers within an AS by creating confederations and route reflectors.

BGP peers exchange complete routing information only after the peers establish a connection. Thereafter, BGP peers exchange routing updates. An update message consists of a network number, a list of autonomous systems that the routing information passed through (the AS path), and other path attributes that describe the route to a set of destination networks. When multiple paths exist, BGP compares the path attributes to choose the preferred path. Even if you disable BGP, the system logs all BGP peer connection requests. For more information about update messages, see [BGP Updates](#) on page 557.

Supernet Advertisements

BGP has no concept of address classes. Each network listed in the network layer reachability information (NLRI) portion of an update message contains a prefix length field, which describes the length of the mask associated with the network. The prefix length field allows for both supernet and subnet advertisement. The supernet advertisement is what makes classless interdomain routing (CIDR) possible (see [CIDR and Aggregate Addresses](#) on page 545).

Bandwidth and Maintenance Reduction

BGP provides two features that reduce the high bandwidth and maintenance costs associated with a large full-mesh topology:

- confederations
- route reflectors

**Note**

Confederations and route reflectors are not supported on iBGP for non-default VRFs.

For information on confederations and route reflectors, see [Routing Information Consolidation](#) on page 545.

BGP 4 Byte AS Support

Each Autonomous System (AS) must have its own unique number. Because the 2-byte AS numbering scheme is unable to meet the increasing demand, the switch supports 4-byte AS numbers. This feature is enabled by supporting RFC 4893, BGP Support for 4-octet AS Number Space.

The switch supports the following three types of peer relationships as a result of 4 byte AS support:

- Old peer to old peer
- Old peer to new peer
- New peer to new peer

An old peer is the one that supports 2-byte AS numbers only and new peer is the one that supports both 2-byte AS numbers and 4-byte AS numbers.

RFC4893 supports two new path attributes:

- AS4_PATH contains the AS path encoded with a 4-octet AS number.
- AS4-AGGR is a new aggregator attribute that carries a 4-octet AS number.




Figure 21: 2-byte and 4-byte Mixed Environment

The preceding figure shows an example of how the switch uses the AS4_PATH attribute in a mixed environment. The figure illustrates how a 2-byte BGP speaker interoperates with a 4-byte BGP speaker.

Router B is a 2-byte BGP speaker. Router A substitutes AS_PATH with the AS_TRANS, a 2-octet AS number defined by RFC4893 for backward compatibility, and encodes the 4-byte AS into AS4_PATH in BGP updates it sends to router B.

Router B does not understand the AS4_PATH but does preserve the information and sends it to router C.

Router C is a 4-byte BGP speaker. Router C merges the information received in AS_PATH and AS4_PATH, and encodes the 4-byte AS when it sends the AS_PATH information to router D.

Old Peer to Old Peer

When the peer relationship between an old peer and another old peer is established, 4 byte AS numbers contained in the AS4_PATH and AS4_AGGREGATOR are transited to other peers.



Important

Do not assign 23456 as an AS number. The Internet Assigned Numbers Authority (IANA) reserved this number for the AS_TRANS attribute and BGP uses it to facilitate communication between peer modes. AS_TRANS uses a 2-byte AS format to represent a 4-byte AS number. The switch interprets the AS_TRANS attribute and propagates it to other peers.

New Peer to New Peer

The new BGP speaker establishes its 4 byte AS support through BGP capability advertisement. A BGP speaker that announces such capability and receives it from its peer, uses 4 byte AS numbers in AS_PATH and AGGREGATOR attributes and assumes these attributes received from its peer are encoded in 4 byte AS numbers.

The new BGP attributes AS4_PATH and AS4_AGGREGATOR received from the new BGP speaker between the new BGP peers in the update message is discarded.

Old Peer to New Peer

An old BGP speaker and a new BGP speaker can form peering relationship only if the new BGP speaker is assigned a 2 byte AS number. This 2 byte number can be any global unique AS number or AS_TRANS.

New BGP speaker sends AS path information to the old BGP speaker in AS_PATH attribute as well as AS4_PATH attribute. If the entire AS_PATH consists of only 2 byte AS numbers then the new BGP speaker does not send AS4_PATH information.

The 4-byte AS number feature does not in any way restrict the use or change the way you configure 2-byte AS numbers. You can also configure 2-byte AS or 4 byte AS numbers in AS path lists, community lists, and route policies.

BGP 4-byte AS Number Notation

BGP 4-byte AS numbers are represented in two ways: AS Plain and AS dot. The default form of representing the AS numbers is AS Plain while you have an option to configure AS dot. AS Plain form of representation is preferred over AS dot representation as a large amount of network providers find the AS dot notation incompatible with the regular expressions used by them. In case of any issues, troubleshooting and analyzing also gets difficult with AS dot notation.

BGP AS Number Format – AS Plain

Table 57: Default Asplain 4-Byte Autonomous System Number Format

Format	Configuration Format	Show Command Output and Regular Expression Match Format
asplain	2-byte: 1 to 65535 4-byte: 65536 to 4294967295	2-byte: 1 to 65535 4-byte: 65536 to 4294967295
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 65536 to 4294967295

BGP AS Number Format - ASdot

Table 58: Asdot 4-Byte Autonomous System Number Format

Format	Configuration Format	Show command output and Regular Expression Match Format
asplain	2-byte: 1 to 65535 4-byte: 65536 to 4294967295	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535

For more information on configuring 4 byte AS numbers, see [Configure 4-byte AS numbers](#) on page 579.

Routing Information Consolidation

Use the information in this section to understand how to reduce the size of routing tables.

CIDR and Aggregate Addresses

Classless interdomain routing (CIDR) is an addressing scheme (also known as supernetting) that eliminates the concept of classifying networks into class types. Earlier addressing schemes identified five classes of networks: Class A, Class B, Class C, Class D, and Class E. This document does not discuss Classes D (used for multicast) and E (reserved and currently not used).

Network 195.215.0.0, an illegal Class C network number, becomes a legal supernet when represented in CIDR notation as 195.215.0.0/16. The /16 is the prefix length and expresses the explicit mask that CIDR requires. In this case, the addition of the prefix /16 indicates that the subnet mask consists of 16 bits (counting from the left).

Using this method, supernet 195.215.0.0/16 represents 195.215.0.0 255.255.0.0. The following table shows the conversion of prefix length to subnet mask.

Table 59: CIDR Conversion

Prefix	Dotted-decimal	Binary	Network class
/1	128.0.0.0	1000 0000 0000 0000 0000 0000 0000 0000	128 Class A
/2	192.0.0.0	1100 0000 0000 0000 0000 0000 0000 0000	64 Class A
/3	224.0.0.0	1110 0000 0000 0000 0000 0000 0000 0000	32 Class A
/4	240.0.0.0	1111 0000 0000 0000 0000 0000 0000 0000	16 Class A
/5	248.0.0.0	1111 1000 0000 0000 0000 0000 0000 0000	8 Class A
/6	252.0.0.0	1111 1100 0000 0000 0000 0000 0000 0000	4 Class A
/7	254.0.0.0	1111 1110 0000 0000 0000 0000 0000 0000	2 Class A
/8	255.0.0.0	1111 1111 0000 0000 0000 0000 0000 0000	1 Class A or 256 Class B
/9	255.128.0.0	1111 1111 1000 0000 0000 0000 0000 0000	128 Class B
/10	255.192.0.0	1111 1111 1100 0000 0000 0000 0000 0000	64 Class B
/11	255.224.0.0	1111 1111 1110 0000 0000 0000 0000 0000	32 Class B
/12	255.240.0.0	1111 1111 1111 0000 0000 0000 0000 0000	16 Class B
/13	255.248.0.0	1111 1111 1111 1000 0000 0000 0000 0000	8 Class B
/14	255.252.0.0	1111 1111 1111 1100 0000 0000 0000 0000	4 Class B
/15	255.254.0.0	1111 1111 1111 1110 0000 0000 0000 0000	2 Class B
/16	255.255.0.0	1111 1111 1111 1111 0000 0000 0000 0000	1 Class B or 256 Class C
/17	255.255.128.0	1111 1111 1111 1111 1000 0000 0000 0000	128 Class C
/18	255.255.192.0	1111 1111 1111 1111 1100 0000 0000 0000	64 Class C
/19	255.255.224.0	1111 1111 1111 1111 1110 0000 0000 0000	32 Class C
/20	255.255.240.0	1111 1111 1111 1111 1111 0000 0000 0000	16 Class C
/21	255.255.248.0	1111 1111 1111 1111 1111 1000 0000 0000	8 Class C
/22	255.255.252.0	1111 1111 1111 1111 1111 1100 0000 0000	4 Class C
/23	255.255.254.0	1111 1111 1111 1111 1111 1110 0000 0000	2 Class C
/24	255.255.255.0	1111 1111 1111 1111 1111 1111 0000 0000	1 Class C

Use CIDR to assign network prefixes of arbitrary lengths, as opposed to the obsolete class system, which assigned prefixes as even multiples of an octet.




Figure 23: First aggregate and prefix length

- The 21 bits that match the common bits form the prefix length. The prefix length is the number of binary bits that form the explicit mask (in dotted-decimal notation) for this IP prefix.
- The remaining aggregate is formed from 192.32.8.0 to the end range, 192.32.9.255.

As shown in [Figure 23](#), the resulting first aggregate 192.32.0.0/21 represents all of the IP prefixes from 192.32.0.0 to 192.32.7.255.

The following figure shows the results after forming the remaining aggregate from 192.32.9.0 to the end range, 192.32.9.255.

The resulting aggregate 192.32.8.0/23 represents all of the IP prefixes from 192.32.8.0 to 192.32.9.255.

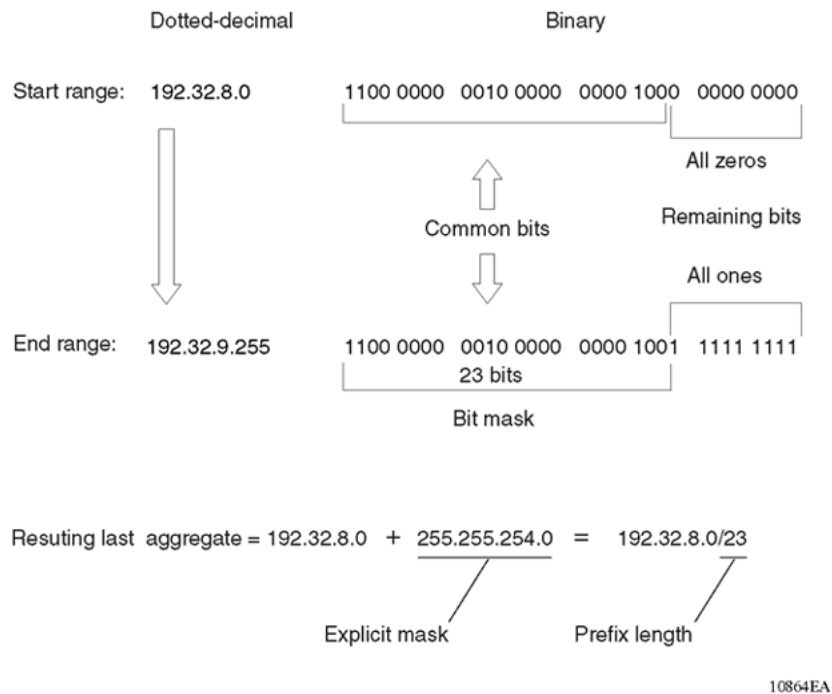


Figure 24: Last aggregate and prefix length

The final result of calculating the supernet address that ranges from 192.32.0.0 to 192.32.9.255 is as follows:

192.32.0.0 (with mask) 255.255.248.0 = 192.32.0.0/21

192.32.8.0 (with mask) 255.255.254.0 = 192.32.8.0/23

Aggregate Routes

Eliminating the idea of network classes provides an easy method to aggregate routes. Rather than advertise a separate route for each destination network in a supernet, BGP uses a supernet address to advertise a single route (called an aggregate route) that represents all the destinations. CIDR also reduces the size of the routing tables used to store advertised IP routes.

The following figure shows an example of route aggregation using CIDR. In this example, a single supernet address 195.215.0.0/16 advertises 256 separate Class C network numbers 195.215.0.0 through 195.215.255.0.




Figure 25: Aggregating routes with CIDR

Confederations

A BGP router configured for iBGP establishes a peer-to-peer session with every other iBGP speaker in the AS. In an AS with a large number of iBGP speakers, this full-mesh topology can result in high bandwidth and maintenance costs.



Note

Confederations are not supported on iBGP for non-default VRFs.

As shown in the following example, a full-mesh topology for an AS with 50 iBGP speakers requires 1225 internal peer-to-peer connections:

Example:

$$n \times (n-1)/2 = n \text{ iBGP sessions}$$

where:

$$50 \times (50-1)/2 = 1225 \text{ number of unique iBGP sessions}$$

You can reduce the high bandwidth and maintenance costs associated with a large full-mesh topology by dividing the AS into multiple smaller autonomous systems (sub-autonomous systems), and then group them into a single confederation (see the following figure).

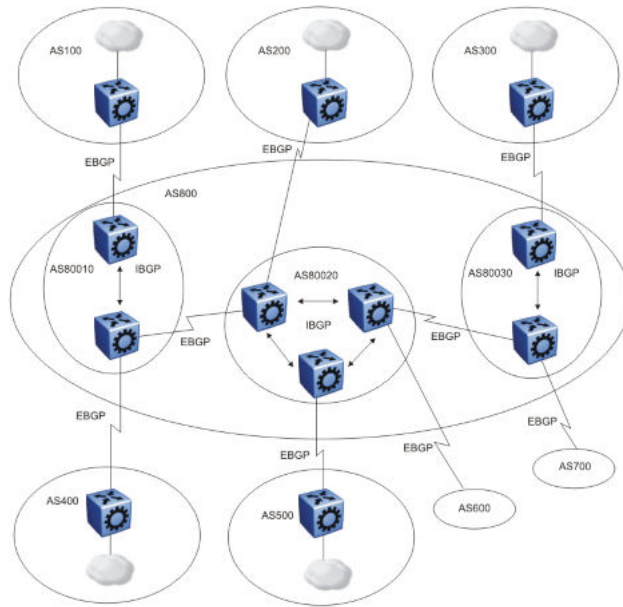


Figure 26: Confederations

As shown in the preceding figure, each sub-AS is fully meshed within itself and has eBGP sessions with other sub-autonomous systems in the same confederation.

Although the peers in different autonomous systems have eBGP sessions with the various sub-AS peers, they preserve the next-hop, Multi-Exit Discriminator (MED), and local preference information and exchange routing updates as if they were iBGP peers. All of the autonomous systems retain a single interior gateway protocol (IGP). When the confederation uses its own confederation identifier, the system displays the group of sub-autonomous systems as a single AS (with the confederation identifier as the AS number).

Route Reflectors

Another way to reduce the iBGP mesh inherent in an AS with a large number of iBGP speakers is to configure a route reflector. Using this method, when an iBGP speaker needs to communicate with other BGP speakers in the AS, the speaker establishes a single peer-to-peer route reflector client session with the iBGP route reflector.



Note

Route reflectors are not supported on iBGP for non-default VRFs.

In an AS, more than one route reflector cluster can exist and more than one route reflector in a cluster. When more than one reflector exists in a cluster, take care to prevent route loops.

The following figure shows a simple iBGP configuration with three iBGP speakers (routers A, B, and C). Without route reflectors, after Router A receives an advertised route from an external neighbor, it must advertise the route to Routers B and C.

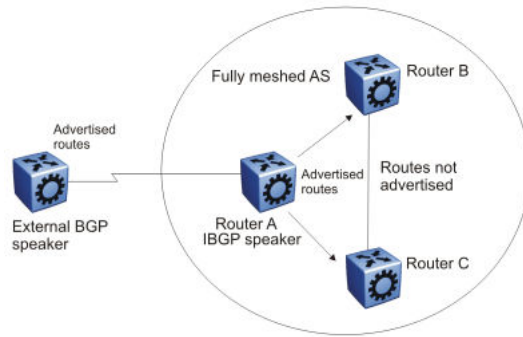


Figure 27: Fully meshed AS with iBGP speakers

Routers B and C do not readvertise the iBGP learned routes to other iBGP speakers. BGP does not allow routers to pass routes learned from internal neighbors on to other internal neighbors, which avoids routing information loops.

As shown in the following figure, when you configure an internal BGP peer (Router B) as a route reflector, all of the iBGP speakers do not need to be fully meshed. In this case, the assigned route reflector passes iBGP learned routes to a set of iBGP neighbors.




Figure 28: AS with route reflector

After Router B, the route reflector, receives routes from Router A (the iBGP speaker), it advertises them to router C. Conversely, after the route reflector receives routes from internal peers, it advertises those routes to Router A. Routers A and C do not need an iBGP session.

Route reflectors separate internal peers into two groups: client peers and nonclient peers. The route reflector and its clients form a cluster. The client peers in the cluster do not need to be fully meshed, and do not communicate with iBGP speakers outside their cluster. Nonclient peers must be fully meshed with each other.

The following figure shows a cluster, where Router A is the route reflector in a cluster with client routers B, C, and D. Routers E, F, and G are fully meshed, nonclient routers.

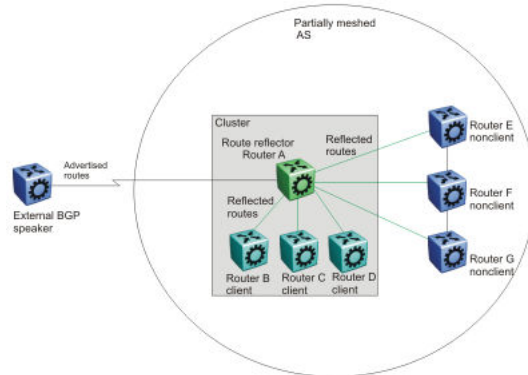


Figure 29: Route reflector with client and nonclient peers

BGP Communities

You can group destinations into communities to simplify policy administration. A community is a group of destinations that share a common administrative property.

Use community control routing policies with respect to destinations. Create communities when you have more than one destination and want to share a common attribute.

The following list identifies specific community types:

- Internet—advertise this route to the Internet community
- no advertise—do not advertise to BGP peers including iBGP peers

You can use a community to control which routing information to accept, prefer, or distribute to other BGP neighbors. If you specify the append option in the route policy, the router adds the specified community value to the existing value of the community attribute. Otherwise, the specified community value replaces a previous community value.

BGP Path Attributes

You can create policies that control routes, work with default routing, control specific and aggregated routes, and manipulate BGP path attributes.

Four categories of BGP path attributes exist:

- Well-known mandatory attributes must be in every BGP update message.
- Well-known discretionary attributes can be in a BGP update message.
- Optional transitive attributes are accepted and passed to other BGP peers.
- Optional non-transitive attributes can be either accepted or ignored, but must not pass along to other BGP peers.

Border routers that utilize built-in algorithms or manually configured policies to select paths use path attributes. BGP uses the following path attributes to control the path a BGP router chooses:

- origin (well-known mandatory)
- AS_path (well-known mandatory)
- next hop (well-known mandatory)
- MED attribute (optional non-transitive)
- local preference (well-known discretionary)
- atomic aggregate (well-known discretionary)
- aggregator (optional transitive)
- community (optional transitive)

For more information about path attributes in BGP updates, see [Path Attributes](#) on page 558.

BGP Route Selection

A BGP router determines the best path to a destination network. This path is then eligible for use in the IP forwarding table and the router also advertises the path to its eBGP peers. To choose the best of multiple BGP routes to a destination, the router executes a best path algorithm.

The algorithm chooses a route in the following order:

- highest weight

Weight is a locally significant parameter associated with each BGP peer. You can use the weight to influence which peer paths the router uses.

- highest local preference

The local preference has global significance within an AS. You can manipulate the preference using route policies to influence path selection.

- prefer locally originated paths

The router prefers a path locally originated using the network, redistribution, or aggregate command over a path learned through a BGP update. The router prefers local paths sourced by network or redistribute commands over local aggregates sourced by the aggregate address command.

- shortest AS path

The AS path parameter specifies the autonomous systems that the network prefix traversed. The AS path commonly determines the best path. For example, a router can choose a path based on whether the network passed through a specific AS. You can configure a route policy to match the AS, and then modify the local preference. Also, you can pad the AS path before the AS advertises it to a peer AS, so that downstream routers are less likely to prefer the advertised network path.

The AS_CONFED_SEQUENCE length will also be considered while picking the best path inside the confederation.

- lowest origin type

The order of preference is IGP, EGP, INC (incomplete).

- lowest MED

The MED parameter influences the preferred path from a remote AS to the advertising AS. This parameter applies when there are multiple exit points from the remote AS to the advertising AS. A lower MED value indicates a stronger path preference than a higher MED value. By default, the MED attribute is ignored as specified by the BGP global parameter Always Compare MED except when the routes come from the same AS. This parameter must be enabled for MEDs to be compared (and for this step of the best path algorithm to execute).

The router compares MEDs regardless of what the first (neighboring) AS specified in the AS_PATH. Deterministic MED, when enabled, means that the first AS of the multiple paths must be the same. Paths received with no MED are assigned a MED of 0, unless the global BGP parameter Missing Is Worst is enabled. If so, received paths are assigned a MED of 4 294 967 294. Missing is Worst is enabled by default. The "no-med-path-is-worst" flag has an impact only when the "First AS" or the "Most Left AS" is the same for multiple routes received. The router changes paths received with a MED of 4 294 967 295 to 4 294 967 294 before insertion into the BGP table.




Note

You cannot enable or disable the MED selection process. BGP aggregation does not occur when routes have different MEDs or next hops.

When MED value is set in route-map configuration, the configured MED value is not applicable if it is already set in the associated Path Attribute.

1. When router A sets MED value of 100 by route-map, it will send Path Attribute with MED=100 to EBGp peer B.
2. Router B sends Path Attribute with MED=100 to IBGP peer C.
3. If the route-map is configured with "set MED 200", then router C does not apply MED=200 to the Path Attribute as it is already set to 100 when it is received from router B.
4. Router D will get Path Attribute with MED=100 so that router C does not influence router D when it selects the best route.



Example: If Prefix: X is set as MED=100 from router A, it will be received at B with MED=100, and will carry same MED=100 value to router C, as it is an IBGP peer. Router C will not propagate MED=100 value to D as MED is a non-transitive attribute, so MED can travel maximum of 1 AS.

- lowest IGP metric to the BGP next-hop

If multiple paths exist whose BGP next-hop is reachable through an IGP, the path with the lowest IGP metric to the BGP next-hop is chosen.

- prefer external paths (learned by eBGP) over internal paths (iBGP)

The system prefers external paths over internal paths.

- if Equal Cost Multipath (ECMP) is enabled, insert up to four paths in the routing table

If you enable ECMP, multiple BGP learned routes that use the same metric to different IP next-hops are installed in the IP forwarding table for traffic load-balancing purposes.

- lowest router ID

The lowest router ID, or Circuitless IP (CLIP) address, is preferred.

BGP and Dampened Routes

The switch supports route dampening (route suppression). When you use route dampening, a route accumulates penalties each time the route fails. After the accumulated penalties exceed a threshold, the router no longer advertises the route. The router enters the suppressed routes into the routing table only after the accumulated penalty falls below the reuse threshold.

Route flap dampening suppresses the advertisement of the unstable route until the route becomes stable. For information about how to enable flap-dampening, see [Configure BGP](#) on page 574. For information about viewing flap dampening configurations, see [View the Global Flap-Dampening Configuration](#) on page 606.

Dampening applies only to routes that are learned through an eBGP. Route flap dampening prevents routing loops and protects iBGP peers from having higher penalties for routes external to the AS.

The following paragraph describes the algorithm that controls route flaps.

After the route flaps the first time

- the router creates a route history entry
- a timer starts (180 seconds)

If the route does not flap again, the router uses this timer to delete the history entry after the 180 seconds expires.

After the route flaps a second time

- The penalty is recalculated based on the decay function.

If the penalty is greater than the cut-off value (1536), the route is suppressed and the reuse time is calculated based on the reuse time function.

- The reuse timer starts.

After the reuse time expires, the suppressed route is announced again (the reuse time is recalculated if the route flaps again). The penalty decays slower for withdrawn routes than for update routes. The route history entry is kept longer if the route is withdrawn. For update history, the delete time is 90 seconds and the withdrawn history delete time is 180 seconds.

BGP Updates

BGP uses update messages to communicate information between two BGP speakers. The update message can advertise a single feasible route to a peer, or withdraw multiple unfeasible routes from service.

The following figure shows the format of an update message.




Figure 30: Update message format

This section describes how BGP uses the update message fields to communicate information between BGP speakers.

Withdrawn Routes Length

The withdrawn routes length parameter (referred to in RFC1771 as the Unfeasible Routes Length field) indicates the total length of the withdrawn routes field in octets. The withdrawn routes length field calculates the length of the NLRI field. For example, a value of 0 indicates that no routes are withdrawn from service, and that the withdrawn routes field is not present in this update message.

Withdrawn Routes

The withdrawn routes parameter is a variable-length parameter that contains a list of IP prefixes for routes that are withdrawn from service. The following figure shows the format of an IP prefix.




Figure 31: IP prefix format

The length indicates the number of bits in the prefix (also called the network mask).

For example, 192.0.2.0/24 is equivalent to 192.0.2.0 255.255.255.0 (the /24 indicates the number of bits in the length parameter to represent the network mask 255.255.255.0).

The prefix parameter contains the IP address prefix itself, followed by enough trailing bits to make the length of the whole field an integer multiple of 8 bits (1 octet).

Total Path Attributes Length

The total path attributes length parameter indicates the total length of the path attributes parameter in octets.

The total path attributes length calculates the length of the NLRI parameter. For example, a value of 0 indicates that no NLRI field is present in this update message.

Path Attributes

The path attributes parameter is a variable-length sequence of path attributes that exists in every BGP update. The path attributes contain BGP attributes associated with the prefixes in the NLRI parameter.

For example, the attribute values allow you to specify the prefixes that the BGP session can exchange, or which of the multiple paths of a specified prefix to use.

The attributes carry the following information about the associated prefixes:

- the path origin
- the AS paths through which the prefix is advertised
- the metrics that display degrees of preference for this prefix

The following figure shows the encoding used with the path attribute parameter.




Figure 32: Path attribute encoding

Attribute Type

As shown in the following figure, the attribute type is a two-octet field that comprises two sub-fields: attribute flags and attribute type code.



Figure 33: Attribute type fields

The attribute flags parameter is a bit string that contains four binary values that describe the attribute, and four unused bits. The following list provides bit descriptions (from the high-order bit to the low-order bit):

- The high-order bit (bit 0) is the optional bit. When this bit is set (the value is 1), the attribute is optional. When this bit is clear (the value is 0), the attribute is well-known. Well-known attributes must be recognized by all BGP implementations and, when appropriate, passed on to BGP peers. Optional attributes are not required in all BGP implementations.
- The second high-order bit (bit 1) is the transitive bit. For well-known attributes, this bit must be set to 1. For optional attributes, it defines whether the attribute is transitive (when set to 1) or non-transitive (when set to 0).
- The third high-order bit (bit 2) is the partial bit. The partial bit defines whether the information in the optional transitive attribute is partial (when set to 1) or complete (when set to 0). For well-known attributes and for optional non-transitive attributes the partial bit must be set to 0.
- The fourth high-order bit (bit 3) is the extended length bit. The extended length bit defines whether the attribute length is one octet (when set to 0) or two octets (when

set to 1). The attribute flag can use the extended length only if the length of the attribute value is greater than 255 octets.

- If the extended length bit of the attribute flags octet is set to 0, the third octet of the path attribute contains the length of the attribute data in octets.
- If the extended length bit of the attribute flags octet is set to 1, then the third and the fourth octets of the path attribute contain the length of the attribute data in octets.
- The lower-order four bits of the attribute flags octet are unused. The lower-order four bits must be zero (and must be ignored when received).

The attribute type code parameter contains the attribute type code, as defined by the Internet Assigned Numbers Authority (IANA). The attribute type code uniquely identifies the attribute from all others. The remaining octets of the path attribute represent the attribute value and are interpreted according to the attribute flags and the attribute type code parameters.

The following table shows the supported attribute type codes.

Table 60: BGP mandatory path attributes

Attribute	Type code	Description
Origin	1	Defines the origin of the path information: <ul style="list-style-type: none"> • Value = 0 --- IGP (the path is valid all the way to the IGP of the originating AS) • Value = 1--- EGP (the last AS in the AS path uses an EGP to advertise the path) • Value = 2--- Incomplete (the path is valid only to the last AS in the AS path)
AS path	2	Contains a list of the autonomous systems that packets must traverse to reach the destinations. This code represents each AS path segment as follows: <ul style="list-style-type: none"> • path segment type • path segment length • path segment value
Next hop	3	Specifies the IP address of the border router to use as a next hop for the advertised destinations (destinations listed in the NLRI field of the update message).
Multixit discriminator	4	Discriminates among multiple exit or entry points to the same neighboring AS on external (internal-AS) links.
Local preference	5	Indicates the preference that AS border routers assign to a chosen route when they advertise it to iBGP peers

Table 60: BGP mandatory path attributes (continued)

Attribute	Type code	Description
Atomic aggregate	6	Ensures that certain NLRI is not deaggregated
Aggregator	7	Identifies which AS performed the most recent route aggregation. This attribute contains the last AS number that formed the aggregate route followed by the IP address of the BGP speaker that formed the aggregate route.

Attribute Length

The attribute length can be one or two octets in length, depending on the value of the extended length parameter in the attributes flag field.

This parameter indicates the length of the attribute value field.

Attribute Value

The attribute value contains the actual value of the specific attribute. The system implements the attribute value according to the values in the attribute flags and the attribute type code parameters.

NLRI

The NLRI parameter is a variable length field that contains a list of prefixes. The packet size that BGP speakers can exchange limits the number of prefixes in the list.

Equal Cost Multipath

Equal Cost Multipath (ECMP) support allows a BGP speaker to perform route or traffic balancing within an AS by using multiple equal-cost routes submitted to the routing table by OSPF, RIP, or static routes.

For more information about ECMP, see [Equal Cost Multipath](#) on page 2109.

MD5 Message Authentication

Authenticate BGP messages by using Message Digest 5 (MD5) signatures. After you enable BGP authentication, the BGP speaker verifies that the BGP messages it receives from its peers are actually from a peer and not from a third party masquerading as a peer.

BGPv4 TCP MD5 message authentication provides the following features:

- A TCP MD5 signature can exist for BGP peers. You can configure authentication and secret keys for each peer. Peers configured with common secret keys can authenticate each other and exchange routing information.
- The switch can concurrently have BGP peers with authentication enabled and other BGP peers with authentication disabled.
- The switch always encrypts the secret keys.

After you enable BGPv4 TCP MD5 authentication, the router computes an MD5 signature for each TCP packet based on the TCP packet and an individual peer secret key. The router adds this MD5 signature to the TCP packet that contains a BGP message and sends it with the packet, but it does not send the secret key.

The receiver of the TCP packet is configured with the secret key and can verify the MD5 signature. A third party that tries to masquerade as the sender, however, cannot generate an authentic signature because it does not have the secret key.

In commands, the term password refers to the secret key. The secret keys provide security. If the keys are compromised, then the authentication itself is compromised. To prevent this, the switch stores the secret keys in encrypted form.

MD5 Signature Generation

BGP peers calculate MD5 signatures in BGP messages based on the following elements:

- TCP pseudo-header
- TCP header, excluding options
- TCP segment data
- TCP MD5 authentication key

If TCP receives an MD5 authentication key, it reduces its maximum segment size by 18 octets, which is the length of the TCP MD5 option. TCP adds an MD5 signature to each transmitted packet. The peer inserts the resulting 16-byte MD5 signature into the following TCP options: kind=19, length=18.

MD5 Signature Verification

After the switch receives a packet, it performs three tests. The following table lists the tests and the event message that TCP logs if a test fails.

Table 61: MD5 signature verification rules on BGP TCP packets

Condition tested	Action on success	Failure event message
Is the connection configured for MD5 authentication?	Verify that the packet contains a kind=19 option.	TCP MD5 No Signature
Is MD5 authentication enabled for this TCP connection?	TCP computes the expected MD5 signature.	TCP MD5 Authentication Disabled
Does the computed MD5 signature match the received MD5 signature?	TCP sends the packet to BGP.	TCP MD5 Invalid Signature

If a packet passes a test, it proceeds to the next test. After a packet passes all three tests, TCP accepts the packet and sends it to BGP.

If a packet fails a test, the switch logs an event, increments the count of TCP connection errors (wfTcpConnMd5Errors), and discards the packet. The TCP connection remains open.

BGP and Route Redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if OSPF routes exist in a router and they must travel through a BGP network, then configure redistribution of OSPF routes through BGP. This sends OSPF routes to a router that uses BGP.

The switch can redistribute routes:

- on an interface basis.
- on a global basis between protocols on a single VRF instance (intraVRF).
- between the same or different protocols on different VRF instances (interVRF).

Configure interface-based redistribution by configuring a route policy and apply it to the interface. Configure the match parameter to the protocol from which to learn the routes.

You can redistribute routes on a global basis, rather than on an interface basis. Use the **ip bgp redistribute** command to accomplish the (intraVRF) redistribution of routes through BGP, so that BGP redistribution occurs globally on all BGP-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

If you configure redistribution globally and on an interface, redistribution through the route policy takes precedence.

You can redistribute routes from a protocol in one VRF to BGP in another VRF. You can use a route policy for redistribution control. If you enable route redistribution between VRF instances, ensure that IP addresses do not overlap.

Use caution when you configure redistribution. An improperly configured parameter can cause the router to advertise learned eBGP routes out of your local AS. If this happens, the local AS can route other networks.

Do not use redistribution if you peer to an Internet Service Provider (ISP) and do not want traffic to transit your local AS.

When you redistribute OSPF routes into BGP, route priorities can create routing loops. Because BGP has a higher route preference than OSPF external type 1 and 2 routes, if you redistribute OSPF external type 1 and 2 routes into BGP, the router uses the BGP routes, which can cause a routing loop.

Route-Maps and BGP Neighbors

BGP Routing Information Base (BGP RIB) stores routing information received from different peers. BGP RIB has two types of BGP routes, External and Internal (Local). The routes learned from BGP neighbors are External routes and all imported routes are considered as Internal (Local) routes.

In BGP RIB, the OSPF routes redistributed into BGP are considered as Internal (Local) and are matched by route-type only when the keyword is set to **local**. When match

route-type is set to **external**, the route-maps applied on BGP neighbors are ignored and the set operation is not performed.

**Note**

This is applied only on the route-maps applied to BGP neighbors in BGP RIB, and not considered when applying a route-map to the redistribute command.

BGP Route Redistribution and DvR

DvR Controllers redistribute routes (direct routes, static routes and the default route) into the DvR domain. You can configure redistribution of DvR host routes into BGP.

For information on DvR, see [Distributed Virtual Routing](#) on page 915.

Tagged Routes

You can enable route filtering to ensure tagged IS-IS routes are not redistributed into BGP. For more information, see [IS-IS Route Tagging](#) on page 1200.

BGP+

The switch extends the BGPv4 process to support the exchange of IPv6 routes using BGPv4 peering. BGP+ is an extension of BGPv4 for IPv6, which is indicated using the Address Family Identifier (AFI) in the BGP header.

The switch supports capabilities for AFI with the following values: 1 (IPv4) and 2 (IPv6). If the switch receives an OPEN message advertising an AFI with a different value, the connection is closed and a BGP notification message is sent to the peer mentioning unsupported capability.

BGP+ is only supported on the global VRF instance.

**Note**

Ensure you configure IPv6 forwarding for BGP+ to work.

Note that the BGP+ support on the switch is not an implementation of BGPv6. Native BGPv6 peering uses the IPv6 Transport layer (TCPv6) for establishing the BGPv6 peering, route exchanges, and data traffic.

The switch supports the exchange of IPv6 reachability information over IPv4 transport. To support BGP+, the switch supports two BGP protocol extensions, standards RFC 4760 (multi-protocol extensions to BGP) and RFC 2545 (MP-BGP for IPv6). These extensions allow BGPv4 peering to be enabled with IPv6 address family capabilities.

The implementation of BGP+ on the switch uses an existing TCPv4 stack to establish a BGPv4 connection. Optionally, nontransitive BGP properties are used to transfer IPv6 routes over the BGPv4 connection. Any BGP+ speaker has to maintain at least one IPv4 address to establish a BGPv4 connection.

Different from IPv4, IPv6 introduces scoped unicast addresses, identifying whether the address is global or link-local. When BGP+ is used to convey IPv6 reachability

information for interdomain routing, it is sometimes necessary to announce a next hop attribute that consists of a global address and a link-local address. For BGP+, no distinction is made between global and site-local addresses.

The BGP+ implementation includes support for BGPv6 policies, including redistributing BGPv6 into OSPFv3, IS-IS, RIPng, and advertising OSPFv3, IS-IS, RIPng, IPv6 static and local routes into BGPv6 (through BGP+). It also supports the aggregation of global unicast IPv6 addresses.

When configuring BGP+ on the router that is enabled only for IPv6 (the router does not have an IPv4 address), then BGP router ID must be manually configured for the router.

BGP+ does not support confederations. You can configure confederations for IPv4 routes only.

The basic configuration of BGP+ is the same as BGPv4 with one additional parameter added and some existing commands altered to support IPv6 capabilities. You can enable and disable IPv6 route exchange by specifying the address family attribute as IPv6. Note that an IPv6 tunnel is required for the flow of IPv6 data traffic.

BGP+ Tunnel

When you use BGP+ you must configure an IPv6 tunnel and static routes at BGP+ peers.

When BGP+ peers advertise route information, they use Update messages to advertise route information.

These RTM routes contain next-hop addresses from the BGP peer that the route was learned from.

The static routes correlate the next-hop addresses represented by the IPv4-mapped IPv6 address to a specific outgoing interface.

Following is one way to express a static route in an IPv6-configured tunnel for BGP+:

ipv6 route 2001:DB8:0:0:0:ffff:192.0.2.0/24 cost 1 tunnel 10 where 2001:DB8:0:0:0:ffff:192.0.2.0 is the IPv4-mapped IPv6 address of the BGP peer at 192.0.2.0

ECMP with BGP+

The ECMP feature supports and complements BGP+ protocol.

The number of equal-cost-paths supported can differ by hardware platform. For more information, see *Fabric Engine Release Notes*.

You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more

efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.

**Note**

To add BGP+ equal cost paths in the routing table, you must enable the following:

- IPv6 ECMP feature globally
- BGP multiple-paths attribute

BGPv6

BGP peering over IPv6 transport uses a BGPv6 peer to exchange IPv6 routes over an IPv6 transport layer. This is different than BGP+, which enables exchange of IPv6 routes over a BGPv4 peer. Also with BGP+, you must use an IPv6 tunnel to install and configure IPv6 routes in an IPv6 Routing Table Manager (RTM). BGP+ uses an IPv4 mapped IPv6 address for the next hop address and requires you to configure IPv6 static routes and install IPv6 routes in an IPv6 RTM where the next hop for the static route is an IPv6 tunnel interface.

BGPv6 supports the following:

- Input/Output policies.
- Redistribution of OSPFv3, IS-IS, IPv6 static route, and IPv6 direct routes into BGPv6.
- Aggregation of global unicast IPv6 addresses.

**Note**

BGP+ also supports the preceding features.

RFC

The switch supports the BGP multiprotocol extension, as described in RFC 4760. Also supports RFC 2545 (MP-BGP for IPv6).

The BGP protocol extensions ensure peering can be enabled with IPv6 address family capabilities.

Route Exchange

BGPv6 does not exchange any IPv4 routes. BGPv6 advertises or learns only IPv6 routes.

The following table shows the differences between BGPv4 and BGPv6 for route exchange.

Table 62: Differences between BGPv4 and BGPv6 for route exchange

Version	GRT/VRF	IPv4 routes exchange	IPv6 routes exchange
BGPv4	GRT	Supported	Supported (BGP+)
	VRF	Supported	Not supported Note: IPv6 over IPv4 tunnels is not yet virtualized.
BGPv6	GRT	Not supported	Supported
	VRF	Not supported	Supported

Specify the address family attribute as IPv6 to enable IPv6 route exchange.

You can enable IPv6 route exchange by specifying the address family attribute as IPv6. Optionally, you can use non-transitive BGP properties to exchange IPv6 routes between the BGPv6 peering. Any BGPv6 speaker must maintain at least one IPv6 address to establish a BGPv6 connection. The IPv6 scoped unicast addresses can identify the address as global or link-local. If you use BGPv6 to convey IPv6 reachability information for interdomain routing, you can also announce a next hop attribute that consists of a global address and a link-local address.



Note

BGPv6 does not support adjacency on link-local.

Authentication

BGPv6 uses IPsec for security. MD-5 authentication is supported for BGPv4 and is not supported for BGPv6.

The following table shows the differences between BGPv4 and BGPv6 for authentication.

Table 63: Differences between BGPv4 and BGPv6 for authentication

	MD5	IPsec	SHA1/SHA2
BGPv4	Supported	Not supported	Not supported
BGPv6	Not supported	Supported Note: IP Sec is not virtualized, hence BGPv6 is supported only in Global Router mode, and not supported in VRF mode.	Not supported

MD5 Authentication

MD5 authentication is not supported in BGPv6 so it is not necessary to enable MD5 authentication.

IPsec

Only IPsec is supported. Therefore, MD5 authentication cannot be configured.

Consistency Checking

Includes consistency checking for MD5 authentication. BGP peer and BGP peer group configuration for IPv6 addresses include a rule to block MD5 authentication. If you attempt to configure MD5 authentication, you will receive an error message.

IPv6 Tunneling

With BGPv6, IPv6 tunneling is not required for IPv6 data traffic flow. An IPv6 tunnel is required for BGP+.

Circuitless IP

Circuitless IP (CLIP) is a virtual (or loopback) interface that you do not associate with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to your switch as long as an actual path exists to reach the device. For example, as shown in the following figure, a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Note also that an iBGP session exists between two additional addresses 195.39.128.1/32 (CLIP 1) and 195.39.128.2/32 (CLIP 2).




Figure 34: Routers with iBGP connections

The system treats the CLIP interface like an IP interface and treats the network associated with the CLIP as a local network attached to the device. This route always exists and the circuit is always up because no physical attachment exists.

The router advertises routes to other routers in the domain either as external routes using the route-redistribution process or after you enable OSPF in a passive mode to advertise an OSPF internal route. You can configure only the OSPF protocol on the CLIP interface. After you create a CLIP interface, the system software programs a local route with the CPU as the destination ID. The CPU processes all packets destined to the CLIP interface address. The system treats other packets with destination addresses

associated with this network (but not to the interface address) as if they are from an unknown host.

A circuitless IP or CLIP address is a logical IP address for network management, as well as other purposes. The CLIP is typically a host address (with a 32 bit subnet mask). Configure the OSPF router ID to the configured CLIP address. By default, the BGP router ID is automatically equivalent to the OSPF router ID.

For information about how to configure CLIP interfaces, see [Configure a CLIP Interface](#) on page 2138 and [Configure a Circuitless IPv4 Interface](#) on page 2161.

BGP Configuration Considerations and Limitations

Use the information in this section to help you configure BGP on your switch, which supports BGPv4 as described in RFC 1771.

BGP Implementation Guidelines

The following list provides guidelines to successfully implement BGP:

- BGP does not operate with an IP router in nonforwarding (host-only) mode. Make sure that the routers you want BGP to operate with are in forwarding mode.
- If you use BGP for a multihomed AS (one that contains more than a single exit point), use OSPF for your IGP and BGP for your sole exterior gateway protocol, or use intra-AS iBGP routing.
- If OSPF is the IGP, use the default OSPF tag construction. Using EGP or modifying the OSPF tags makes network administration and proper configuration of BGP path attributes difficult.
- For routers that support both BGP and OSPF, the OSPF router ID and the BGP identifier must be the same IP address. The BGP router ID automatically uses the OSPF router ID.
- In configurations where BGP speakers reside on routers that have multiple network connections over multiple IP interfaces (the typical case for iBGP speakers), consider using the address of the circuitless (virtual) IP interface as the local peer address. In this configuration, you ensure that BGP is reachable as long as an active circuit exists on the router.
- By default, BGP speakers do not advertise or inject routes into the IGP. You must configure route policies to enable route advertisement.
- Coordinate routing policies among all BGP speakers within an AS so that every BGP border router within an AS constructs the same path attributes for an external path.
- Configure accept and announce policies on all iBGP connections to accept and propagate all routes. Make consistent routing policy decisions on external BGP connections.

Minimum Requirements

You must configure the following minimum parameters:

- router ID
- local AS number

- enable BGP globally
- BGP neighbor peer session: remote IP addresses
- enable BGP peers
- When you use both BGP and OSPF, the OSPF and BGP router ID must be the same.

The router ID must be a valid IP address of an IP interface on the router or a CLIP address. BGP update messages use this IP address. By default, the BGP router ID automatically uses the OSPF router ID.

You cannot configure the BGP router ID if you configure BGP before you configured the OSPF router ID. You must first disable BGP, configure the OSPF route ID, and then enable BGP globally.

You can add BGP policies to the BGP peer configuration to influence route decisions. BGP policies apply to the peer through the soft-reconfiguration commands.

After you configure the switch for BGP, some parameter changes can require you to enable or disable the BGP global state or the neighbor admin-state.

You can dynamically modify BGP policies. On the global level, the BGP redistribution command has an apply parameter that causes the policy to take effect after you issue the command.

BGP Neighbor Maximum Prefix Configuration

By default, the maximum prefix parameter limits 12 000 NLRI messages for each neighbor. The maximum prefix parameter limits the number of routes that the switch can accept.

The maximum prefix parameter prevents large numbers of BGP routes from flooding the network if you implement an incorrect configuration. You can assign a value to the maximum prefix limit, including 0 (0 means unlimited routes). When you configure the maximum prefix value, consider the maximum number of active routes that your equipment configuration can support.

BGP and OSPF Interaction

RFC1745 defines the interaction between BGP and OSPF when OSPF is the IGP within an autonomous system. For routers that use both protocols, the OSPF router ID and the BGP ID must be the same IP address. You must configure a BGP route policy to allow BGP advertisement of OSPF routes.

Interaction between BGPv4 and OSPF can advertise supernets to support CIDR. BGPv4 supports interdomain supernet advertisements; OSPF can carry supernet advertisements within a routing domain.

BGP and Internet Peering

By using BGP, you can perform Internet peering directly between the switch and another edge router. In such a scenario, you can use each switch for aggregation and link it with a Layer 3 edge router, as shown in the following figure.




Figure 35: BGP and Internet peering

In cases where the Internet connection is single-homed, to reduce the size of the routing table, as a best practice, advertise Internet routes as the default route to the IGP.

For route scaling information, see *Fabric Engine Release Notes*.

Routing Domain Interconnection with BGP

You can implement BGP so that autonomous routing domains, such as OSPF routing domains, connect. This connection allows the two different networks to begin communicating quickly over a common infrastructure, thus providing additional time to plan the IGP merger. Such a scenario is particularly effective when you need to merge two OSPF area 0.0.0.0s, as shown in the following figure.




Figure 36: Routing domain interconnection with BGP

BGP and Edge Aggregation

You can perform edge aggregation with multiple point of presence or edge concentrations. The switch supports 12 pairs (peering services). You can use BGP to inject dynamic routes rather than using static routes or RIP (see the following figure).




Figure 37: BGP and edge aggregation

BGP and ISP Segmentation

You can use the platform as a peering point between different regions or autonomous systems (AS) that belong to the same ISP. In such cases, you can define a region as an OSPF area, an AS, or a part of an AS.

You can divide the AS into multiple regions that each run different IGPs. Interconnect regions logically by using a full iBGP mesh. Each region then injects its IGP routes into iBGP and also injects a default route inside the region. For destinations that do not belong to the region, each region defaults to the BGP border router.

Use the community parameter to differentiate between regions. To provide Internet connectivity, this scenario requires you to make your Internet connections part of the central iBGP mesh (see the following figure).




Figure 38: Multiple regions separated by iBGP

In the preceding figure, consider the following:

- The AS is divided into three regions that each run different and independent IGPs.
- Regions logically interconnect by using a full-mesh iBGP, which also provides Internet connectivity.
- Internal non-BGP routers in each region default to the BGP border router, which contains all routes.
- If the destination belongs to another region, the traffic is directed to that region; otherwise, the traffic is sent to the Internet connections according to BGP policies.

To configure multiple policies between regions, represent each region as a separate AS. Implement eBGP between autonomous systems, and implement iBGP within each AS. In such instances, each AS injects its IGP routes into BGP, where they are propagated to all other regions and the Internet.

The following figure shows the use of eBGP to join several autonomous systems.




Figure 39: Multiple regions separated by eBGP

You can obtain AS numbers from the Inter-Network Information Center (NIC) or use private AS numbers. If you use private AS numbers, be sure to design your Internet connectivity carefully. For example, you can introduce a central, well-known AS to provide interconnections between all private autonomous systems and the Internet. Before it propagates the BGP updates, this central AS strips the private AS numbers to prevent them from leaking to providers.

The following figure illustrates a design scenario in which you use multiple OSPF regions to enable peering with the Internet.




Figure 40: Multiple OSPF regions peering with the Internet

BGP Peers

The following list provides rules related to BGP peers:

- Only metric (=MED) attribute is applied to the output policy if its BGP peer is IBGP
- metric (=MED) and community attributes are applied to output policy if its BGP peer is EBGP
- To influence EBGP and IBGP peers with all applicable BGP attributes, configure **route-map** as an option to **neighbor** command, for example, `neighbor 192.0.2.2 out-route-map policy1`

BGP and Route Aggregation

When you configure the attribute-map with the aggregate command, community, metric, AS Path, and next-hop attributes are set, while the origin attribute is not set.

BGP Session Flapping when IPv6 Forwarding is Enabled or Disabled

In a BGP session that is established with IPv4 and IPv6 capability, disabling or enabling IPv6 forwarding results in BGP session flapping due to capability negotiation. The flapping session in turn affects the IPv4 routing through BGP and the BGP session gets terminated. Ultimately, a capability negotiation takes place to re-establish the IPv4 and IPv6 capable session.

BGP Configuration Using CLI

Configure the Border Gateway Protocol (BGP) to create and maintain an interdomain routing system that guarantees loop-free routing information between autonomous systems (AS).

For information about how to configure route policies for BGP, see [Configure IP Route Policies](#) on page 3261.

Configure BGP

Before You Begin

- To configure the suppress-map, advertise-map, or attribute-map options, the route policy for those options must exist.
- For initial BGP configuration, you must know the AS number.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.



Note

Route refresh is not currently supported on non-default VRFs.

About This Task

Configure BGP globally to enable BGP on the switch and determine how BGP operates.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Specify the AS number and enable BGP:

```
router bgp [WORD <0-11>] [enable]
```



Note

- This command applies only on VRF 0. To enable BGP globally on other VRFs, use the **ip bgp enable** command. You must configure BGP locally before you configure it globally.
- You can also configure AS number on non-default VRFs. For more information, see [Configure an AS Number for a Non-default VRF](#) on page 601.

3. Access Router BGP Configuration mode:

```
router bgp
```

4. Configure BGP variables or accept the default values.

Example

Specify the AS number and enable BGP:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#router bgp 3 enable
```

Variable Definitions

The following table defines parameters for the **router bgp** command.

Variable	Value
<i>WORD <0-11></i>	Specifies the AS number. You cannot enable BGP until you change the local AS to a value other than 0.
<i>enable</i>	Enables BGP on the router.

Use the data in the following table to use the BGP variables in BGP and VRF Router Configuration mode.

Variable	Value
<i>aggregate-address WORD<1-256></i>	Specifies an IP address and its length in the form {a.b.c.d/len}, or an IPv6 address and its length in the form {ipv6addr/len}.
<i>auto-peer-restart enable</i>	Enables the process that automatically restarts a connection to a BGP neighbor. The default value is enable.
<i>auto-summary</i>	When enabled, BGP summarizes networks based on class limits, for example, Class A, B, and C networks. The default value is enable.
<i>bgp always-compare-med</i>	Enables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. The system prefers a path with a lower MED over a path with a higher MED. The default value is disable.
<i>bgp aggregation</i>	Enables the aggregation feature on the interface.
<i>bgp client-to-client reflection</i>	Enables or disables route reflection between two route reflector clients. This variable applies only if the route reflection value is enable. The default value is disable. You can enable route reflection even when clients are fully meshed. This variable only applies to VRF 0.
<i>bgp cluster-id {A.B.C.D}</i>	Configures a cluster ID. This variable applies only if the route reflection value is enable, and if multiple route reflectors are in a cluster. {A.B.C.D} is the IP address of the reflector router. This variable only applies to VRF 0.

Variable	Value
<i>bgp confederation identifier</i> <i><0-4294967295> [peers</i> <i>WORD<0-255>]</i>	Configures a BGP confederation. <i>identifier<0-4294967295></i> specifies the confederation identifier. Use 0-65535 for 2-byte AS and <0-4294967295> for 4-byte AS. <i>peers WORD<0-255></i> lists adjoining autonomous systems that are part of the confederation in the format (5500,65535,0,10,...,...). Use quotation marks (") around the list of autonomous systems. Note: Use this command only on VRF 0.
<i>bgp default local-preference</i> <i><0-2147483647></i>	Specifies the default value of the local preference attribute. The default value is 0. You must disable BGP before you can change the default value.
<i>bgp deterministic-med enable</i>	Enables deterministic MED.
<i>bgp multiple-paths <1-8></i>	Configures the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths the routing table can store. The default value is 1. Note: Configuring the <i>bgp multiple-paths</i> variable does not affect existing routes. The routing table does not show ECMP routes; instead only one route is shown in the routing table. To view Equal-Cost Multipath (ECMP) routes, receive the routes after executing the <i>bgp multiple-paths</i> variable, or toggle the BGP state. The number of equal-cost-paths supported can differ by hardware platform. For more information, see <i>Fabric Engine Release Notes</i> .
<i>comp-bestpath-med-confed</i> <i>enable</i>	When enabled, compares MED attributes within a confederation. The default value is disable. This variable only applies to VRF 0.
<i>debug-screen <off on></i>	Displays debug messages on the console, or saves them in a log file. Disable BGP screen logging (off) or enable BGP screen logging (on).
<i>default-information</i> <i>originate</i>	Enables the advertisement of a default route to peers, if the route exists in the routing table. The default value is disable.
<i>default-information ipv6-</i> <i>originate</i>	Enables the advertisement of an IPv6 default route to peers, if the route exists in the routing table. The default value is disable.

Variable	Value
<code>default-metric</code> <code><-1-2147483647></code>	Configures a value to send to a BGP neighbor to determine the cost of a route a neighbor uses. A default metric value solves the problems associated with redistributing routes that use incompatible metrics. For example, whenever metrics do not convert, using a default metric provides a reasonable substitute and redistribution proceeds. Use this option in conjunction with the redistribute commands so the current routing protocol uses the same metric for all redistributed routes. The default value is 0.
<code>flap-dampening enable</code>	Enables route suppression for routes that flap on and off. The default value is disable.
<code>global-debug mask</code> <code>WORD<1-100></code>	<p>Displays specified debug information for BGP global configurations. The default value is none.</p> <ul style="list-style-type: none"> • <code><WORD 1-100></code> is a list of mask choices separated by commas with no space between choices. <p>Mask choices are:</p> <ul style="list-style-type: none"> • <code>none</code> disables all debug messages. • <code>all</code> enables all debug messages. • <code>error</code> enables display of debug error messages. • <code>packet</code> enables display of debug packet messages. • <code>event</code> enables display of debug event messages. • <code>trace</code> enables display of debug trace messages. • <code>warning</code> enables display of debug warning messages. • <code>state</code> enables display of debug state transition messages. • <code>init</code> enables display of debug initialization messages. • <code>filter</code> enables display of debug messages related to filtering. • <code>update</code> enables display of debug messages related to sending and receiving updates.
<code>ibgp-report-import-rt enable</code>	Configures BGP to advertise imported routes to an interior BGP (iBGP) peer. This variable enables or disables advertisement of nonBGP imported routes to other iBGP neighbors. The default value is enable.
<code>ignore-illegal-rtrid enable</code>	When enabled, BGP overlooks an illegal router ID. For example, you can configure this variable to enable or disable the acceptance of a connection from a peer that sends an open message using a router ID of 0 (zero). The default value is enable.

Variable	Value
<i>neighbor-debug-all mask WORD<1-100></i>	Displays specified debug information for BGP neighbors. The default value is none. For mask options, see the <i>global-debug mask WORD<1-100></i> variable.
<i>no-med-path-is-worst enable</i>	Enables BGP to treat an update without a MED attribute as the worst path. The default value is disable.
<i>quick-start enable</i>	Enables the quick-start flag for exponential backoff.
<i>route-reflector enable</i>	Enables the reflection of routes from iBGP neighbors. The default value is disable. This variable only applies to VRF 0.
<i>route-refresh</i>	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. This variable only applies to VRF 0.
<i>router-id {A.B.C.D}</i>	Specifies the BGP router ID in IP address format. This variable only applies to VRF 0.
<i>synchronization</i>	Enables the router to accept routes from BGP peers without waiting for an update from the IGP. The default value is enable.
<i>traps enable</i>	Enables BGP traps.
<i>vrf-as WORD<0-11></i>	Configures an AS number on a specific VRF instance. Use <i>0-65535</i> for a 2-byte AS and <i><0-4294967295></i> for a 4-byte AS. The default value of 0, or configuring the local-as in the VRF to 0, is equivalent to deleting the local-as configured on user-defined VRFs, and in both cases the local-as on the VRF becomes the local-as on the GlobalRouter.

Job Aid

Use debug command values to control debug messages for global BGP message types, and for message types associated with a specified BGP peer or peer group.

**Tip**

The following tips can help you use the debug commands:

- Display debug commands for multiple mask choices by entering the mask choices separated by commas, with no space between choices.
- To end (disable) the display of debug messages, use the mask choice of none.
- You can save debug messages in a log file, or you can display the messages on your console using the debug-screen command.

For more information about the logged debug messages, see *Fabric Engine Alarms and Logs Reference*.

Configure 4-byte AS numbers

Configure Autonomous System (AS) numbers using the 4-byte format and represent the numbers in octets.

Before You Begin

- You cannot modify the global BGP configuration unless BGP is disabled.
- Configure the local AS number at Global Router (VRF0) only.
- Make sure that you define AS numbers in policies the same way that you configure them for the router. The AS list for the route policies accepts AS number only in the **asplain** format. If you create policies using **asplain** and configure the switch with **asdot**, the match will not occur.

About This Task

Use BGP 4-byte AS numbers to ensure the continuity of loop-free inter-domain routing information between autonomous systems and to control the flow of BGP updates as 2-byte AS numbers will deplete soon. AS Plain notation format is the default and the preferred form of representing 4-byte AS numbers over the AS dot notation format.

You have an option to configure AS dot notation format as well. With AS dot notation, analyzing and troubleshooting any issues encountered becomes difficult as it is incompatible with the regular expressions used by most of the network providers.

If you enable 4-byte AS numbers, or the dotted octet notation, for the Global Router (VRF0), the configuration is inherited by user-defined VRFs. You cannot enable 4-byte AS numbers on individual user-defined VRFs.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Disable BGP to change the AS number format.
`no router bgp enable`
3. Enable the 4-byte AS numbering format.
`router bgp as-4-byte enable`
4. To use the dotted octet notation, enable as-dot.
`router bgp as-dot enable`
5. Configure the 4-byte AS number and enable BGP. If you have enabled as-dot, enter the AS number in octet.
`router bgp WORD<0-11> enable`
6. Access Router BGP Configuration mode:
`router bgp`

7. (Optional) Configure BGP confederation identifier.
`bgp confederation identifier <0-4294967295>`
8. (Optional) Configure BGP confederation peers.
`bgp confederation peers WORD<0-255>`

Example

Disable BGP to change the AS number format.

```
Switch:1(config)#no router bgp enable
```

Enable the 4-byte AS numbering format.

```
Switch:1(config)#router bgp as-4-byte enable
```

To use the dotted octet notation, enable as-dot.

```
Switch:1(config)#router bgp as-dot enable
```

Configure the 4-byte AS number and enable BGP.

```
Switch:1(config)#router bgp 65536 enable
```

Variable Definitions

The following table defines parameters for the **router bgp** command.

Variable	Value
<code>as-4-byte</code> <code><enable></code>	Enables the switch for using 4 byte numbers for an autonomous system (AS). The default value is disable.
<code>as-dot</code> <code><enable></code>	Enables or disables representing AS numbers in octets. The default is disable so the switch uses the plain notation format. If you enable the 4-byte-as and as-dot parameters, enter numbers in the range of 1.0 to 65535.65535. The default value is disable. Note: This parameter is not supported with BGP+.
<code>WORD</code> <code><0-11></code> <code>enable</code>	Sets the local autonomous system (AS) number. You cannot change local-as when BGP is set to enable. <ul style="list-style-type: none"> To set a 2-byte local AS number, enter a local-as number in the range of 0 to 65535. To set a 4-byte local-as number, enable the 4-byte as variable and enter a number in the range of 0 to 4294967295. Note: If as-4-byte is set to false, the range for AS number is 0–65535 and if as-4-byte is set to true, the range is 0–4294967295. If you enable as-dot, enter the AS number in octets in the range of 1.0 to 65535.65535. Note: This parameter is not supported with BGP+.

Configure Aggregate Routes

Before You Begin

- Disable BGP before you enable aggregation.
- You need the appropriate aggregate address and mask.
- If required, policies exist.
- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.



Note

Route refresh is not currently supported on non-default VRFs.

About This Task

Configure aggregate routes so that the router advertises a single route (aggregate route) that represents all destinations. Aggregate routes also reduce the size of routing tables.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

2. Enable BGP aggregation:

```
bgp aggregation enable
```

3. Add an aggregate route to the routing table:

```
aggregate-address WORD<1-256> {advertise-map WORD<0-1536>} [as-set]
[attribute-map WORD<0-1536>] [summary-only] [suppress-map WORD<0-1536>]
```

4. Exit to Global Configuration mode:

```
exit
```

5. Enable BGP:

```
router bgp [<0-65535>] [enable]
```

Example

Add an aggregate route to the routing table and enable BGP:

```
Switch:1(router-bgp)#aggregate-address 2001:DB8::/32 advertise-map map1 attribute-map map2
Switch:1(router-bgp)# router bgp 4 enable
```

Variable Definitions

The following table defines parameters for the **aggregate-address** command.

Variable	Value
<i>advertise-map WORD<0-1536></i>	Specifies the route map name for route advertisements.
<i>as-set</i>	Enables autonomous system information. The default value is disable.
<i>attribute-map WORD<0-1536></i>	Specifies the route map name.
<i>WORD <1-256></i>	Specifies an IP address and its length in the appropriate form. The value must be entered in the format a.b.c.d/len or ipv6addr/len.

Variable	Value
<i>summary-only</i>	Enables the summarization of routes not included in routing updates. This variable creates the aggregate route and suppresses advertisements of more specific routes to all neighbors. The default value is disable.
<i>suppress-map WORD<0-1536></i>	Specifies the route map name for the suppressed route list.

The following table defines parameters for the **router bgp** command.

Variable	Value
<i><0-65535></i>	Specifies the AS number. You cannot enable BGP until you change the local AS to a value other than 0.
<i>enable</i>	Enables BGP on the router.

Configure Allowed Networks

Before You Begin

- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.

About This Task

Configure network addresses to determine the network addresses that BGP advertises. The allowed addresses determine the BGP networks that originate from the switch.

Procedure

- Enter BGP Router Configuration mode:

```
enable

configure terminal

router bgp
```
- Specify IGP network prefixes for BGP to advertise:

```
network <WORD 1-256> [metric <0-65535>]
```

Example

Specify IGP network prefixes for BGP to advertise:

```
Switch:1(router-bgp)#network 2001:DB8::/32 metric 32
```

Variable Definitions

The following table defines parameters for the **network** command.

Variable	Value
<i>WORD</i> <1-256>	Specifies an IP address and its length in the appropriate form.
<i>metric</i> <0-65535>	Specifies the metric to use when the system sends an update for the routes in the network table. The metric configures the MED for the routes advertised to eBGP peers. The range is 0–65535. sec

Configure BGP Peers or Peer Groups

About This Task

Configure peers and peer groups to simplify BGP configuration and make updates more efficient.

BGP speakers can have many neighbors configured with similar update policies. For example, many neighbors use the same distribute lists, filter lists, outbound route maps, and update source. Group the neighbors that use the same update policies into peer groups and peer associations.



Note

- If required, route policies exist.
- You configure BGPv4 on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.
- Route refresh is not currently supported on non-default VRFs.
- Not all parameters are supported on non-default VRFs.

Many of the command variables in this procedure use default values. You can accept the default values or change them to customize the configuration.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

2. Create a peer or peer group:

```
neighbor WORD<0-1536>
```

3. Apply a route policy to all incoming routes:

For BGPv4: `neighbor WORD<0-1536> in-route-map WORD<0-256>`

For BGPv6: `neighbor WORD<0-1536> ipv6-in-route-map WORD<0-256>`

4. Apply a route policy to all outgoing routes:

For BGPv4: `neighbor WORD<0-1536> out-route-map WORD<0-256>`

For BGPv6: `neighbor WORD<0-1536> ipv6-out-route-map WORD<0-256>`

5. (Optional) Configure the source IP address:

`neighbor WORD<0-1536> update-source WORD<1-256>`

6. Enable MD5 authentication (for BGPv4):

`neighbor WORD<0-1536> MD5-authentication enable`

7. Specify an MD5 authentication password (for BGPv4):

`neighbor password <nbr_ipaddr|peer-group-name> WORD<0-1536>`

8. Change the default values for other command variables as required.

9. Enable the configuration:

`neighbor WORD<0-1536> enable`

Example

Create a peer or a peer group:

```
Switch(router-bgp)# neighbor peergroupa
```

Apply a route policy (in-route-map or out-route-map) to all incoming or outgoing routes:

```
Switch(router-bgp)# neighbor peergroupa in-route-map map1 out-route-map map2
```

Configure the source IP address:

```
Switch(router-bgp)# neighbor peergroupa update-source 192.0.2.1
```

Enable MD5 authentication:

```
Switch(router-bgp)# neighbor peergroupa MD5-authentication enable
```

Specify an MD5 authentication password:

```
Switch(router-bgp)# neighbor password peergroupa password
```

Enable the configuration:

```
Switch(router-bgp)# neighbor peergroupa enable
```

Variable Definitions

The following table defines parameters for the **neighbor** command.

Variable	Value
<i>address-family <ipv6></i>	Enables the IPv6 address family on BGP neighbor.
<i>advertisement-interval <5-120></i>	Specifies the time interval, in seconds, that transpires between each transmission of an advertisement from a BGP neighbor. The default value is 5 seconds. The route advertisement interval feature is implemented using the time stamp that indicates when each route is advertised. The time stamp is marked to each route so that the route advertisement interval is compared to the time stamp and BGP is then able to make a decision about whether the route advertisement can be sent or it should be delayed when a better route is received. This feature does not work for a withdraw route because the route entry is already removed when the processing route advertisement is sent and the time stamp marked in the route entry cannot be obtained.
<i>allow-as-in</i>	Allows BGP to inject updates.
<i>default-ipv6-originate</i>	Enables IPv6 BGP neighbor default originate.
<i>default-originate</i>	Enables the switch to send a default route advertisement to the specified neighbor. A default route does not need to be in the routing table. The default value is disable. Do not use this command if default-information originate is globally enabled.
<i>ebgp-multihop</i>	Enables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable.
<i>enable</i>	Enables the BGP neighbor.
<i>fall-over bfd</i>	Enable fall-over Bidirectional Forwarding Detection (BFD).
<i>in-route-map WORD<0-256></i>	Applies a route policy rule to all incoming routes that are learned from, or sent to, the peers or peer groups of the local router. The local BGP router is the BGP router that allows or disallows routes and configures attributes in incoming or outgoing updates. <i>WORD<0-256></i> is an alphanumeric string length (0–256 characters) that indicates the name of the route map or policy.
<i>ipv6-in-route-map WORD <0-256></i>	Creates IPv6 in route map. <i>WORD <0-256></i> specifies the route map name in the range of 0 to 256 characters.

Variable	Value
<i>ipv6-max-prefix</i> <0-2147483647>	Configures a limit on the number of routes that the router can accept from a neighbor. The default value is 12000 routes. A value of 0 (zero) indicates that no limit exists.
<i>ipv6-out-route-map</i> <i>WORD</i> <0-256>	Creates IPv6 out route map. <i>WORD</i> <0-256> specifies the route map name in the range of 0 to 256 characters.
<i>max-prefix</i> <0-2147483647>	Configures a limit on the number of routes that the router can accept from a neighbor. The default value is 12000 routes. A value of 0 (zero) indicates that no limit exists.
<i>MD5-authentication enable</i>	Enables TCP MD5 authentication between two peers. The default value is disable.
<i>neighbor-debug mask</i> <i>WORD</i> <1-100>	<p>Displays specified debug information for a BGP peer. The default value is none.</p> <p><<i>WORD</i> 1-100> is a list of mask choices separated by commas with no space between choices. For example: {<<i>mask</i>>, <<i>mask</i>>, <<i>mask</i>>...}.</p> <p>Mask choices are:</p> <ul style="list-style-type: none"> • <i>none</i> disables all debug messages. • <i>all</i> enables all debug messages. • <i>error</i> enables display of debug error messages. • <i>packet</i> enables display of debug packet messages. • <i>event</i> enables display of debug event messages. • <i>trace</i> enables display of debug trace messages. • <i>warning</i> enables display of debug warning messages. • <i>state</i> enables display of debug state transition messages. • <i>init</i> enables display of debug initialization messages. • <i>filter</i> enables display of debug messages related to filtering. • <i>update</i> enables display of debug messages related to sending and receiving updates.
<i>next-hop-self</i>	<p>When enabled, specifies that the next-hop attribute in an iBGP update is the address of the local router or the router that generates the iBGP update. The default value is disable.</p> <p>You can only configure this variable if the neighbor is disabled.</p>

Variable	Value
<i>out-route-map WORD<0-256></i>	Applies a route policy rule to all outgoing routes that are learned from, or sent to, the peers or peer groups of the local router. The local BGP router is the BGP router that allows or disallows routes and configures attributes in incoming or outgoing updates. <i>WORD<0-256></i> is an alphanumeric string length (0–256 characters) that indicates the name of the route map or policy.
<i>peer-group <WORD 0-1536></i>	Adds a BGP peer to the specified subscriber group. You must create the specified subscriber group before you use this command.
<i>remote-as <WORD 0-11></i>	Configures the remote AS number of a BGP peer or a peer-group. You must disable the admin-state before you can configure this variable. <i><WORD 0-11></i> is an alphanumeric string length (0–11 characters) that indicates the AS number.
<i>remove-private-as enable</i>	Strips private AS numbers when an update is sent. The default value is enable.
<i>retry-interval <1-65535></i>	Configures the time interval, in seconds, for the ConnectRetry timer. The default value is 120 seconds. You can configure the retry interval for BGP neighbors only; you cannot configure the retry interval for BGP peer groups.
<i>route-reflector-client</i>	Configures the specified neighbor or group of neighbors as a route reflector client. The default value is disable. All configured neighbors become members of the client group and the remaining iBGP peers become members of the nonclient group for the local route reflector. Note: This variable only applies to VRF 0.
<i>route-refresh</i>	Enables route refresh for the BGP peer. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. Note: This variable only applies to VRF 0.
<i>send-community</i>	Enables the switch to send the update message community attribute to the specified peer. The default value is disable.
<i>site-of-origin</i>	Specifies a site of origin that is added to the extended communities list in each route from a specific peer.

Variable	Value
<code>soft-reconfiguration-in enable</code>	Enables the router to relearn routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable.
<code>timers <0-21845> <0-65535></code>	Configures timers, in seconds, for the BGP speaker for this peer. <0-21845> is the keepalive time. The default is 60. As a best practice, configure a value of 30 seconds. <0-65535> is the hold time. The default is 180. Switch(router-bgp)# neighbor peergroupa timers 4 6
<code>update-source WORD<1-256></code>	Specifies the source IPv4 address {A.B.C.D.} or IPv6 address to use when the system sends BGP packets to this peer or peer group. You must disable the admin-state before you can configure this variable.
<code>weight <0-65535></code>	Specifies the weight of a BGP peer or peer group, or the priority of updates the router can receive from that BGP peer. The default value is 0. If you have particular neighbors that you want to use for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.
<code>WORD<0-1536></code>	Specifies the peer IP address or the peer group name.

Configure a BGP Peer or Peer Group Password

About This Task

Use this procedure to configure a BGP peer or peer group password for Transmission Control Protocol (TCP) MD5 authentication between two peers.



Note

You configure BGP peer on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip bgp`. The VRF must have an RP Trigger of BGP. Route refresh is not currently supported on non-default VRFs.

Procedure

1. Enter BGP Router Configuration mode:


```
enable
configure terminal
router bgp
```
2. Assign a BGP peer or peer group password:


```
neighbor password <nbr_ipaddr|peer-group=name> WORD <0-1536>
```

Example

Assign a BGP peer or peer group password:

```
Switch:1(router-bgp)# neighbor password peergroupa password1
```

Variable Definitions

The following table defines parameters for the **neighbor password <nbr_ipaddr|peer-group-name>** command.

Variable	Value
<code>password <nbr_ipaddr peer-group-name> WORD <0-1536></code>	Specifies a password for TCP MD5 authentication between two peers. WORD <0-1536> is an alphanumeric string length from 0 to 1536 characters.

Configure Route Redistribution to BGP

Before You Begin

- If required, a route policy exists.
- You can configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.



Note

Route refresh is not currently supported on non-default VRFs.

- Before you redistribute DvR host routes to BGP, you must disable BGP aggregation and BGP auto-summarization of networks, using the commands **no ip bgp aggregation enable** and **no ip bgp auto-summary** respectively.

Disabling these settings ensures that all the DvR host routes are correctly advertised into BGP and are not summarized.



Note

When applying a route map to an inter-vrf redistribution, the route map and any associated IP prefix lists must be configured first on the source VRF before configuring the redistribute policy on the destination VRF.

Inter-vrf redistribution is not supported on IPv6 routes.

About This Task

Configure a redistribution entry to announce routes of a certain source protocol type into the BGP domain such as: DvR routes, static routes, Routing Information Protocol (RIP) routes, or direct routes. Use a route policy to control the redistribution of routes.



Note

When a route map with attributes set to `origin` and `local-pref` is applied to the BGP redistribute command, the attributes are not applied to the redistributed routes.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

2. Create a redistribution instance:

```
redistribute <direct|dvr|ipv6-direct|ipv6-isis|ipv6-static|isis|ospf|ospfv3|rip|ripng|static>
```



Note

Redistribution of ripng routes into BGP is supported only on VRF 0.

3. (Optional) Specify a route policy to govern redistribution:

```
redistribute <direct|dvr|ipv6-direct|ipv6-isis|ipv6-static|isis|ospf|ospfv3|rip|ripng|static> route-map WORD<0-64> [vrf-src WORD<1-16>]
```

4. (Optional) Configure the route metric:

```
redistribute <direct|dvr|ipv6-direct|ipv6-isis|ipv6-static|isis|ospf|ospfv3|rip|ripng|static> metric <0-65535> [vrf-src WORD<1-16>]
```

5. (Optional) Configure the route metric-type:

```
redistribute <direct|dvr|ipv6-direct|ipv6-isis|ipv6-static|isis|ospf|ospfv3|rip|ripng|static> metric-type live-metric [vrf-src WORD<1-16>]
```

6. (Optional) Enable route filtering using tags:

```
redistribute isis auto-filter-tag
```

7. Enable the instance:

```
redistribute <direct|dvr|ipv6-direct|ipv6-isis|ipv6-static|isis|ospf|ospfv3|rip|ripng|static> enable [vrf-src WORD<1-16>]
```

8. Exit BGP Router Configuration mode:

```
exit
```

9. Apply the redistribution instance configuration:

For IPv4: `ip bgp apply redistribute <direct|dvr|isis|ospf|rip|static> [vrf WORD<1-16>] [vrf-src <WORD 1-16>]`

For IPv6: `ipv6 bgp apply redistribute <direct|dvr|isis|ospf|rip|static> [vrf <WORD 1-16>]`

10. Apply BGP redistribution to a specific VRF:

`ip bgp apply redistribute vrf WORD<1-16>`

Changes do not take effect until you apply them.

11. View all routes (including DvR host routes) that are redistributed into BGP:

View routes redistributed from GRT to BGP:

For IPv4: `show ip bgp imported-routes`

For IPv6: `show bgp ipv6 imported-routes`

View routes redistributed to BGP for a specific VRF instance:

For IPv4: `show ip bgp imported-routes [vrf WORD<1-64>] [vrfids WORD<0-512>]`

For IPv6: `show bgp ipv6 imported-routes [WORD<1-256>] [vrf WORD<1-16>] [vrfids WORD<0-255>]`

Examples

Redistribute direct routes from the VRF instance `source1` into BGP, in the GRT context.

Create a redistribution instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router bgp
Switch(router-bgp)#redistribute direct vrf-src source1
```

If required, specify a route policy to govern redistribution:

```
Switch(router-bgp)# redistribute direct route-map policy1 vrf-src source1
```

If required, configure the route metric:

```
Switch:1(router-bgp)# redistribute direct metric 4 vrf-src source1
```

Enable the instance:

```
Switch:1(router-bgp)# redistribute direct enable vrf-src source1
```

Exit BGP Router Configuration mode:

```
Switch:1(router-bgp)# exit
```

Apply the redistribution instance configuration:

```
Switch:1(config)# ip bgp apply redistribute direct vrf-src source1
```

Redistribute DvR routes from the GRT to BGP:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router bgpSwitch:1(router-bgp)#redistribute dvr
Switch:1(router-bgp)#redistribute dvr enable
Switch:1(router-bgp)#exit
Switch:1(config)#ip bgp apply redistribute dvr
```

View the host routes (including DvR host routes) that are redistributed from the GRT to BGP:

```
Switch:1(config)#show ip bgp imported-routes vrf vrf1

=====
BGP Imported Routes - VRF vrf1
=====
ROUTE                                METRIC  COMMUNITY  LOCALPREF  NEXTHOP      ORIGIN
-----
192.0.2.1/255.255.255.0              0       0           100        198.51.100.1  INC
192.0.2.2/255.255.255.0              0       0           100        198.51.100.1  INC
192.0.2.3/255.255.255.0              0       0           100        198.51.100.1  INC
...
...
...
3 out of 763 Total Num of imported routes displayed
```

Redistribute DvR routes to BGP for the specific VRF instance vrf1:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router vrf vrf1
Switch:1(router-vrf)#ip bgp redistribute dvr
Switch:1(router-vrf)#ip bgp redistribute dvr enable
Switch:1(router-vrf)#exit
Switch:1(config)#ip bgp apply redistribute dvr vrf vrf1
```

View the DvR host routes that are redistributed to BGP for vrf vrf1:

```
Switch:1(config)#show ip bgp imported-routes vrf vrf1

=====
BGP Imported Routes - VRF vrf1
=====
ROUTE                                METRIC  COMMUNITY  LOCALPREF  NEXTHOP      ORIGIN
-----
192.0.2.4/255.255.255.0              0       0           100        203.0.113.1   INC
192.0.2.5/255.255.255.0              0       0           100        203.0.113.1   INC
192.0.2.6/255.255.255.0              0       0           100        203.0.113.1   INC
192.0.2.7/255.255.255.0              0       0           100        203.0.113.1   INC
192.0.2.8/255.255.255.0              0       0           100        203.0.113.1   INC
...
...
...
5 out of 675 Total Num of imported routes displayed
```

This example demonstrates redistribution of inter-VRF routes (both direct and DvR routes) to BGP, with a route policy configured.

Redistribute inter-VRF DvR routes between VRFs (with VRF IDs 10 and 30), to BGP.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf 10
Switch:1(router-vrf)#ip prefix-list "test10" 192.0.2.0/24 ge 25 le 32
Switch:1(router-vrf)#route-map "test10" 1
Switch:1(router-vrf)#permit
Switch:1(router-vrf)#enable
Switch:1(router-vrf)#match network "test10"
Switch:1(router-vrf)#set metric 99
Switch:1(router-vrf)#exit

Switch:1(config)#router vrf 30
Switch:1(router-vrf)#ip bgp redistribute direct vrf-src 10
Switch:1(router-vrf)#ip bgp redistribute direct enable vrf-src 10
Switch:1(router-vrf)#ip bgp redistribute dvr vrf-src 10
Switch:1(router-vrf)#ip bgp redistribute dvr route-map "test10" vrf-src 10
Switch:1(router-vrf)#ip bgp redistribute dvr enable vrf-src 10
Switch:1(router-vrf)#exit

Switch:1(config)#ip bgp apply redistribute direct vrf 30 vrf-src 10
Switch:1(config)#ip bgp apply redistribute dvr vrf 30 vrf-src 10
```

Variable Definitions

The following table defines parameters for the **redistribute** and **ip bgp apply redistribute** commands.

Variable	Value
<i>auto-filter-tag</i>	Enables route filtering to ensure tagged routes are not redistributed into BGP. The default is disabled.
<i><direct dvr ipv6-direct ipv6-isis ipv6-static isis ospf ospfv3 rip ripng static ></i>	Specifies the type of routes to redistribute (the protocol source).
<i>enable</i>	Enables the BGP route redistribution instance.
<i>metric <0-65535></i>	Configures the metric to apply to redistributed routes.
<i>metric-type live-metric</i>	Configures the metric type to apply to redistributed routes. When you enable the live-metric option, when BGP redistributes static, RIP, OSPF, IS-IS, or DvR routes, the metric value is taken from the routing table and is set to the Path attributes as a MED value. By default, this option is disabled, which means the BGP MED value is not derived from the metric in the routing table.
<i>route-map WORD<0-64></i>	Configures the route policy to apply to redistributed routes.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the name of a VRF instance.
<code>vrf-src WORD<1-16></code>	Specifies the source VRF instance by name for route redistribution.

Configure Redistribution to BGP+ for VRF 0 (GRT)

Before You Begin

If required, a route policy exists.

About This Task

Configure an IPv6 redistribute entry to announce IPv6 routes of a certain source protocol type into the BGP domain, for example, static, OSPF, IS-IS, RIPv6, or direct routes. Use a route policy to control the redistribution of routes.



Note

When a route map with attributes set to `origin` and `local-pref` is applied to the BGP redistribute command, the attributes are not applied to the redistributed routes.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

2. Create a redistribution instance:

```
redistribute <ipv6-direct|ipv6-isis|ipv6-static|ospfv3|ripng>
```

3. (Optional) If required, specify a route policy to govern redistribution:

```
redistribute <ipv6-direct|ipv6-isis|ipv6-static|ospfv3|ripng> route-  
map WORD <0-64>
```

4. (Optional) If required, configure a route metric:

```
redistribute <ipv6-direct|ipv6-isis|ipv6-static|ospfv3|ripng> metric  
<0-65535>
```

5. Enable the instance:

```
redistribute <ipv6-direct|ipv6-isis|ipv6-static|ospfv3|ripng> enable
```

Unlike IPv4 redistribution, you do not need to manually apply the IPv6 redistribution instance. After you enable the IPv6 redistribution instance, it is automatically applied.

Example

Specify a route policy to govern redistribution by using the following command:

```
Switch:1>enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router bgp
Switch:1(router-bgp)#redistribute ipv6-direct route-map policy2
```

Variable Definitions

The following table defines parameters for the **redistribute <ipv6-direct|ipv6-isis|ipv6-static|ospfv3|ripng>** command.

Variable	Value
<i>enable</i>	Enables the BGP route redistribution instance. The default value is none.
<i>metric <0-65535></i>	Configures the metric to apply to redistributed routes. The default value is 0.
<i>route-map <WORD 0-64></i>	Configures the route policy to apply to redistributed routes. The default value is none.

Job Aid

Use the data in the following table to know how route policies are used for BGP from IPv6 perspective.

Table 64: BGP for IPv6 route policy support

	REDISTRIBUTE					ACCEPT	ANNOUNCE
	IPv6 Direct	IPv6 Static	OSPFv3	IPv6 IS-IS	RIPng	BGP	BGP
MATCH							
as-path						Yes	Yes
community	Yes	Yes	Yes	Yes	Yes	Yes	Yes
community-exact						Yes	Yes
extcommunity						Yes	Yes
interface							
local-preference							
metric				Yes	Yes		
network				Yes	Yes		
next-hop				Yes	Yes		
protocol							
route-source						Yes	

Table 64: BGP for IPv6 route policy support (continued)

	REDISTRIBUTE					ACCEPT	ANNOUNCE
	IPv6 Direct	IPv6 Static	OSPFv3	IPv6 IS-IS	RIPng	BGP	BGP
route-type			Yes				Yes
tag							
vrf							
vrfids							
SET							
as-path						Yes	Yes
as-path-mode						Yes	Yes
automatic-tag							
community						Yes	Yes
community-mode						Yes	Yes
injectlist	Yes	Yes	Yes	Yes	Yes		
ip-preference							
local-preference						Yes	Yes
mask							
metric	Yes	Yes	Yes	Yes	Yes	Yes	Yes
metric-type							
metric-type-internal							
next-hop						Yes	Yes
nssa-pbit							
origin							Yes
origin-egp-as							
tag							
weight						Yes	

Configure AS Path Lists

Before You Begin

- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.



Note

Route refresh is not currently supported on non-default VRFs.

About This Task

Configure an AS path list to restrict the routing information a router learns or advertises to and from a neighbor. The AS path list acts as a filter that matches AS paths.

Procedure

- Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

- Create the path list:

```
ip as-list <1-1024> memberid <0-65535> <permit|deny> as-path WORD<0-1536>
```

Use this command for each member by specifying different member IDs.

Example

Create the path list:

```
Switch:1(config)# ip as-list 234 memberid 3456 permit as-path "5"
```

Variable Definitions

The following table defines parameters for the **ip as-list** command.

Variable	Value
<0-65535>	Specifies an integer value between 0–65535 that represents the regular expression entry in the AS path list.
<1-1024>	Specifies an integer value from 1–1024 that represents the AS-path list ID you want to create or modify.
<permit deny>	Permits or denies access for matching conditions.
WORD<0-1536>	Specifies the AS number as an integer value between 0–1536. Place multiple AS numbers within quotation marks (").

Configure Community Lists

Before You Begin

- You configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.



Note

Route refresh is not currently supported on non-default VRFs.

About This Task

Configure community lists to specify permitted routes by using their BGP community. This list acts as a filter that matches communities or AS numbers.

Procedure

- Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

- Create a community list:

```
ip community-list <1-1024> memberid <0-65535> <permit|deny> community-string WORD<0-256>
```

Example

Create a community list:

```
Switch:1(config)# ip community-list 1 memberid 4551 permit community-string internet
```

Variable Definitions

The following table defines parameters for the **ip community-list** command.

Variable	Value
<0-65535>	Specifies an integer value from 0–65535 that represents the member ID in the community list.
<1-1024>	Specifies an integer value from 1–1024 that represents the community list ID.

Variable	Value
<code><permit deny></code>	Configures the access mode, which permits or denies access for matching conditions.
<code>WORD<0-256></code>	<p>Specifies the community as an alphanumeric string value with a string length from 0–256 characters. Enter this value in one of the following formats:</p> <ul style="list-style-type: none"> • (AS num:community-value) • (well-known community string) <p>Well known communities include: internet, no-export, no-advertise, local-as (known as NO_EXPORT_SUBCONFED).</p>

Configure Extended Community Lists

Before You Begin

- Configure BGP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip bgp**. The VRF must have an RP Trigger of BGP.



Note

Route refresh is not currently supported on non-default VRFs.

About This Task

Configure community lists to specify permitted routes by BGP extended community attributes, including route targets and sites of origin (SOO). This list acts as a filter that matches route targets and SOO.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```

2. Create an extended community list based on the route target attribute:

```
ip extcommunity-list <1-1024> memberId <0-65535> rt {<0-65535>
<0-2147483647>|<A.B.C.D> <0-65535>} [soo {<0-65535> <0-2147483647>|
<A.B.C.D> <0-65535>}]
```

You can optionally configure the SOO attributes at the end of the same command or you can configure the SOO separately using the syntax in the following step.

3. Create an extended community list based on the SOO attribute:

```
ip extcommunity-list <1-1024> memberId <0-65535> soo {<0-65535>
<0-2147483647>|<A.B.C.D> <0-65535>}
```

Example

Create an extended community list based on the route target attribute:

```
Switch:1(config)# ip extcommunity-list 1 memberid 234 rt 192.0.2.1 5 soo 32 45
```

Variable Definitions

The following table defines parameters for the **ip extcommunity-list** command.

Variable	Value
<1-1024>	Specifies an integer value from 1-1024 that represents the community list ID you want to create or modify.
memberId <0-65535>	Specifies an integer value from 0-65535 that represents the member ID in the community list.
rt <0-65536> <0-2147483647> rt <A.B.C.D> <0-65535>	Specifies the route target in the format {AS number:assigned number} (that is, {0-65535}:{0-2147483647}) or {ipaddress:assigned number} (that is, {a.b.c.d}:{0-65535}).
soo <0-65535> <0-2147483647> soo <A.B.C.D> <0-65535>	Specifies the site of origin in the format {AS number:assigned number} (that is, {0-65535}:{0-2147483647}) or {ipaddress:assigned number} (that is, {a.b.c.d}:{0-65535}).

Configure an AS Number for a Non-default VRF

Before You Begin

- Disable BGP synchronization.

About This Task

The Autonomous System (AS) number configured on the global Virtual Routing Forwarding (VRF) instance, called the GlobalRouter (GRT), is inherited by all user-created VRFs by default, however, you can override the AS number for the specific VRF instance using the following procedure.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Set the AS number:

```
ip bgp vrf-as WORD<0-11>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrfred
Switch:1(router-vrf)#ip bgp vrf-as 3
```

Variable Definitions

The following table defines parameters for the **ip bgp vrf-as** command.

Variable	Value
<i>WORD</i> <0–11>	<p>Configures the local autonomous system (AS) number for the specific VRF instance. You cannot change local-as when BGP is set to enable.</p> <ul style="list-style-type: none"> To configure a 2-byte local AS number, enter a local-as number in the range of 0 to 65535. To configure a 4-byte local-as number, enable the 4-byte as variable and enter a number in the range of 0 to 4294967295. <p>Note: If as-4-byte is configured to false, the range for AS number is 0–65535 and if as-4-byte is configured to true, the range is 0–4294967295.</p> <p>If you enable as-dot, enter the AS number in octets in the range of 1.0 to 65535.65535.</p> <p>The AS number in a specific VRF instance inherits the AS number in the GlobalRouter in the following instances:</p> <ul style="list-style-type: none"> Configuring the AS number in a specific VRF instance to 0 (ip bgp vrf-as 0). Deleting the AS number in a specific VRF instance (no ip bgp vrf-as or default ip bgp vrf-as).

BGP Verification Using CLI

About This Task

Use **show** commands to verify Border Gateway Protocol (BGP) configuration and to monitor or troubleshoot BGP operation.



Note

If the next hop of a BGP route is resolved using an IS-IS route, show commands can display the IS-IS internal next hop from the 127.1.x.y class rather than the IS-IS sys name.

View BGP Aggregate Information

About This Task

Display information about current aggregate addresses.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about current aggregates:

```
show ip bgp aggregates [<prefix/len>] [vrf WORD <1-16>] [vrfids
WORD<0-255>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp aggregates** command.

Variable	Value
<prefix/len>	Specifies the IP address and the mask length.
vrf WORD<1-16>	Specifies a VRF instance by name.
vrfids WORD<0-255>	Specifies a range of VRFs by ID number.

View IPv6 BGP+ Aggregate Information

About This Task

Display information about current IPv6 aggregate addresses.

Use BGP 4 byte AS numbers to ensure the continuity of loop-free inter-domain routing information between ASs and to control the flow of BGP updates as 2 byte AS numbers will deplete soon.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about current IPv6 aggregates:

```
show bgp ipv6 aggregates [<WORD 1-256>] [vrf <WORD 1-16>] [vrfids
<0-255>]
```

Variable Definitions

The following table defines parameters for the **show bgp ipv6 aggregates** command.

Variable	Value
WORD <1-256>	Specifies the IPv6 prefix and the prefix length (the length can be 0 to 128).
vrf WORD <1-16>	Specifies a VRF instance by name (the string length ranges from 1-16 characters).
vrfids <0-255>	Specifies a range of VRFs by ID number (the ID ranges from 0-255).

View CIDR Routes

About This Task

Display information about classless interdomain routing (CIDR) routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about CIDR routes:

```
show ip bgp cidr-only [<prefix/len>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp cidr-only** command.

Variable	Value
<prefix/len>	Specifies an exact match of the prefix. This variable is an IP address and an integer value from 0–32 in the format a.b.c.d/xx.
vrf WORD<1-16>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
vrfids WORD<0-512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

View BGP Configuration

About This Task

View information about the BGP configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about the current BGP configuration:

```
show ip bgp conf [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>show ip bgp conf vrf vrf1
=====
                        BGP Configuration - VRF vrf1
=====

                        BGP version - 4
                        local-as - 22610
                        Identifier - 27.82.217.1
                        BGP on/off - ON
                        BGP origin - AUTO-SENSE
                        as-4-byte - enable
                        as-4-byte-origin - AUTO-SENSE
                        as-dot - disable
                        aggregation - enable
                        always-cmp-med - disable
                        auto-peer-restart - enable
```

```

auto-summary - enable
comp-bestpath-med-confed - disable
default-local-preference - 100
default-metric - -1
deterministic-med - disable
flap-dampening - disable
debug-screen - Off
global-debug - none
ibgp-report-import-rt - enable
ignore-illegal-rtrid - enable
max-equalcost-routes - 1
no-med-path-is-worst - enable
route-refresh - disable
orig-def-route - disable
orig-v6-def-route - disable
quick-start - disable
synchronization - enable

--More-- (q = quit)

```

Variable Definitions

The following table defines parameters for the **show ip bgp conf** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF instance by name (the string length ranges from 1-16 characters).
<code>vrfids WORD<0-512></code>	Specifies a range of VRFs by ID number (the ID ranges from 0-512).

View BGP Confederation

About This Task

Display information about BGP confederations.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about current BGP confederations:

```
show ip bgp confederation
```

Example

```

Switch:1>show ip bgp confederation
confederation identifier 0
confederation peer as

```

View Flap-Dampened Routes

About This Task

Display information about flap-dampened routes to determine unreliable routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about flap-dampened routes:

```
show ip bgp dampened-paths {A.B.C.D} [<prefix/len>] [longer-prefixes]
[vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp dampened-paths** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the source IP address in the format a.b.c.d.
<i>longer-prefixes</i>	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<i><prefix/len></i>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
<i>vrf WORD<1-16></i>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
<i>vrfids WORD<0-512></i>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

View the Global Flap-Dampening Configuration

About This Task

Display global information about flap-dampening.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display global information about flap-dampening:

```
show ip bgp flap-damp-config [prefix/len] [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```

Example

```
Switch:1>show ip bgp flap-damp-config vrf vrf1
```

```
=====
                        BGP Flap Dampening - VRF vrf1
=====
                        Status - enable
                        PolicyName - N/A
                        CutoffThreshold - 1536
                        ReuseThreshold - 512
                        Decay - 2
                        MaxHoldDown - 180
```

Variable Definitions

The following table defines parameters for the **show ip bgp flap-damp-config** command.

Variable	Value
<i><prefix/len></i>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
<i>vrf WORD<1–16></i>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
<i>vrfids WORD<0–512></i>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

View Imported Routes

About This Task

Display information about BGP imported routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGP imported routes:

```
show ip bgp imported-routes [<prefix/len>] [longer-prefixes] [vrf
WORD<1–16>] [vrfids WORD<0–512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp imported-routes** command.

Variable	Value
<i>longer-prefixes</i>	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<i><prefix/len></i>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
<i>vrf WORD<1–16></i>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
<i>vrfids WORD<0–512></i>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

View BGPv6 Imported Routes

About This Task

Display information about BGPv6 imported routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGPv6 imported routes:

```
show bgp ipv6 imported-routes [<prefix/len>] [longer-prefixes] [vrf WORD<1-16>] [vrfids WORD<0-255>]
```

Variable Definitions

The following table defines parameters for the **show bgp ipv6 imported-routes** command.

Variable	Value
<i>longer-prefixes</i>	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<i><prefix/len></i>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
<i>vrf WORD<1-16></i>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
<i>vrfids WORD<0-255></i>	Specifies a range of VRFs by ID number (the ID ranges from 0–255).

View BGP Neighbors Information

About This Task

Display information about BGP neighbors.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGP neighbors:

```
show ip bgp neighbors [{A.B.C.D}] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
3. Display information about BGP peer advertised routes:

```
show ip bgp neighbors {A.B.C.D} advertised-routes [<prefix/len>] [longer-prefixes] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
4. Display information about BGP peer routes:

```
show ip bgp neighbors {A.B.C.D} routes [<prefix/len>] [community <enable|disable>] [longer-prefixes] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
5. Display statistics for BGP peers:

```
show ip bgp neighbors {A.B.C.D} stats [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>show ip bgp neighbors vrf vrf1
```

```

=====
BGP Neighbor Info - VRF vrf1
=====

BGP neighbor is 200.200.200.63 remote AS 63, Internal Peer, MP-BGP-capable, BGP state
[Established] UP Time 0 day(s), 07:27:24 remote router ID 63.1.1.1

      vrf instance - 0
      admin-state - BGP ON
      connect-retry-interval - 120
      ebgp-multihop - disable
      hold-time - 30
      keepalive-time - 10
      hold-time-configured - 180
      keepalive-time-configured - 60
      max-prefix - 12000
      nexthop-self - disable
      originate-def-route - disable
      MD5-authentication - disable
      neighbor-debug - all
      remove-private-as - disable
      as-override - disable
      allow-as-in - 0
      route-advertisement-interval - 5
      route-reflector-client - disable
      send-community - disable
      soft-reconfiguration-in - disable
      updt-source-interface - 0.0.0.0
      weight - 100
      Route Policy In -
      Route Policy Out -
      address-family vpnv4 - disable
      route-refresh - disable

Total bgp neighbors -
1

```

Variable Definitions

The following table defines parameters for the **show ip bgp neighbors** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the IP address.
<i>community <enable disable></i>	Enables or disables the display of community attributes.
<i>longer-prefixes</i>	Shows long prefixes. The longer-prefixes indicate the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
<i>prefix/len</i>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
<i>vrf WORD<1–16></i>	Specifies a VRF instance by name.
<i>vrfids WORD<0–512></i>	Specifies a range of VRFs by ID number.

View BGPv6 Neighbors Information

View information about BGPv6 neighbors.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View information about BGPv6 neighbors:

```
show bgp ipv6 neighbors [WORD<1-256>] [vrf <WORD 1-16>] [vrfids <0-255>]
```
3. View information about BGPv6 peer advertised routes:

```
show bgp ipv6 neighbors WORD<1-256> advertised-routes [WORD<1-256>] [longer-prefixes] [vrf <WORD 1-16>] [vrfids <0-255>]
```
4. View information about BGPv6 peer routes:

```
show bgp ipv6 neighbors WORD<1-256> routes [WORD<1-256>] [community <enable|disable>] [vrf <WORD 1-16>] [vrfids <0-255>]
```

Example

The following examples shows the summary output for **show bgp ipv6 neighbors** command, and the *advertised-routes* and *routes* variable options.

```
Switch:1>show bgp ipv6 neighbors vrf vrf1

=====
                        BGPv6 Neighbor Info - VRF vrf1
=====
BGPv6 neighbor is 2015:cdba:0:0:0:0:3257:9652 remote AS 200, External Peer,
BGP state [Established] UP Time 0 day(s), 00:50:30
remote router ID 0.0.0.6

                        vrf instance - 0
                        admin-state - BGP ON
connect-retry-interval - 120
                        ebgp-multihop - disable
                        hold-time - 180
                        keepalive-time - 60
                        hold-time-configured - 180
keepalive-time-configured - 60
                        ipv6-max-prefix - 8000
                        nexthop-self - disable
                        originate-defv6-route - disable
                        neighbor-debug - all
                        remove-private-as - disable

route-advertisement-interval - 5
route-reflector-client - disable
                        send-community - disable
soft-reconfiguration-in - enable
                        updt-source-interface - 0:0:0:0:0:0:0:0
                        weight - 100
                        IPv6Route Policy In -
                        IPv6Route Policy Out -
                        address-family ipv6 - enable
                        route-refresh - enable

Total bgpv6 neighbors: 1
```

```
Switch:1>show bgp ipv6 neighbors 2015:cdba:0:0:0:0:3257:9655 advertised-routes vrf vrf1
```



```

The total number of routes advertised to the neighbor is 2
=====
                        BGPv6 Neighbor Advertised Routes - VRF vrfl
=====
NETWORK/MASK                NEXTHOP ADDRESS                LOC PREF ORG   STATUS
-----
2001:cdba:0:0:0:0:0/64      2001:cdba:0:0:0:0:3257:9651      100    INC    Best
2007:cdba:0:0:0:0:0/64      2001:cdba:0:0:0:0:3257:9651      100    INC    Used
-----

Switch:1>show bgp ipv6 neighbors 2015:cdba:0:0:0:0:3257:9655 routes vrf vrfl
-----
The total number of accepted routes from the neighbor is 2
=====
                        BGPv6 Neighbor Routes - VRF vrfl
=====
NETWORK/MASK                PEER-REM-ADDR                NEXTHOP-ADDRESS                ORG LOC-PREF STATUS
-----
1100:0:0:0:0:0:0/64        2015:cdba:0:0:0:0:3257:9655  2015:cdba:0:0:0:0:3257:9655  INC 100 Used AS_PATH:(150)
2015:cdba:0:0:0:0:0/64     2015:cdba:0:0:0:0:3257:9655  2015:cdba:0:0:0:0:3257:9655  INC 100 Best AS_PATH:(150)
-----

```

context

Variable Definitions

The following table defines parameters for the **show bgp ipv6 neighbors** command.

Variable	Value
<i>WORD<1-256></i>	Specifies the IPv6 address.
<i>advertised-routes</i>	Specifies an IPv6 neighbors advertised routes.
<i>routes</i>	Specifies an IPv6 neighbors routes.
<i>WORD<1-256></i>	Specifies an IPv6 address/length.
<i>longer-prefixes</i>	Shows long prefixes. The longer-prefixes indicate the mask length from any specified prefix to 128. For example, show from prefix :X::X/X/len to X:X::X/X/ 128.
<i>community <enable disable></i>	Enables or disables the display of community attributes.
<i>vrf</i>	Specifies a VRF instance by name (the string length ranges from 1-16 characters).
<i>vrfids</i>	Specifies a range of VRFs by ID number (the ID ranges from 0-255).

View BGP Network Configurations

About This Task

Display information about BGP network configurations.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display information about BGP network configurations:

```
show ip bgp networks [<prefix/len>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp networks** command.

Variable	Value
<prefix/len>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
vrf WORD<1-16>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
vrfids WORD<0-512>	Specifies a range of VRFs by ID number (the ID ranges from 0–512).

View IPv6 BGP+ Network Configurations

About This Task

Display information about BGP+ network configurations.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGP+ network configurations:

```
show bgp ipv6 networks <WORD 1-256> [vrf <WORD 1-16>] [vrfids <0-255>]
```

Variable Definitions

The following table defines parameters for the **show bgp ipv6 networks** command.

Variable	Value
<WORD 1-256>	Specifies the IPv6 prefix and the prefix length (must be an integer value between 0 and 128).
vrf	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
vrfids	Specifies a range of VRFs by ID number (the ID ranges from 0–255).

View BGP Peer Group Information

About This Task

Display information about BGP peer groups.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGP peer groups:

```
show ip bgp peer-group [WORD<0-1536>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp peer-group** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF instance by name (the string length ranges from 1-16 characters).
<code>vrfids WORD<0-512></code>	Specifies a range of VRFs by ID number (the ID ranges from 0-512).
<code>WORD<0-1536></code>	Specifies the name of the peer group (the string length ranges from 0-1536 characters).

View BGP Redistributed Routes

About This Task

Display information about BGP redistributed routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGP redistributed routes:

```
show ip bgp redistributed-routes [<prefix/len>] [vrf WORD<1-16>]  
[vrfids WORD<0-512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp redistributed-routes** command.

Variable	Value
<code><prefix/len></code>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0-32).
<code>vrf WORD<1-16></code>	Specifies a VRF instance by name (the string length ranges from 1-16 characters).
<code>vrfids WORD<0-255></code>	Specifies a range of VRFs by ID number (the ID ranges from 0-255).

View BGPv6 Redistributed Routes

About This Task

Display information about BGPv6 redistributed routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGPv6 redistributed routes:

```
show bgp ipv6 redistributed-routes [vrf <WORD 1-16>] [vrfids <0-255>]
```

Variable Definitions

The following table defines parameters for the **show bgp ipv6 redistributed-routes** command.

Variable	Value
vrf	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
vrfids	Specifies a range of VRFs by ID number (the ID ranges from 0–255).

View a Summary of BGP Configurations

About This Task

Display summarized information about BGP.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display summarized information about BGP:

```
show ip bgp summary [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

The following example shows partial output for the **show ip bgp summary** command.

```
Switch:1>show ip bgp summary vrf vrf1

=====
                        BGP Summary - VRF vrf1
=====

                        BGP version - 4
                        local-as - 22610
                        Identifier - 27.82.217.1
                        Decision state - Idle
                        The total number of routes is 0

BGP NEIGHBOR INFO :

NEIGHBOR      RMTAS      STATE      HLDTM  KPALV  HLD CFG  KPCFG  WGHT  CONRTY  ADVINT  UPTIME
```

```

-----
192.0.2.1      22620      Active      0      0      180      60      100      120      5      0 day(s), 07:25:09
Total bgp neighbors: 1

BGP CONFEDERATION INFO :
confederation identifier 0
confederation peer as

--More-- (q = quit)

```

Variable Definitions

The following table defines parameters for the **show ip bgp summary** command.

Variable	Value
<i>vrf</i> <i>WORD</i> <1-16>	Specifies a VRF instance by name.
<i>vrfids</i> <i>WORD</i> <0-512>	Specifies a range of VRFs by ID number.

View a Summary of BGPv6 Configurations

About This Task

View a summary of BGP peering over IPv6 transport.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View BGPv6 summary:

```
show bgp ipv6 summary [vrf <WORD 1-16>] [vrfids <0-255>]
```

Example

The following example shows partial output for the **show bgp ipv6 summary** command.

```

Switch:1>show bgp ipv6 summary vrf vrf1

=====
                        BGP ipv6 Summary - VRF vrf1
=====

                        BGP version - 4
                        local-as - 200
                        Identifier - 0.0.0.6
                        Decision state - Idle
                        The total number of routes is 1

BGPv6 NEIGHBOR INFO :

NEIGHBOR                RMTAS  STATE      HLDTM  KPALV  HLDCFG  KPCFG  WGHT  CONRTY  ADVINT
-----
2001:DB8:0:0:0:0:ffff    50    Established 180    60    180    60    100    120
5

Total bgpv6 neighbors: 1

BGP CONFEDERATION INFO :

```

```

confederation identifier 0
confederation peer as

BGPv6 NETWORK INFO :

=====
                        BGPv6 Networks - VRF vrf1
=====
)

```

Variable Definitions

The following table defines parameters for the **show bgp ipv6 summary** command.

Variable	Value
<i>vrf</i>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
<i>vrfids</i>	Specifies a range of VRFs by ID number (the ID ranges from 0–255).

View BGP Routes

About This Task

Display information about BGP routes.



Note

BGP stores route information on the AVL tree and this command retrieves that information. Information in the AVL tree is not sorted. The information returned by this command will not be displayed in any particular order.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about BGP routes:

```
show ip bgp route [<prefix/len>] [community <enable|disable>] [ip
{A.B.C.D}] [longer-prefixes] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable Definitions

The following table defines parameters for the **show ip bgp route** command.

Variable	Value
<i>community <enable disable></i>	Enables or disables the display of community attributes.
<i>ip {A.B.C.D}</i>	Specifies an IP address.
<i>longer-prefixes</i>	Shows long prefixes. Longer-prefixes indicates the mask length from a specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).

Variable	Value
<code><prefix/len></code>	Shows paths with this prefix. The prefix is the IP address and exact mask length (must be an integer value from 0–32).
<code>vrf WORD<1-16></code>	Specifies a VRF instance by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRFs by ID number.

View BGPv6 Routes

About This Task

Display information about BGPv6 routes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Enter Privileged EXEC mode:
enable
3. Display information about BGP routes:
show bgp ipv6 route [`<WORD 1-256>`] [`longer-prefixes`] [`community <enable|disable>`] [`ipv6 <WORD 1-256>`] [`vrf <WORD 1-16>`] [`vrfids <0-255>`]

Variable Definitions

The following table defines parameters for the **show bgp ipv6 route** command.

Variable	Value
<code>[<WORD 1-256>]</code>	Specifies the IPv6 prefix and the prefix length (must be an integer value between 0 and 128).
<code>community <enable disable></code>	Enables or disables the display of community attributes.
<code>ipv6 <WORD 1-256>]</code>	Specifies an IPv6 address.
<code>longer-prefixes</code>	Shows long prefixes. The longer-prefixes indicate the mask length from any specified prefix to 128 (for example, show from prefix X:X::X/X/len to X:X::X/X/ 128).
<code>vrf</code>	Specifies a VRF instance by name (the string length ranges from 1–16 characters).
<code>vrfids</code>	Specifies a range of VRFs by ID number (the ID ranges from 0–255).

BGP Configuration Using EDM

About This Task

Configure Border Gateway Protocol (BGP) to create an inter-domain routing system that guarantees loop-free routing information between autonomous systems.

For information about how to configure route policies, see [Configure a Route Policy](#) on page 3275.

Configure BGP

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Enable BGP so that BGP runs on the router. Configure general BGP parameters to define how BGP operates on the system.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Generals** tab.
4. In AdminStatus, select **enable**.
5. Configure the local autonomous system (AS) ID.
6. In the **Aggregate** area, enable or disable route aggregation as required.
7. Configure the BGP options as required.
8. In the **DebugMask** area, select the type of information to show for BGP debugging purposes.
9. Configure BGP confederations as required.
10. Configure BGP route reflectors as required.
11. Select **Apply**.

Generals *Field Descriptions*

Use the data in the following table to use the **Generals** tab.

Name	Description
VrfBgpOrigin	Specifies the origin of the BGP configuration in the VRF. config indicates the BGP instance is created manually. autoSense indicates the BGP instance was generated from Auto-sense.
bgpVersion	Specifies the version of BGP that operates on the router. Note: This parameter only applies to VRF 0.
Identifier	Specifies the BGP router ID number.
AdminStatus	Enables or disables BGP on the router. The default is disable. You cannot enable AdminStatus until you change the LocalAS value to a nonzero value.
4ByteAs	Enables or disables 4-byte AS numbers. The default is disable. Note: This parameter only applies to VRF 0.
LocalAs	Configures the local AS number. You cannot change the LocalAs value if AdminStatus is enable. The switch does not support this parameter with BGP +. Note: If the inserted LocalAs is 0, then the LocalAs in that VRFcontext loses its significance and it becomes the LocalAs configured in GlobalRouter (the equivalence to CLI commands ip bgp vrf-as 0 and no ip bgp vrf-as or default ip bgp vrf-as).

Name	Description
AsDot	<p>Enables or disables representing AS numbers in octets. The default is disable so the switch uses the plain notation format</p> <p>The AS dot notation is easier to read and remember than the AS plain notation, but it can be difficult to convert from AS plain to AS dot. The IETF prefers the AS plain notation.</p> <p>The switch does not support this parameter with BGP +.</p> <p>Note: This parameter only applies to VRF 0.</p>
Aggregate	Enables or disables aggregation. The default is enable.
DefaultMetric	<p>Configures the metric sent to BGP neighbors. The default metric determines the cost of a route a neighbor uses. Use this parameter in conjunction with the redistribute parameters so that BGP uses the same metric for all redistributed routes.</p> <p>The default is -1. The range is -1-2147483647.</p>
DefaultLocalPreference	Specifies the default local preference. The local preference indicates the preference that AS border routers assign to a chosen route when they advertise it to iBGP peers. The default is 100. The range is 0-2147483647.
AlwaysCompareMed	Enables or disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. The system prefers a path with a lower MED over a path with a higher MED. The default is disable.
DeterministicMed	Enables or disables deterministic MED. Deterministic MED compares the MEDs after routes advertised by different peers in the same AS are chosen. The default is disable.
AutoPeerRestart	Enables or disables the process that automatically restarts a connection to a BGP neighbor. The default is enable.
AutoSummary	Enables or disables automatic summarization. If you enable this variable, BGP summarizes networks based on class limits (for example, Class A, B, or C networks). The default is enable.
NoMedPathIsWorst	Enables or disables NoMedPathIsWorst. If you enable this variable, BGP treats an update without a MED attribute as the worst path. The default is enabled.
BestPathMedConfed	Enables or disables the comparison of MED attributes within a confederation. The default is disable.

Name	Description
DebugMask	<p>Displays the specified debug information for BGP global configurations. The default value is none. Other options are</p> <ul style="list-style-type: none"> • none disables all debug messages. • event enables the display of debug event messages. • state enables display of debug state transition messages. • update enables display of debug messages related to updates transmission and reception. • error enables the display of debug error messages. • trace enables the display of debug trace messages. • init enables the display of debug initialization messages. • all enables all debug messages. • packet enables the display of debug packet messages. • warning enables the display of debug warning messages. • filter enables the display of debug messages related to filtering.
IgnoreIllegalRouterId	Enables BGP to overlook an illegal router ID. For example, this variable enables the acceptance of a connection from a peer that sends an open message using a router ID of 0. The default is enable.
Synchronization	Enables or disables the router to accept routes from BGP peers without waiting for an update from the IGP. The default is enable.
MaxEqualcostRoutes	Configures the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths the routing table can store. The default value is 1; the range is 1–8.
IbgpReportImportRoute	Configures BGP to report imported routes to an interior BGP (iBGP) peer. This variable also enables or disables reporting of non-BGP imported routes to other iBGP neighbors. The default is enable.
FlapDampEnable	Enables or disables route suppression for routes that go up and down (flap). The default is disable.
QuickStart	Enables or disables the Quick Start feature, which forces the BGP speaker to begin establishing peers immediately, instead of waiting for the auto-restart timer to expire. The default is disable.
TrapEnable	Enables or disables the BGP traps. The default is disable.

Name	Description
ConfederationASIdentifier	Specifies a BGP confederation identifier in the range of 0–65535. Note: This parameter applies only to VRF 0.
ConfederationPeers	Lists adjoining autonomous systems that are part of the confederation in the format (5500,65535,0,10,...).. This value can use 0–255 characters. Note: This parameter applies only to VRF 0.
RouteReflectionEnable	Enables or disables the reflection of routes from iBGP neighbors. The default is enable. Note: This parameter applies only to VRF 0.
RouteReflectorClusterId	Configures a reflector cluster ID IP address. This variable applies only if you enable RouteReflectionEnable, and if multiple route reflectors are in a cluster. Note: This parameter applies only to VRF 0.
ReflectorClientToClientReflection	Enables or disables route reflection between two route reflector clients. This variable applies only if RouteReflectionEnable is enable. The default is enable. Note: This parameter applies only to VRF 0.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. Note: This parameter only applies to VRF 0.

Configure 4-byte AS numbers

Before You Begin

- You cannot modify the global BGP configuration unless BGP is disabled.

- Make sure that you define AS numbers in policies the same way that you configure them for the router. The choices are asplain (regular expression) or asdot (dot notation). If you create policies using asplain and configure the switch with asdot, the match will not occur.

About This Task

Configure AS numbers using the 4-byte format and represent the numbers in octets.

Use BGP 4-byte AS numbers to ensure the continuity of loop-free inter-domain routing information between autonomous systems and to control the flow of BGP updates as 2 byte AS numbers will deplete soon. AS Plain notation format is the default and the preferred form of representing 4-byte AS numbers over the AS dot notation format.

You have an option to configure AS dot notation format as well. With AS dot notation, analyzing and troubleshooting any issues encountered becomes difficult as it is incompatible with the regular expressions used by most of the network providers.

If you enable 4-byte AS numbers, or the dotted octet notation, for the Global Router (VRF0), the configuration is inherited by user-defined VRFs. You cannot enable 4-byte AS numbers on individual user-defined VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Generals** tab.
4. To change the AS number format, select **disable** for **AdminStatus**.
5. Select **Apply**.
6. In **4-byteAs**, select **enable**.
7. In **AsDot**, select **enable**.
8. In **LocalAs**, type the 4-byte AS number in octets.
9. In **AdminStatus**, select **enable**.
10. Select **Apply**.

Generals *Field Descriptions*

Use the data in the following table to use the **Generals** tab.

Name	Description
VrfBgpOrigin	Specifies the origin of the BGP configuration in the VRF. config indicates the BGP instance is created manually. autoSense indicates the BGP instance was generated from Auto-sense.
bgpVersion	Specifies the version of BGP that operates on the router. Note: This parameter only applies to VRF 0.
Identifier	Specifies the BGP router ID number.
AdminStatus	Enables or disables BGP on the router. The default is disable. You cannot enable AdminStatus until you change the LocalAS value to a nonzero value.
4ByteAs	Enables or disables 4-byte AS numbers. The default is disable. Note: This parameter only applies to VRF 0.
LocalAs	Configures the local AS number. You cannot change the LocalAs value if AdminStatus is enable. The switch does not support this parameter with BGP +. Note: If the inserted LocalAs is 0, then the LocalAs in that VRFcontext loses its significance and it becomes the LocalAs configured in GlobalRouter (the equivalence to CLI commands ip bgp vrf-as 0 and no ip bgp vrf-as or default ip bgp vrf-as).

Name	Description
AsDot	<p>Enables or disables representing AS numbers in octets. The default is disable so the switch uses the plain notation format</p> <p>The AS dot notation is easier to read and remember than the AS plain notation, but it can be difficult to convert from AS plain to AS dot. The IETF prefers the AS plain notation.</p> <p>The switch does not support this parameter with BGP +.</p> <p>Note: This parameter only applies to VRF 0.</p>
Aggregate	Enables or disables aggregation. The default is enable.
DefaultMetric	<p>Configures the metric sent to BGP neighbors. The default metric determines the cost of a route a neighbor uses. Use this parameter in conjunction with the redistribute parameters so that BGP uses the same metric for all redistributed routes.</p> <p>The default is -1. The range is -1-2147483647.</p>
DefaultLocalPreference	Specifies the default local preference. The local preference indicates the preference that AS border routers assign to a chosen route when they advertise it to iBGP peers. The default is 100. The range is 0-2147483647.
AlwaysCompareMed	Enables or disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. The system prefers a path with a lower MED over a path with a higher MED. The default is disable.
DeterministicMed	Enables or disables deterministic MED. Deterministic MED compares the MEDs after routes advertised by different peers in the same AS are chosen. The default is disable.
AutoPeerRestart	Enables or disables the process that automatically restarts a connection to a BGP neighbor. The default is enable.
AutoSummary	Enables or disables automatic summarization. If you enable this variable, BGP summarizes networks based on class limits (for example, Class A, B, or C networks). The default is enable.
NoMedPathIsWorst	Enables or disables NoMedPathIsWorst. If you enable this variable, BGP treats an update without a MED attribute as the worst path. The default is enabled.
BestPathMedConfed	Enables or disables the comparison of MED attributes within a confederation. The default is disable.

Name	Description
DebugMask	<p>Displays the specified debug information for BGP global configurations. The default value is none. Other options are</p> <ul style="list-style-type: none"> • none disables all debug messages. • event enables the display of debug event messages. • state enables display of debug state transition messages. • update enables display of debug messages related to updates transmission and reception. • error enables the display of debug error messages. • trace enables the display of debug trace messages. • init enables the display of debug initialization messages. • all enables all debug messages. • packet enables the display of debug packet messages. • warning enables the display of debug warning messages. • filter enables the display of debug messages related to filtering.
IgnoreIllegalRouterId	Enables BGP to overlook an illegal router ID. For example, this variable enables the acceptance of a connection from a peer that sends an open message using a router ID of 0. The default is enable.
Synchronization	Enables or disables the router to accept routes from BGP peers without waiting for an update from the IGP. The default is enable.
MaxEqualcostRoutes	Configures the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths the routing table can store. The default value is 1; the range is 1–8.
IbgpReportImportRoute	Configures BGP to report imported routes to an interior BGP (iBGP) peer. This variable also enables or disables reporting of non-BGP imported routes to other iBGP neighbors. The default is enable.
FlapDampEnable	Enables or disables route suppression for routes that go up and down (flap). The default is disable.
QuickStart	Enables or disables the Quick Start feature, which forces the BGP speaker to begin establishing peers immediately, instead of waiting for the auto-restart timer to expire. The default is disable.
TrapEnable	Enables or disables the BGP traps. The default is disable.

Name	Description
ConfederationASIdentifier	Specifies a BGP confederation identifier in the range of 0–65535. Note: This parameter applies only to VRF 0.
ConfederationPeers	Lists adjoining autonomous systems that are part of the confederation in the format (5500,65535,0,10,...).. This value can use 0–255 characters. Note: This parameter applies only to VRF 0.
RouteReflectionEnable	Enables or disables the reflection of routes from iBGP neighbors. The default is enable. Note: This parameter applies only to VRF 0.
RouteReflectorClusterId	Configures a reflector cluster ID IP address. This variable applies only if you enable RouteReflectionEnable, and if multiple route reflectors are in a cluster. Note: This parameter applies only to VRF 0.
ReflectorClientToClientReflection	Enables or disables route reflection between two route reflector clients. This variable applies only if RouteReflectionEnable is enable. The default is enable. Note: This parameter applies only to VRF 0.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates it contains in its database that are eligible for the peer that issues the request. Note: This parameter only applies to VRF 0.

View BGP Global Stats

About This Task

View BGP global stats.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Global Stats** tab.

Global Stats *Field Descriptions*

Use the data in the following table to use the **Global Stats** tab.

Name	Description
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.
Starts	Displays the number of times the BGP connection started.
Stops	Displays the number of times the BGP connection stopped.
Opens	Displays the number of times BGP opens TCP.
Closes	Displays the number of times BGP closes TCP.
Fails	Displays the number of times TCP attempts failed.
Fatals	Displays the number of times TCP crashes due to fatal error.
ConnExps	Displays the number of times the TCP retry timer expired.
HoldExps	Displays the number of times the hold timer expired.
KeepExps	Displays the number of times the keepalive timer expired.
RxOpens	Displays the number of open instances BGP receives.
RxKeeps	Displays the number of keepalive instances BGP receives.
RxUpdates	Displays the number of update instances BGP receives.
RxNotifys	Displays the number of notification instances BGP receives.
TxOpens	Displays the number of open instances BGP transmitted.

Name	Description
TxKeeps	Displays the number of keepalive instances BGP transmitted.
TxUpdates	Displays the number of updates instances BGP transmits.
TxNotifys	Displays the number of notification instances BGP transmits.
BadEvents	Displays the number of invalid events FSM received.
SyncFails	Displays the number of times FDB sync failed.
TrEvent	Displays the trace event.
RxECodeHeader	Displays the total header errors received.
RxECodeOpen	Displays the total open errors received.
RxECodeUpdate	Displays the total update errors received.
RxECodeHoldtimer	Displays the total hold timer errors received.
RxECodeFSM	Displays the total FSM errors received.
RxECodeCease	Displays the total cease errors received.
RxHdrCodeNoSync	Displays the header not synchronized errors received.
RxHdrCodeInvalidMsgLen	Displays the header invalid message length errors received.
RxHdrCodeInvalidMsgType	Displays the header invalid message type errors received.
RxOpCodeBadVer	Displays the open errors received for Bad Version.
RxOpCodeBadAs	Displays the open errors received for le Bad AS Number.
RxOpCodeBadRtrID	Displays the open errors received for Bad BGP Rtr ID.
RxOpCodeUnsuppOption	Displays the open errors received for Unsupported Option.
RxOpCodeAuthFail	Displays the open errors received for Auth Failures.
RxOpCodeBadHold	Displays the open errors received for Bad Hold Value.
RxUpdCodeMalformedAttrList	Displays the update errors received for Malformed Attr List.
RxUpdCodeWelKnownAttrUnrecog	Displays the update errors received for Welknown Attr Unrecog.
RxUpdCodeWelknownAttrMiss	Displays the update errors received for Welknown Attr Missing.

Name	Description
RxUpdCodeAttrFlagError	Displays the update errors received for Attr Flag Error.
RxUpdCodeAttrLenError	Displays the update errors received for Attr Len Error.
RxUpdCodeBadORIGINAttr	Displays the update errors received for Bad ORIGIN Attr.
RxUpdCodeASRoutingLoop	Displays the update errors received for AS Routing Loop.
RxUpdCodeBadNHAttr	Displays the update errors received for Bad NEXT-HOP Attr.
RxUpdCodeOptionalAttrError	Displays the update errors received for Optional Attr Error.
RxUpdCodeBadNetworkField	Displays the update errors received for Bad Network Field.
RxUpdCodeMalformedASPath	Displays the update errors received for Malformed AS Path.
TxECodeHeader	Displays the total Header errors transmitted.
TxECodeOpen	Displays the total Open errors transmitted.
TxECodeUpdate	Displays the total Update errors transmitted.
TxECodeHoldtimer	Displays the total Hold timer errors transmitted.
TxECodeFSM	Displays the total FSM errors transmitted.
TxECodeCease	Displays the total Cease errors transmitted.
TxHdrCodeNoSync	Displays the header Not Synchronized errors transmitted.
TxHdrCodeInvalidMsgLen	Displays the header Invalid msg len errors transmitted.
TxHdrCodeInvalidMsgType	Displays the header Invalid msg type errors transmitted.
TxOpCodeBadVer	Displays the open errors transmitted for Bad Version.
TxOpCodeBadAs	Displays the open errors transmitted for Bad AS Number.
TxOpCodeBadRtrID	Displays the open errors transmitted for Bad BGP Rtr ID.
TxOpCodeUnsuppOption	Displays the open errors transmitted for Unsupported Option.
TxOpCodeAuthFail	Displays the open errors transmitted for Auth Failures.
TxOpCodeBadHold	Displays the open errors transmitted for Bad Hold Value.
TxUpdCodeMalformedAttrList	Displays the update errors transmitted for Malformed Attr List.

Name	Description
TxUpdCodeWelknownAttrUnrecog	Displays the update errors transmitted for Welknown Attr Unrecog.
TxUpdCodeWelknownAttrMiss	Displays the update errors transmitted for Welknown Attr Missing.
TxUpdCodeAttrFlagError	Displays the update errors transmitted for Attr Flag Error.
TxUpdCodeAttrLenError	Displays the update errors transmitted for Attr Len Error.
TxUpdCodeBadORIGINAttr	Displays the update errors transmitted for Bad ORIGIN Attr.
TxUpdCodeASRoutingLoop	Displays the update errors transmitted for AS Routing Loop
TxUpdCodeBadNHAttr	Displays the update errors transmitted for Bad NEXT-HOP Attr
TxUpdCodeOptionalAttrError	Displays the update errors transmitted for Optional Attr Error.
TxUpdCodeBadNetworkField	Displays the update errors transmitted for Bad Network Field.
TxUpdCodeMalformedASPath	Displays the update errors transmitted for Malformed AS Path.

Configure Aggregate Routes

Before You Begin

- Enable aggregate routes globally.
- You need the appropriate aggregate address and mask.
- If required, ensure the required policies exist.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure aggregate routes so that the router advertises a single route (aggregate route) that represents all destinations. Aggregate routes also reduce the size of routing tables.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.

3. Select the **Aggregates** tab.
4. Select **Insert**.
5. Configure the aggregate **Address** and **PrefixLen**.
6. Select **AsSetGenerate** or **SummaryOnly** as required.
7. Configure policies for the aggregate route.
8. Select **Insert**.

Aggregates *Field Descriptions*

Use the data in the following table to use the **Aggregates** tab.

Name	Description
Address	Specifies the aggregate IP address.
PrefixLen	Specifies the aggregate PrefixLen.
AsSetGenerate	Enables or disables AS-set path information generation. The default is disable.
SummaryOnly	Enables or disables the summarization of routes in routing updates. Enable this parameter to create the aggregate route and suppress advertisements of more-specific routes to all neighbors. The default is disable.
SuppressPolicy	Specifies the route policy (by name) used for the suppressed route list. Enable this parameter to create the aggregate route and suppress advertisements of the specified routes.
AdvertisePolicy	Specifies the route policy (by name) used for route advertisements. The route policy selects the routes that create AS-set origin communities.
AttributePolicy	Specifies the route policy (by name) used to determine aggregate route attributes.

Configure Aggregate IPv6 Routes

Before You Begin

- Aggregate routes are enabled.
- You have determined the appropriate aggregate prefix and length.
- If required, policies exist.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure IPv6 aggregate routes so that the router advertises a single route (aggregate route) that represents all destinations. Aggregate routes also reduce the size of routing tables.

To configure aggregate routes for IPv4, see [Configure Aggregate Routes](#) on page 631.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **BGP+**.
3. Select the **Aggregates** tab.
4. Select **Insert**.
5. Specify the aggregate **Address** and **PrefixLen**.
6. (Optional) Configure **AsSetGenerate** and **SummaryOnly** as required.
7. (Optional) Configure policies for the aggregate route.
8. Select **Insert**.

Aggregates Field Descriptions

Use the data in the following table to use the **Aggregates** tab.

Name	Description
Address	Specifies the aggregate address. The default is none.
PrefixLen	Specifies the length of the prefix (in bits).
AsSetGenerate	Enables or disables AS-set path information generation. The default is disable.
SummaryOnly	Enables or disables the summarization of routes in routing updates. Enable this parameter to create the aggregate route and suppress advertisements of more-specific routes to all neighbors. The default is disable.
SuppressPolicy	Specifies the route policy (by name) used for the suppressed route list. Enable this parameter to create the aggregate route and suppress advertisements of the specified routes.
AdvertisePolicy	Specifies the route policy (by name) used for route advertisements. The route policy selects the routes that create AS-set origin communities.
AttributePolicy	Specifies the route policy (by name) used to determine aggregate route attributes.

Configure Allowed Networks

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure network addresses to determine the network addresses that BGP advertises. The allowed addresses determine the BGP networks that originate from the switch.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Network** tab.
4. Select **Insert**.
5. Configure the network address, mask, and metric.
6. Select **Insert**.

Network Field Descriptions

Use the data in the following table to use the **Network** tab.

Name	Description
NetworkAfAddr	Specifies the network prefix that BGP advertises.
NetworkAfPrefixLen	Specifies the prefix length of the network address.
NetworkAfMetric	Specifies the metric to use when the system sends an update for the routes in the network table. The metric configures the MED for the routes advertised to eBGP peers. The range is 0–65535.

Configure Allowed IPv6 Networks

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure IPv6 network addresses to determine the network addresses that BGP advertises. The allowed addresses determine the BGP networks that originate from the switch.

To configure allowed IPv4 networks, see [Configure Allowed Networks](#) on page 634.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **BGP+**.
3. Select the **Network** tab.
4. Select **Insert**.
5. Configure the network address, prefix length, and metric.
6. Select **Insert**.

Network Field Descriptions

Use the data in the following table to use the **Network** tab.

Name	Description
NetworkAfAddr	Specifies the network prefix that BGP advertises. The default is none.
NetworkAfPrefixLen	Specifies the network prefix length. The default is none.
NetworkAfMetric	Specifies the metric used when an update is sent for the routes in the network table. The metric configures the MED for the routes advertised to EBGp peers. The range is 0 to 65535. The default is 0.

Configure BGP Peers

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure BGP peers to connect two routers to each other for the purpose of exchanging routing information. BGP peers exchange complete routing information only after they establish the peer connection.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Peers** tab.
4. Select **Insert**.
5. Configure the peer as required.
6. Select **Insert**.
7. In the **Enable** column, double-click the value, and then select **true**.
By default, new peer configuration parameters are disabled.
8. Select **Apply**.
9. To modify a peer configuration, double-click the value, and then select a new value.
10. Select **Apply**.

Peers Field Descriptions

Use the data in the following table to use the **Peers** tab.

Name	Description
Instance	Specifies the BGP peer instance.
LocalAddrType	Specifies the local IP address type of the entered BGP peer.
LocalAddr	Specifies the local IP address of the entered BGP peer.
RemoteAddrType	Specifies the remote IP address type of the entered BGP peer.
RemoteAddr	Specifies the remote IP address of the entered BGP peer.
AdminStatus	Specifies the administrative status of the BGP peer.
GroupName	Specifies the peer group name to which the peer belongs (optional).
AccessType Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Configures a BGP neighbor as part of the trusted or unrestricted segment of a Segmented VRF. The default is the trusted segment. The switch adds existing BGP neighbors in the VRF context to the trusted segment automatically. EDM only displays this field if you launch the VRF context view. This field does not apply to the GlobalRouter.
PeerState	Specifies the BGP peer connection state.
RemoteAs	Configures a remote AS number for the peer or peer-group.
Enable	Controls whether the peer connection is enabled or disabled. The default is disabled.
EbgpMultiHop	Enables or disables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable.

Name	Description
RoutePolicyIn	Specifies the policy (by name) that applies to all network routes learned from this peer.
RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing route updates.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
UpdateSourceInterface	Specifies the source IP address to use when the switch sends eBGP packets to this peer or peer group.
ConnectRetryInterval	Specifies the time interval, in seconds, for the connect retry timer. The suggested value for this timer is 120 seconds.
HoldTimeConfigured	Specifies the time interval, in seconds, for the hold time for this BGP speaker with this peer. This value is in an open message sent to this peer by this BGP speaker. To determine the hold time with the peer, the switch compares this value with the HoldTime value in an open message received from the peer. The HoldTime must be at least three seconds. If the value is zero, the hold time does not establish with the peer. The suggested value for this timer is 180 seconds.
KeepAliveConfigured	Specifies the time interval, in seconds, for the KeepAlive timer configured for this BGP speaker with this peer. KeepAliveConfigured determines the keep alive message frequency relative to HoldTimeConfigured; KeepAlive indicates the actual time interval for the keep alive messages. The maximum value for this timer is one-third of HoldTimeConfigured. If KeepAliveConfigured is zero, no periodic keep alive messages are sent to the peer after the peers establish a BGP connection. Configure a value of 60 seconds.
MD5Authentication	Enables and disables MD5 authentication.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between each transmission of an advertisement from a BGP neighbor. The default value is 30 seconds. The route advertisement interval feature is implemented using the time stamp that indicates when each route is advertised. The time stamp is marked to each route so that the route advertisement interval is compared to the time stamp and BGP is then able to make a decision about whether the route advertisement can be sent or should be delayed when a better route is received. This feature does not work for a withdraw route because the route entry is already removed when the processing route advertisement is sent and the time stamp marked in the route entry cannot be obtained.
DefaultOriginate	When enabled, specifies that the current route originated from the BGP peer. This parameter enables or disables sending the default route information to the specified neighbor or peer. The default value is false.

Name	Description
DefaultOriginatelpv6	When enabled, specifies that the current IPv6 route originated from the BGP peer. This parameter enables or disables sending the default IPv6 route information to the specified neighbor or peer. The default value is false.
Weight	Specifies the peer or peer group weight, or the priority of updates the system can receive from this BGP peer. The default value is 100.
MaxPrefix	Configures a limit on the number of routes accepted from a neighbor. The default value is 12000 routes. A value of 0 means no limit exists.
NextHopSelf	Specifies that the next-hop attribute in an iBGP update is the address of the local router or the router that generates the iBGP update. The default is disable.
RouteReflectorClient	Specifies that this peer is a route reflector client. Note: This parameter only applies to VRF 0.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
DebugMask	Displays the specified debug information for the BGP peer. The default value is none. <ul style="list-style-type: none"> • None disables all debug messages. • Event enables the display of debug event messages. • State enables display of debug state transition messages. • Update enables display of debug messages related to updates transmission and reception. • Error enables the display of debug error messages. • Trace enables the display of debug trace messages. • Init enables the display of debug initialization messages. • All enables all debug messages. • Packet enables the display of debug packet messages. • Warning enables the display of debug warning messages. • Filter enables the display of debug messages related to filtering.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer. The default value is disable.
Vpnv4Address	Specifies the vpnv4 routes.

Name	Description
IpvpnLiteCap	Enable or disable IP VPN-lite capability on the BGP neighbor peer.
Ipv6Cap	Enable or disable the IPv6 capability on the BGP neighbor peer. The default value is disable.
SooAddress	Specifies the site-of-origin (SoO) address of the BGP peer.
SooAsNumber	Specifies the site-of-origin (SoO) Autonomous System (AS) number of the BGP peer.
SooAssignedNum	Specifies the site-of-origin (SoO) assigned number of the BGP peer.
SooType	Specifies the site-of-origin (SoO) type of the BGP peer.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates in the database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer. The default is disable.
AllowAsIn	Specifies the number of AS-in allowed for the BGP peer.
Ipv6RoutePolicyIn	Specifies the policy (by name) that applies to all network IPv6 routes learned from this peer.
Ipv6RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing IPv6 route updates.
Ipv6MaxPrefix	Configures a limit on the number of IPv6 routes accepted from a neighbor. A value of 0 means no limit exists.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for this BGP peer.

Configure BGPv6 Peers

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure BGPv6 peers to connect two routers to each other for the purpose of exchanging routing information. BGPv6 peers exchange complete routing information only after they establish the peer connection.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **BGP+**.
3. Select the **Peers** tab.
4. Select **Insert**.
5. Configure the peer, as required.
6. Select **Insert**.
7. In the **Enable** column, double-click the value, and then select **enable**.
By default, new peer configuration parameters are disabled.
8. Select **Apply**.
9. To modify a peer configuration, double-click the value, and then select a new value.
10. Select **Apply**.

Peers Field Descriptions

Use the data in the following table to use the **Peers** tab.

Name	Description
RemoteAddr	Specifies the remote IPv6 address of the entered BGP+ peer.
GroupName	Specifies the peer group name to which the peer belongs (optional).
PeerState	Specifies the BGPv6 peer connection state.
RemoteAs	Configures a remote AS number for the peer or peer-group in the range 0 to 65535.
Enable	Controls whether the peer connection is enabled or disabled. The default is disabled.
EbgpMultiHop	Enables or disables a connection to a BGPv6 peer that is more than one hop away from the local router. The default value is disable.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
UpdateSourceInterface	Specifies the source IP address to use when the switch sends eBGP packets to this peer or peer group.
ConnectRetryInterval	Specifies the time interval, in seconds, for the connect retry timer. The suggested value for this timer is 120 seconds. The range is 1 to 65535.

Name	Description
HoldTimeConfigured	Specifies the time interval, in seconds, for the hold time for this BGP speaker with this peer. This value is in an open message sent to this peer by this BGP speaker. To determine the hold time with the peer, the switch compares this value with the HoldTime value in an open message received from the peer. The HoldTime must be at least three seconds. If the value is zero, the hold time does not establish with the peer. The suggested value for this timer is 180 seconds. The range is 0 to 65535.
KeepAliveConfigured	Specifies the time interval, in seconds, for the KeepAlive timer configured for this BGP speaker with this peer. KeepAliveConfigured determines the keep alive message frequency relative to HoldTimeConfigured; KeepAlive indicates the actual time interval for the keep alive messages. The maximum value for this timer is one-third of HoldTimeConfigured. If KeepAliveConfigured is zero, no periodic keep alive messages are sent to the peer after the peers establish a BGP connection. Configure a value of 60 seconds. The range is 0 to 21845.
MD5Authentication	Enables and disables MD5 authentication.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between each transmission of an advertisement from a BGPv6 neighbor. The default value is 30 seconds and the range is 5 to 120 seconds. The route advertisement interval feature is implemented using the time stamp that indicates when each route is advertised. The time stamp is marked to each route so that the route advertisement interval is compared to the time stamp and BGP is then able to make a decision about whether the route advertisement can be sent or should be delayed when a better route is received. This feature does not work for a withdraw route because the route entry is already removed when the processing route advertisement is sent and the time stamp marked in the route entry cannot be obtained.
DefaultOriginateIpv6	When enabled, specifies that the current IPv6 route originated from the BGP peer. This parameter enables or disables sending the default IPv6 route information to the specified neighbor or peer. The default value is false.
Weight	Specifies the peer or peer group weight, or the priority of updates the system can receive from this BGP peer. The default value is 100 and the range is 0 to 65535.
NextHopSelf	Specifies that the next-hop attribute in an iBGP update is the address of the local router or the router that generates the iBGP update. The default is disable.

Name	Description
RouteReflectorClient	Specifies that this peer is a route reflector client. Note: This parameter only applies to VRF 0.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is disable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
DebugMask	Displays the specified debug information for the BGP peer. The default value is none. <ul style="list-style-type: none"> • None disables all debug messages. • Event enables the display of debug event messages. • State enables display of debug state transition messages. • Update enables display of debug messages related to updates transmission and reception. • Error enables the display of debug error messages. • Trace enables the display of debug trace messages. • Init enables the display of debug initialization messages. • All enables all debug messages. • Packet enables the display of debug packet messages. • Warning enables the display of debug warning messages. • Filter enables the display of debug messages related to filtering.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer. The default value is disable.
Vpnv4Address	Enables BGP address families for IPv4 (BGP) and Layer 3 VPN (MP-BGP) support. Enable this parameter for VPN/VRF Lite routes.
IvpnLiteCap	Enable or disable IP VPN-lite capability on the BGP neighbor peer.
Ipv6Cap	Enable or disable the IPv6 capability on the BGP neighbor peer. The default value is disable.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates in the database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride Note:	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer. The default is disable.

Name	Description
AllowAsIn Note:	Specifies the number of AS-in allowed for the BGP peer. The range is 1–10.
Ipv6RoutePolicyIn	Specifies the policy (by name) that applies to all network IPv6 routes learned from this peer.
Ipv6RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing IPv6 route updates.
Ipv6MaxPrefix	Configures a limit on the number of IPv6 routes accepted from a neighbor. A value of 0 means no limit exists.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for this peer.

Configure Peer Groups

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure or edit peer groups to create update policies for neighbors in the same group.

Procedure

- In the navigation pane, expand **Configuration > IP**.
- Select **BGP**.
- Select the **Peer Groups** tab.
You can modify an existing parameter by double-clicking the value.
- Select **Insert**.
- Configure the peer group as required.
- Select **Insert**.
- Select **Apply**.

Peer Groups *Field Descriptions*

Use the data in the following table to use the **Peer Groups** tab.

Name	Description
Index	Specifies the index of this peer group.
GroupName	Specifies the peer group to which this neighbor belongs (optional).
AccessType Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Configures a BGP peer group as part of the trusted or unrestricted segment of a Segmented VRF. The default is the trusted segment. The switch adds existing BGP neighbors in the VRF context to the trusted segment automatically. EDM only displays this field if you launch the VRF context view. This field does not apply to the GlobalRouter.
Enable	Enables or disables the peer group.
RemoteAs	Configures a remote AS number for the peer-group.
DefaultOriginate	When enabled, the BGP speaker (the local router) sends the default route 0.0.0.0 to a group of neighbors for use as a default route. The default is disabled.
DefaultOriginateIpv6	When enabled, the BGP speaker (the local router) sends the default route to a group of neighbors for use as a default route. The default is disabled.
EbgpMultiHop	When enabled, the switch accepts and attempts BGP connections to external peers that reside on networks that do not directly connect. The default is disabled.
AdvertisementInterval	Specifies the time interval, in seconds, that elapses between BGP routing updates. The default value is 30 seconds.
KeepAlive	Specifies the time interval, in seconds, between sent BGP keep alive messages to remote peers. The default value is 60.
HoldTime	Configures the hold time for the group of peers in seconds. Use a value that is three times the value of the KeepAlive time. The default value is 180.
Weight	Assigns an absolute weight to a BGP network. The default value is 100.
MaxPrefix	Limits the number of routes accepted from this group of neighbors. A value of zero indicates no limit. The default value is 12,000 routes.
NextHopSelf	Specifies that the switch must set the NextHop attribute to the local router address before sending updates to remote peers.
RoutePolicyIn	Specifies the route policy that applies to all networks learned from this group of peers.
RoutePolicyOut	Specifies the route policy that applies to all outgoing updates to this group of peers.

Name	Description
RouteReflectorClient	Specifies that this peer group is a route reflector client. Note: This parameter only applies to VRF 0.
SoftReconfigurationIn	When enabled, the router relearns routes from the specified neighbor or group of neighbors without restarting the connection after the policy changes in the inbound direction. The default value is enable. Enabling SoftReconfigurationIn stores all BGP routes in local memory (even non-best routes).
MD5Authentication	Enables and disables MD5 authentication. The default is disable.
RemovePrivateAs	Strips (when enabled) private AS numbers when the switch sends an update. The default is enable.
SendCommunity	Enables or disables sending the community attribute of the update message to the specified peer group. The default value is disable.
AfUpdateSourceInterfaceType	Specifies the interface type.
AfUpdateSourceInterface	Specifies the IP address used for circuitless IP (CLIP) for this peer group.
Vpnv4Address	Enables BGP address families for IPv4 (BGP) and Layer 3 VPN (MP-BGP) support. Enable this parameter for VPN/VRF Lite routes.
IpvpnLiteCap	Specifies (when enabled) that IP VPN Lite capability can be enabled or disabled on the BGP neighbor peer. The default is disable.
RouteRefresh	Enables or disables route refresh. If enabled, a route refresh request received by a BGP speaker causes the speaker to resend all route updates in the database that are eligible for the peer that issues the request. This parameter only applies to VRF 0.
AsOverride	Specifies that the AS Override parameter can be enabled or disabled for the BGP peer group. The default is disable.
AllowedAsIn	Specifies the number of AS-in allowed for the BGP peer group.
IPv6Cap	Enable or disable the IPv6 capability on the BGP neighbor peer. The default value is disable.
Ipv6RoutePolicyIn	Specifies the policy (by name) that applies to all network IPv6 routes learned from this peer.
Ipv6RoutePolicyOut	Specifies the policy (by name) that applies to all outgoing IPv6 route updates.

Name	Description
Ipv6MaxPrefix	Configures a limit on the number of IPv6 routes accepted from a neighbor. A value of 0 means no limit exists.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for the BGP peer group.

View IPv6 Community Attributes

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

View IPv6 community attributes for specific routes to utilize the update message fields to communicate information between BGP speakers. Use the Path Attribute values to specify the prefixes that the BGP session can exchanged, or which of the multiple paths of a specified prefix to use.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **BGP+**.
3. Select the **Bgp Route Summary** tab.
4. Select a route for which you want to view the route summary information.
5. Select the **Route Comm Attr** option on the menu.

The **BGP Path Attributes** tab opens with the BGP IPv6 community attribute information.

BGP Path Attributes Field Descriptions

Use the data in the following table to use the **BGP Path Attributes** tab.

Name	Description
Origin	Specifies the ultimate origin of the path information.
NextHopAddr	Specifies the address of the border router that is used to access the destination network. This address is the nexthop address received in the UPDATE packet associated with this prefix.

Name	Description
Med	This metric is used to discriminate between multiple exit points to an adjacent autonomous system. When the MED value is absent but has a calculated default value, this object will contain the calculated value.
LocalPref	Specifies the value used during route decision process in the BGP protocol. Applicable to BGP only.
AggregatorAS	Specifies the AS number of the last BGP4 speaker that performed route aggregation. If the AGGREGATOR path attribute is absent, this object will not be present in the conceptual row.
AggregatorAddr	Specifies the IP address of the last BGP4 speaker that performed route aggregation. If the AGGREGATOR path attribute is absent, this object will not be present in the conceptual row.
String	<p>This is a string representing the autonomous system path to the network which was received from the peer which advertised it. The format of the string is implementation-dependent, and is designed for operator readability.</p> <p>Note: SnmpAdminString is only capable of representing a maximum of 255 characters. This may lead to the string being truncated in the presence of a large AS Path.</p>

Display Dampened Routes Information

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.
- Enable dampened routes.

About This Task

Display dampened path information to see which routes are suppressed.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.

3. Select the **Dampened Routes** tab.

Dampened Routes Field Descriptions

Use the data in the following table to use the **Dampened Routes** tab.

Name	Description
IpAddrPrefix	Specifies the IP address prefix in the NLRI field. This variable is an IP address that contains the prefix with a length specified by IpAddrPrefixLen. Bits beyond the length specified by IpAddrPrefixLen are set to zero.
IpAddrPrefixLen	Specifies the length, in bits, of the IP address prefix in the NLRI field.
Peer	Specifies the IP address of the peer from which the router learns the path information.
FlapPenalty	Specifies the penalty based on number of route flaps.
FlapCount	Specifies the number of times a route flapped (went down and up) since the last time the penalty was reset to zero.
RouteDampened	Indicates whether this route is suppressed or announced.
ReuseTime	Specifies the system-configured time for route reuse.

Configure Route Redistribution to BGP

Before You Begin

- If required, configure a route policy.
- When you configure BGP on a specific VRF instance, the VRF must have an RP trigger of BGP.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- Before you redistribute DvR host routes to BGP, ensure that you disable BGP aggregation and BGP auto-summarization of networks. Disabling these options ensures that all DvR host routes are advertised into BGP correctly, and are not summarized.

About This Task

Configure redistribute entries for BGP to announce routes of a certain source type to BGP, for example, DvR, direct, static, Routing Information Protocol (RIP), and Open Shortest Path First (OSPF). If you do not configure a route policy, then the switch uses the default action based on metric, metric type, and subnet. Use a route policy to perform detailed redistribution.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. Configure the source protocol.
6. (Optional) If required, choose a route policy.
7. Configure the metric to apply to redistributed routes.
8. (Optional) Select **AutoFilterTag** to enable route filtering using tags.
9. Enable the redistribution instance.
10. Select **Insert**.

Redistribute Field Descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrfId	Specifies the destination VRF instance (read-only).
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrfId	Specifies the source VRF instance.
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables (or disables) a BGP redistribute entry for a specified source type.
RoutePolicy	Configures the route policy to use for the detailed redistribution of external routes from a specified source into the BGP domain.
Metric	Configures the metric for the redistributed route. The default value is 0. Use a value that is consistent with the destination protocol.
MetricType	Configures the metric type to apply to redistributed routes. When you enable the live-metric option, when BGP redistributes static, RIP, OSPF, IS-IS, or DvR routes, the metric value is taken from the routing table and is set to the Path attributes as a MED value. By default, this option is disabled, which means the BGP MED value is not derived from the metric in the routing table.
AutoFilterTag	Enables route filtering to ensure tagged routes are not redistributed into BGP. The default is disabled.

Configure Redistribution to BGPv6

Before You Begin

- If required, configure a route policy.
- When you configure BGPv6 on a specific VRF instance, the VRF must have an RP trigger of BGP.

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Configure redistribute entries for BGPv6 to announce routes of a certain source type to BGPv6. If you do not configure a route policy, then the switch uses the default action based on metric, metric type, and subnet. Use a route policy to perform detailed redistribution.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **BGP+**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. Configure the source protocol.
6. (Optional) If required, choose a route policy.
7. Configure the metric to apply to redistributed routes.
8. Enable the redistribution instance.
9. Select **Insert**.

Redistribute Field Descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrfId	Specifies the destination VRF instance (read-only).
Protocol Note:	Specifies the protocols that receive the redistributed routes.
SrcVrfId	Specifies the source VRF instance (read-only).
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables (or disables) a BGPv6 redistribute entry for a specified source type.
RoutePolicy	Configures the route policy to use for the detailed redistribution of external routes from a specified source into the BGPv6 domain.

Name	Description
Metric	Configures the metric for the redistributed route. The default value is 0. Use a value that is consistent with the destination protocol.
MetricType	Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2.

View BGP+ or BGPv6 Route Summary Information

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

You can display current IPv6 BGP+ route information.

Procedure

- In the navigation pane, expand **Configuration > IPv6**.
- Select **BGP+**.
- Select the **Bgp Route Summary** tab to view the BGP route summary information

Bgp Route Summary Field Descriptions

Use the data in the following table to use the **Bgp Route Summary** tab.

Name	Description
Prefix	Specifies the IP address prefix in the Network Layer Reachability Information (NLRI) field. This is an IP address that contains the prefix with a length specified by IpAddrPrefixLen. Any bits beyond the length specified by IpAddrPrefixLen are set to zero.
PrefixLen	Specifies the length, in bits, of the IP address prefix in the NLRI field.

Name	Description
LocalAddr	The local address of this entry's BGP connection.
RemoteAddr	Specifies the IP address of the peer from which path information was learned.

View BGP Route Summary

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Display BGP route summary.

Procedure

- In the navigation pane, expand **Configuration > IP**.
- Select **BGP**.
- Select the **Bgp Route Summary** tab.

Bgp Route Summary Field Descriptions

Use the data in the following table to use the **Bgp Route Summary** tab.

Name	Description
Prefix	Configures the IP address of the route.
PrefixLen	Specifies the IP address and the mask length (the length can be 0–32).
LocalAddr	Specifies the local IP address of the entered BGP route.
RemoteAddr	Specifies the remote IP address of the entered BGP route.

Configure an AS Path List

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure an AS path list to restrict the routing information a router learns or advertises to and from a neighbor. The AS path list acts as a filter that matches AS paths.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Policy**.
3. Select the **As Path List** tab.
4. Select **Insert**.
5. Enter the appropriate information for your configuration.
6. Select **Insert**.

As Path List *Field Descriptions*

Use the data in the following table to use the **As Path List** tab.

Name	Description
Id	Specifies the AS path list.
MemberId	Specifies the AS path access list member ID.
Mode	Specifies the action to take if the system selects a policy for a specific route. Select permit (allow the route) or deny (ignore the route).
AsRegularExpression	Specifies the expression to use for the AS path.

Configure a Community Access List

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of BGP.

About This Task

Configure community lists to specify permitted routes by using their BGP community. This list acts as a filter that matches communities or AS numbers.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Policy**.
3. Select the **Community List** tab.
4. Select **Insert**.
5. Configure the list as required.

6. Select **Insert**.

Community List *Field Descriptions*

Use the data in the following table to use the **Community List** tab.

Name	Description
Id	Specifies the community list. The range is 0–1024.
MemberId	Specifies the community list member ID. The range is 0–65535.
Mode	Specifies the action to take if the system selects a policy for a specific route. Select permit (allow the route) or deny (ignore the route).
Community	Specifies the community access list community string.

BGP Configuration Examples

This section provides configuration examples to configure Border Gateway Protocol (BGP).

IPv6 Tunnel Configurations for BGP+

You must configure an IPv6 tunnel and static routes at BGP+ peers when you use BGP+.

When BGP+ peers advertise route information, they use Update messages to advertise route information. And, when route information is encapsulated in Update messages, BGP+ peers convert their own IPv4 peer addresses to IPv4-mapped IPv6 addresses and insert them into the next-hop field in the Update message.

When the BGP+ software module receives Update messages, it adds route information to the IPv6 Routing Manager (RTM). These RTM routes contain next-hop addresses from the BGP peer that the route was learned from. The next-hop addresses are represented as IPv4-mapped IPv6 addresses.

But, because the IPv6 RTM cannot correlate the IPv4-mapped IPv6 address to a specific outgoing interface, you must create a manually-configured static route to make the link between the BGP peer and the IPv6 tunnel interface so that traffic can reach networks advertised by the peer.

Following is one way to express a static route in an IPv6-configured tunnel for BGP+:

```
ipv6 route 0:0:0:0:0:ffff:192.0.2.0/24 cost 1 tunnel 10
```

Configure the IPv6 tunnel endpoint and the BGP peer to reside on the same switch.

If the IPv6 tunnel endpoint and the BGP peer must reside on different switches you can terminate the tunnel on a different switch, but you must consider the following:

- Because the IPv6 tunnel endpoint does not reside on the same switch as the BGP peer, the BGP device cannot use the tunnel as the outgoing interface. That is, to reach the IPv6-configured tunnel endpoint, if the BGP peer resides on a different switch from the IPv6 tunnel endpoint, the next-hop for the manually-configured IPv4-mapped IPv6 static route is the native IPv6 interface next-hop address.
- The node where the tunnel terminates must contain all of the information needed to route the packets between the remote IPv6 network clouds.

**Note**

In order for the tunnel endpoint switch to be aware of all of the necessary IPv6 routes, you may need to redistribute the BGP routes into OSPFv3.

IPv4-mapped IPv6 Addresses

IPv4-mapped IPv6 addresses are IPv4 addresses that the system has mapped into the IPv6 address space.

The system uses these IPv4-mapped IPv6 addresses for devices that are only IPv4-capable.

These IPv4-mapped address have the first 80 bits set to zeros, followed by the next 16 bits set to ones, and the last 32 bits have IPv4 addresses.

When converted to an IPv4-mapped IPv6 address, an IPv4 device address of 192.0.2.1 would be represented as one of the following:

- 0:0:0:0:FFFF:192.0.2.1
- ::FFFF:192.0.2.1

The following figure illustrates the components in an IPv4-mapped IPv6 address.

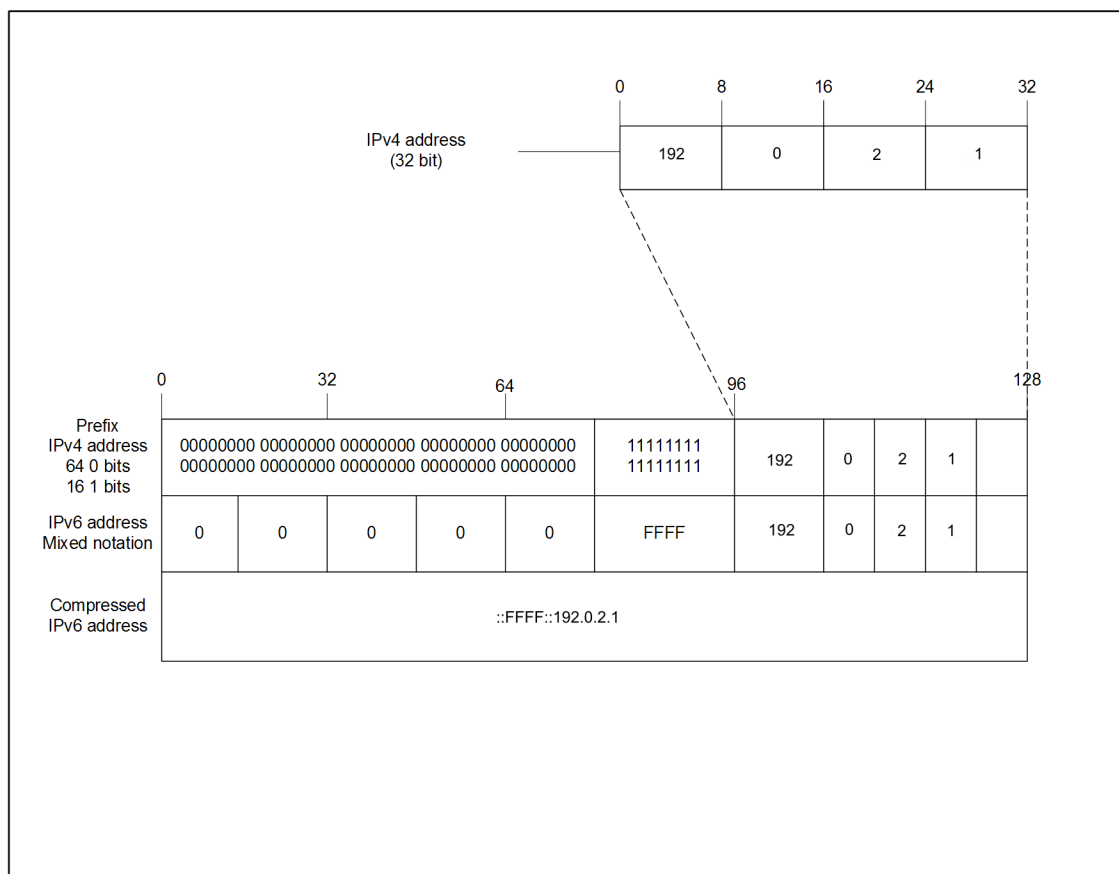


Figure 41: IPv4-mapped IPv6 address components

eBGP+ peership between two switches with IPv6 Tunneling

The following figure shows a sample network that contains eBGP+ peers using IPv6 tunneling.




Figure 42: eBGP+ peers with IPv6 tunneling

The configuration in the figure, *eBGP+ peers with IPv6 tunneling*, assumes that the BGP peer IP address is the next hop.

When you configure the static route for the BGP+ tunnel, you must designate the BGP peer IP address as the next hop in most cases.

You can configure multiple static routes, using the same tunnel, but you must ensure reachability when you create the static routes.

R1 Configuration

```
interface GigabitEthernet 1/2
brouter port 1/2 vlan 2090 subnet 192.0.2.1/255.255.255.0 mac-offset 2
exit
# BGP CONFIGURATION - GlobalRouter
#

router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "192.0.2.2"
neighbor 192.0.2.2 remote-as 65001
neighbor 192.0.2.2 address-family ipv6
neighbor 192.0.2.2 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 192.0.2.1 address 2001:DB8::/32 destination 200.1.1.2

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:0:ffff:192.0.2.1 cost 1 tunnel 10
#
```

R2 Configuration

```
interface GigabitEthernet 1/32
brouter port 1/32 vlan 2090 subnet 192.0.2.2/255.255.255.0 mac-offset 2
exit
# BGP CONFIGURATION - GlobalRouter
#

router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "192.0.2.1"
neighbor 192.0.2.1 remote-as 65000
neighbor 192.0.2.1 address-family ipv6
neighbor 192.0.2.1 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 192.0.2.2 address 2001:DB8::/32 destination
```

```

192.0.2.1
#
# IPV6 STATIC ROUTE CONFIGURATION
#
ipv6 route 0:0:0:0:0:ffff:192.0.2.1 cost 1 tunnel 10
#

```

iBGP+ peership on CLIP between two switches with IPv6 Tunneling

The following figure shows a sample network that contains iBGP+ peers using IPv6 tunneling.




Figure 43: iBGP+ peers on CLIP interfaces with IPv6 tunneling

You must enable OSPF on the interface and globally as well.

If you cannot enable OSPF, you must configure static routes to provide reachability to the BGP+ peer.

The static route must point to the next hop for the routes to be installed in the IPv6 RTM.

The next hop must be the BGP peer IP address.

The IPv4 interfaces do not need to connect directly, but the routing table on each switch must include the IPv4 interface of the other switch.

iBGP between the CLIP interfaces needs to run OSPF as a routing protocol so that the BGP neighbor can remain reachable.

eBGP connections cannot use a CLIP interface as an end point.

R1 Configuration

```

interface GigabitEthernet 1/2
brouter port 1/2 vlan 2090 subnet 192.0.2.1/255.255.255.0 mac-offset
2
exit
# OSPF CONFIGURATION - GlobalRouter
#
router ospf enable

# OSPF PORT CONFIGURATION
#

```



```

interface gigabitethernet 1/2
ip ospf enable
exit

# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#

interface loopback 1
ip address 1 1.1.1.1/255.255.255.255
ip ospf 1

# BGP CONFIGURATION - GlobalRouter
#

router bgp
no synchronization
exit
router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "2.2.2.2"
neighbor 2.2.2.2 remote-as 65000
neighbor 2.2.2.2 next-hop-self
neighbor 2.2.2.2 update-source 1.1.1.1
neighbor 2.2.2.2 address-family ipv6
neighbor 2.2.2.2 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 192.0.2.1 address 2001:DB8::/32 destination
192.0.2.2

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:0:0:ffff:2.2.2.2/128 cost 1 tunnel 10
#

```

R2 Configuration

```

interface GigabitEthernet 1/32
brouter port 1/32 vlan 2090 subnet 192.0.2.2/255.255.255.0 mac-offset
2
exit
# OSPF CONFIGURATION - GlobalRouter
#

router ospf enable

# OSPF PORT CONFIGURATION
#

interface gigabitethernet 1/32
ip ospf enable
exit

```

```
# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#

interface loopback 1
ip address 1 2.2.2.2/255.255.255.255
ip ospf 1

# BGP CONFIGURATION - GlobalRouter
#

router bgp
no synchronization
exit
router bgp as-dot enable
router bgp 65000 enable
router bgp
neighbor "1.1.1.1"
neighbor 1.1.1.1 remote-as 65000
neighbor 1.1.1.1 next-hop-self
neighbor 1.1.1.1 update-source 2.2.2.2
neighbor 1.1.1.1 address-family ipv6
neighbor 1.1.1.1 enable
exit
# IPV6 CONFIGURATION
#

ipv6 forwarding

# IPV6 TUNNEL CONFIGURATION
#

ipv6 tunnel 10 source 192.0.2.2 address 2001:DB8::/32 destination
192.0.2.1

#
# IPV6 STATIC ROUTE CONFIGURATION
#

ipv6 route 0:0:0:0:0:ffff:1.1.1.1/128 cost 1 tunnel 10
#
```

Native IPv6 eBGP Peership Between Two Switches on VRFs

The following figure shows a sample network that contains native IPv6 eBGP peers on VRFs.




Figure 44: Native IPv6 eBGP peers on VRFs

Configure the local AS first on GRT (and it is inherited by all VRFs), and then enable BGP on GRT/VRF.

You must configure the **address-family ipv6** option for IPv6 peers, otherwise, peer-ship is formed, but no routing updates between them will take place.

You must configure the **ebgp-multihop** option for the given eBGP peer that is not on one of local subnets (remote peers), otherwise, peer-ship will not be formed.



Note

The switch does not accept any configuration command for BGP in router-vrf configuration mode unless a BGP instance associated to the VRF context is created. You can use `ip bgp` command in router-vrf configuration mode to create a BGP instance on VRF.

R1 Configuration

```
#
# VRF CONFIGURATION
#

ip vrf vrf1 vrfid 1
router vrf vrf1
exit
ip vrf vrf2 vrfid 2
router vrf vrf2
exit
#

# PORT CONFIGURATION - PHASE I
#

interface GigabitEthernet 1/1
encapsulation dot1q
exit

#
# VLAN CONFIGURATION
#

vlan members remove 1 1/1,1/46
vlan create 100 type port-mstprstp 0
vlan members 100 1/1 portmember
interface Vlan 100
ip address 100.1.1.1 255.255.255.0 1
ipv6 interface mac-offset 1
ipv6 interface enable
ipv6 interface address 2001:0:100:0:0:0:0:1/64
exit
vlan create 101 type port-mstprstp 0
vlan members 101 1/1 portmember
interface Vlan 101
vrf vrf1
ip address 101.1.1.1 255.255.255.0 2
ipv6 interface mac-offset 2
ipv6 interface enable
ipv6 interface address 2001:0:101:0:0:0:0:1/64
exit
vlan create 102 type port-mstprstp 0
vlan members 102 1/1 portmember
interface Vlan 102
vrf vrf2
ip address 102.1.1.1 255.255.255.0 3
```

```
ipv6 interface mac-offset 3
ipv6 interface enable
ipv6 interface address 2001:0:102:0:0:0:0:1/64
ipv6 forwarding

exit

#
# PORT CONFIGURATION - PHASE II
#

interface GigabitEthernet 1/1
default-vlan-id 100
no shutdown
exit

#
# CIRCUITLESS IPV6 INTERFACE CONFIGURATION - GlobalRouter
#

interface loopback 1
ipv6 interface address 1:1:1:1:0:0:0:1/128

exit          #
# CIRCUITLESS IPV6 INTERFACE CONFIGURATION - VRF
#

interface loopback 2
ipv6 interface address 11:1:1:1:0:0:0:1/128 vrf vrf1
exit
interface loopback 3
ipv6 interface address 12:1:1:1:0:0:0:1/128 vrf vrf2
exit

#
# BGP CONFIGURATION - GlobalRouter
#

router bgp
no synchronization
exit
router bgp 1000 enable
router bgp
network 1:1:1:1:0:0:0:1/128 metric 100000
neighbor "2001:0:100:0:0:0:0:2"
neighbor 2001:0:100:0:0:0:0:2 remote-as 10000
neighbor 2001:0:100:0:0:0:0:2 next-hop-self
neighbor 2001:0:100:0:0:0:0:2 ebgp-multihop
neighbor 2001:0:100:0:0:0:0:2 address-family ipv6
neighbor 2001:0:100:0:0:0:0:2 update-source 2001:0:100:0:0:0:0:1
neighbor 2001:0:100:0:0:0:0:2 enable
exit#
# BGP CONFIGURATION - VRF
#

router vrf vrf1
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp network 11:1:1:1:0:0:0:1/128 metric 100000
ip bgp neighbor "2001:0:101:0:0:0:0:2"
ip bgp neighbor 2001:0:101:0:0:0:0:2 remote-as 10000
ip bgp neighbor 2001:0:101:0:0:0:0:2 next-hop-self
ip bgp neighbor 2001:0:101:0:0:0:0:2 ebgp-multihop
```

```

ip bgp neighbor 2001:0:101:0:0:0:2 address-family ipv6
ip bgp neighbor 2001:0:101:0:0:0:2 update-source 2001:0:101:0:0:0:1
ip bgp neighbor 2001:0:101:0:0:0:2 enable
exit
router vrf vrf2
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp network 12:1:1:1:0:0:1/128 metric 100000
ip bgp neighbor "2001:0:102:0:0:0:2"
ip bgp neighbor 2001:0:102:0:0:0:2 remote-as 10000
ip bgp neighbor 2001:0:102:0:0:0:2 next-hop-self
ip bgp neighbor 2001:0:102:0:0:0:2 ebgp-multihop
ip bgp neighbor 2001:0:101:0:0:0:2 address-family ipv6
ip bgp neighbor 2001:0:102:0:0:0:2 update-source 2001:0:102:0:0:0:1
ip bgp neighbor 2001:0:102:0:0:0:2 enable
exit

```

R2 Configuration

```

#
# VRF CONFIGURATION
#

ip vrf vrf1 vrfid 1
router vrf vrf1
exit
ip vrf vrf2 vrfid 2
router vrf vrf2
exit

#
# PORT CONFIGURATION - PHASE I
#

interface GigabitEthernet 1/1
encapsulation dot1q

exit

#
# VLAN CONFIGURATION
#

vlan members remove 1 1/1
vlan create 100 type port-mstprstp 0
vlan members 100 1/1 portmember
interface Vlan 100

ip address 100.1.1.2 255.255.255.0 1
ipv6 interface mac-offset 1
ipv6 interface enable
ipv6 interface address 2001:0:100:0:0:0:2/64
ipv6 forwarding

exit
vlan create 101 type port-mstprstp 0
vlan members 101 1/1 portmember
interface Vlan 101
vrf vrf1
ip address 101.1.1.2 255.255.255.0 2
ipv6 interface mac-offset 2
ipv6 interface enable
ipv6 interface address 2001:0:101:0:0:0:2/64

```

```
ipv6 forwarding

exit
vlan create 102 type port-mstprstp 0
vlan members 102 1/1 portmember
interface Vlan 102
vrf vrf2
ip address 102.1.1.2 255.255.255.0 3

ipv6 interface mac-offset 3
ipv6 interface enable
ipv6 interface address 2001:0:102:0:0:0:0:2/64
ipv6 forwarding

exit

#
# PORT CONFIGURATION - PHASE II
#

interface GigabitEthernet 1/1
default-vlan-id 100
no shutdownexit

#
# CIRCUITLESS IPV6 INTERFACE CONFIGURATION - GlobalRouter
#

interface loopback 1
ipv6 interface address 2:2:2:2:0:0:0:2/128
exit

#
# CIRCUITLESS IPV6 INTERFACE CONFIGURATION - VRF
#

interface loopback 2
ipv6 interface address 21:2:2:2:0:0:0:2/128 vrf vrf1
exit
interface loopback 3
ipv6 interface address 22:2:2:2:0:0:0:2/128 vrf vrf2
exit

#
# BGP CONFIGURATION - GlobalRouter
#

router bgp
no synchronization
exit
router bgp 10000 enable
router bgp
neighbor "2001:0:100:0:0:0:0:1"
neighbor 2001:0:100:0:0:0:0:1 remote-as 1000
neighbor 2001:0:100:0:0:0:0:1 next-hop-self
neighbor 2001:0:100:0:0:0:0:1 ebgp-multihop
neighbor 2001:0:100:0:0:0:0:1 address-family ipv6
neighbor 2001:0:100:0:0:0:0:1 update-source 2001:0:100:0:0:0:0:2
neighbor 2001:0:100:0:0:0:0:1 enableexit

#
# BGP CONFIGURATION - VRF
```

```
#
router vrf vrf1
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp neighbor "2001:0:101:0:0:0:0:1"
ip bgp neighbor 2001:0:101:0:0:0:0:1 remote-as 1000
ip bgp neighbor 2001:0:101:0:0:0:0:1 next-hop-self
ip bgp neighbor 2001:0:101:0:0:0:0:1 ebgp-multihop
ip bgp neighbor 2001:0:101:0:0:0:0:1 address-family ipv6
ip bgp neighbor 2001:0:101:0:0:0:0:1 update-source 2001:0:101:0:0:0:2
ip bgp neighbor 2001:0:101:0:0:0:0:1 enable
exit
router vrf vrf2
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp neighbor "2001:0:102:0:0:0:0:1"
ip bgp neighbor 2001:0:102:0:0:0:0:1 remote-as 1000
ip bgp neighbor 2001:0:102:0:0:0:0:1 next-hop-self
ip bgp neighbor 2001:0:102:0:0:0:0:1 ebgp-multihop
ip bgp neighbor 2001:0:102:0:0:0:0:1 address-family ipv6
ip bgp neighbor 2001:0:102:0:0:0:0:1 update-source 2001:0:102:0:0:0:2
ip bgp neighbor 2001:0:102:0:0:0:0:1 enable
exit
```

iBGP over User-created VRFs Configuration Example

This section shows examples of configured internal Border Gateway Protocol (iBGP) IPv4 and IPv6 peers over user-created Virtual Routing and Forwarding (VRF) instances.



Note

The Autonomous System (AS) number configured on the global VRF is inherited by all user-created VRFs, however, you can override the AS number for a specific user-created VRF. For more information, see [Configure an AS Number for a Non-default VRF](#) on page 601.

IPv4 iBGP Peers Configuration

Configuration on switch 1:

```
#
# VRF CONFIGURATION
#
ip vrf vrf1 vrfid 1
router vrf vrf1
exit
ip vrf vrf2 vrfid 2
router vrf vrf2
exit
#
# PORT CONFIGURATION - PHASE I
#
interface GigabitEthernet 1/1
encapsulation dot1q
exit
#
# VLAN CONFIGURATION
#
```

```
vlan members remove 1 1/1,1/46
vlan create 100 type port-mstprstp 0
vlan members 100 1/1 portmember
interface Vlan 100
ip address 10.10.10.1 255.255.255.0 1
exit
vlan create 101 type port-mstprstp 0
vlan members 101 1/1 portmember
interface Vlan 101
vrf vrf1
ip address 11.10.10.1 255.255.255.0 2
exit
vlan create 102 type port-mstprstp 0
vlan members 102 1/1 portmember
interface Vlan 102
vrf vrf2
ip address 12.10.10.1 255.255.255.0 3
exit
#
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 100
no shutdown
exit
#
# CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#
interface loopback 1
ip address 10.1.1.10/32
exit
#
# CIRCUITLESS IP INTERFACE CONFIGURATION - VRF
#
interface loopback 2
ip address 11.1.1.11/32 vrf vrf1
exit
interface loopback 3
ip address 12.1.1.12/32 vrf vrf2
exit
#
# BGP CONFIGURATION - GlobalRouter
#
router bgp
no synchronization
exit
router bgp 1000 enable
router bgp
network 10.1.1.10/32 metric 100000
neighbor "10.10.10.2"
neighbor 10.10.10.2 remote-as 1000
neighbor 10.10.10.2 next-hop-self
neighbor 10.10.10.2 update-source 10.10.10.1
neighbor 10.10.10.2 enable
exit
#
# BGP CONFIGURATION - VRF
#
router vrf vrf1
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp network 11.1.1.11/32 metric 100000
ip bgp neighbor "11.10.10.2"
```



```
ip bgp neighbor 11.10.10.2 remote-as 1000
ip bgp neighbor 11.10.10.2 next-hop-self
ip bgp neighbor 11.10.10.2 update-source 11.10.10.1
ip bgp neighbor 11.10.10.2 enable
exit
router vrf vrf2
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp network 12.1.1.12/32 metric 100000
ip bgp neighbor "12.10.10.2"
ip bgp neighbor 12.10.10.2 remote-as 1000
ip bgp neighbor 12.10.10.2 next-hop-self
ip bgp neighbor 12.10.10.2 update-source 12.10.10.1
ip bgp neighbor 12.10.10.2 enable
exit
```

Configuration on switch 2:

```
#
# VRF CONFIGURATION
#
ip vrf vrf1 vrfid 1
router vrf vrf1
exit
ip vrf vrf2 vrfid 2
router vrf vrf2
exit
#
# PORT CONFIGURATION - PHASE I
#
interface GigabitEthernet 1/1
encapsulation dot1q
exit
#
# VLAN CONFIGURATION
#
vlan members remove 1 1/1,1/46
vlan create 100 type port-mstprstp 0
vlan members 100 1/1 portmember
interface Vlan 100
ip address 10.10.10.2 255.255.255.0 1
exit
vlan create 101 type port-mstprstp 0
vlan members 101 1/1 portmember
interface Vlan 101
vrf vrf1
ip address 11.10.10.2 255.255.255.0 2
exit
vlan create 102 type port-mstprstp 0
vlan members 102 1/1 portmember
interface Vlan 102
vrf vrf2
ip address 12.10.10.2 255.255.255.0 3
exit
#
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 100
no shutdown
exit
#
# BGP CONFIGURATION - GlobalRouter
```

```
#
router bgp
no synchronization
exit
router bgp 1000 enable
router bgp
neighbor "10.10.10.1"
neighbor 10.10.10.1 remote-as 1000
neighbor 10.10.10.1 next-hop-self
neighbor 10.10.10.1 update-source 10.10.10.2
neighbor 10.10.10.1 enable
exit
#
# BGP CONFIGURATION - VRF
#
router vrf vrf1
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp neighbor "11.10.10.1"
ip bgp neighbor 11.10.10.1 remote-as 1000
ip bgp neighbor 11.10.10.1 next-hop-self
ip bgp neighbor 11.10.10.1 update-source 11.10.10.2
ip bgp neighbor 11.10.10.1 enable
exit
router vrf vrf2
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp neighbor "12.10.10.1"
ip bgp neighbor 12.10.10.1 remote-as 1000
ip bgp neighbor 12.10.10.1 next-hop-self
ip bgp neighbor 12.10.10.1 update-source 12.10.10.2
ip bgp neighbor 12.10.10.1 enable
exit
```

IPv6 iBGP Peers Configuration

Configuration on switch 1:

```
#
# VRF CONFIGURATION
#
ip vrf vrf1 vrfid 1
router vrf vrf1
exit
ip vrf vrf2 vrfid 2
router vrf vrf2
exit
#
# PORT CONFIGURATION - PHASE I
#
interface GigabitEthernet 1/1
encapsulation dot1q
exit
#
# VLAN CONFIGURATION
#
vlan members remove 1 1/1,1/46
vlan create 100 type port-mstprstp 0
vlan members 100 1/1 portmember
interface Vlan 100
ipv6 interface mac-offset 1
ipv6 interface enable
```

```
ipv6 interface address 2001:DB8:0::1/64
exit
vlan create 101 type port-mstprstp 0
vlan members 101 1/1 portmember
interface Vlan 101
vrf vrf1
ipv6 interface mac-offset 2
ipv6 interface enable
ipv6 interface address 2001:DB8:1::1/64
exit
vlan create 102 type port-mstprstp 0
vlan members 102 1/1 portmember
interface Vlan 102
vrf vrf2
ipv6 interface mac-offset 3
ipv6 interface enable
ipv6 interface address 2001:DB8:2::1/64
exit
#
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 100
no shutdown
exit
#
# CIRCUITLESS IPV6 INTERFACE CONFIGURATION - GlobalRouter
#
interface loopback 1
ipv6 interface address 2001:DB8:2000::1/128
exit
#
# CIRCUITLESS IPV6 INTERFACE CONFIGURATION - VRF
#
interface loopback 2
ipv6 interface address 2001:DB8:2001::1/128 vrf vrf1
exit
interface loopback 3
ipv6 interface address 2001:DB8:2002::1/128 vrf vrf2
exit
#
# BGP CONFIGURATION - GlobalRouter
#
router bgp
no synchronization
exit
router bgp 1000 enable
router bgp
network 2001:DB8:2000::1/128 metric 100000
neighbor "2001:DB8:0::2"
neighbor 2001:DB8:0::2 remote-as 1000
neighbor 2001:DB8:0::2 next-hop-self
neighbor 2001:DB8:0::2 address-family ipv6
neighbor 2001:DB8:0::2 update-source 2001:DB8:0::1
neighbor 2001:DB8:0::2 enable
exit
#
# BGP CONFIGURATION - VRF
#
router vrf vrf1
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp network 2001:DB8:2001::1/128 metric 100000
```

```

ip bgp neighbor "2001:DB8:1::2"
ip bgp neighbor 2001:DB8:1::2 remote-as 1000
ip bgp neighbor 2001:DB8:1::2 next-hop-self
ip bgp neighbor 2001:DB8:1::2 update-source 2001:DB8:1::1
ip bgp neighbor 2001:DB8:1::2 address-family ipv6
ip bgp neighbor 2001:DB8:1::2 enable
exit
router vrf vrf2
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp network 2001:DB8:2002::1/128 metric 100000
ip bgp neighbor "2001:DB8:2::2"
ip bgp neighbor 2001:DB8:2::2 remote-as 1000
ip bgp neighbor 2001:DB8:2::2 next-hop-self
ip bgp neighbor 2001:DB8:2::2 update-source 2001:DB8:2::1
ip bgp neighbor 2001:DB8:2::2 address-family ipv6
ip bgp neighbor 2001:DB8:2::2 enable
exit

```

Configuration on switch 2:

```

#
# VRF CONFIGURATION
#
ip vrf vrf1 vrfid 1
router vrf vrf1
exit
ip vrf vrf2 vrfid 2
router vrf vrf2
exit
#
# PORT CONFIGURATION - PHASE I
#
interface GigabitEthernet 1/1
encapsulation dot1q
exit
#
# VLAN CONFIGURATION
#
vlan members remove 1 1/1,1/46
vlan create 100 type port-mstprstp 0
vlan members 100 1/1 portmember
interface Vlan 100
ipv6 interface mac-offset 1
ipv6 interface enable
ipv6 interface address 2001:DB8:0::2/64
exit
vlan create 101 type port-mstprstp 0
vlan members 101 1/1 portmember
interface Vlan 101
vrf vrf1
ipv6 interface mac-offset 2
ipv6 interface enable
ipv6 interface address 2001:DB8:1::2/64
exit
vlan create 102 type port-mstprstp 0
vlan members 102 1/1 portmember
interface Vlan 102
vrf vrf2
ipv6 interface mac-offset 3
ipv6 interface enable
ipv6 interface address 2001:DB8:2::2/64
exit

```

```
#
# PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 100
no shutdown
exit
#
# BGP CONFIGURATION - GlobalRouter
#
router bgp
no synchronization
exit
router bgp 1000 enable
router bgp
neighbor "2001:DB8:0::1"
neighbor 2001:DB8:0::1 remote-as 1000
neighbor 2001:DB8:0::1 next-hop-self
neighbor 2001:DB8:0::1 address-family ipv6
neighbor 2001:DB8:0::1 update-source 2001:DB8:0::2
neighbor 2001:DB8:0::1 enable
exit
#
# BGP CONFIGURATION - VRF
#
router vrf vrf1
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp neighbor "2001:DB8:1::1"
ip bgp neighbor 2001:DB8:1::1 remote-as 1000
ip bgp neighbor 2001:DB8:1::1 next-hop-self
ip bgp neighbor 2001:DB8:1::1 update-source 2001:DB8:1::2
ip bgp neighbor 2001:DB8:1::1 address-family ipv6
ip bgp neighbor 2001:DB8:1::1 enable
exit
router vrf vrf2
ip bgp
no ip bgp synchronization
ip bgp enable
ip bgp neighbor "2001:DB8:2::1"
ip bgp neighbor 2001:DB8:2::1 remote-as 1000
ip bgp neighbor 2001:DB8:2::1 next-hop-self
ip bgp neighbor 2001:DB8:2::1 update-source 2001:DB8:2::2
ip bgp neighbor 2001:DB8:2::1 address-family ipv6
ip bgp neighbor 2001:DB8:2::1 enable
exit
```



Chassis Operations

[Chassis Operations Fundamentals](#) on page 672

[Chassis Operations Configuration Using the CLI](#) on page 714

[Chassis operations configuration using EDM](#) on page 745

The following sections provide information for chassis operations such as hardware and software compatibility.

Chassis Operations Fundamentals

This section provides conceptual information for chassis operations such as hardware and software compatibility and power management. Read this section before you configure the chassis operations.

Entity MIB – Physical Table

Table 65: Entity MIB product support

Feature	Product	Release introduced
Entity MIB - Physical Table	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 65: Entity MIB product support (continued)

Feature	Product	Release introduced
Entity MIB enhancements and integration for the following: <ul style="list-style-type: none"> Physical Table Alias Mapping Table Physical Contains Table Last Change Time object 	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Entity MIB - Logical Table includes the following logical interface entities: <ul style="list-style-type: none"> VLANs MLTs Circuitless IP (CLIP) Fabric Extend interfaces (logical layer 2 and layer 3 IS-IS interfaces) 	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.4
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.4
	VSP 7400 Series	VOSS 8.4

The Entity MIB – Physical Table assists in the discovery of functional components on the switch. The Entity MIB – Physical Table supports a physical interface table that includes information about the chassis, power supply, fan, I/O cards, console, and management port.

Some hardware platforms support removable interface modules while others offer a fixed configuration. The names used for these modules can vary depending on the hardware platform.

The following table identifies the entity index range for the switch components.

Table 66: Entity index range for switch components

Component	Entity index range
Chassis	1
Power supply slot	3 to 8
Fan tray and fan slot	9 to 16

Table 66: Entity index range for switch components (continued)

Component	Entity index range
I/O slot	17 to 30
SF Slot	31 to 36
I/O card or module	37 to 50
SF Card	51 to 56
Console port	57
Console port 2	58
Management port	64
Management port 2	65
Power supply	68 to 73
Fan tray	74 to 81
Fan module	82 to 105
Port	192 to 1023
Pluggable Module and Sensor	19201 to 102314

For more information about Entity MIB – Physical Table, see [View Physical Entities](#) on page 749.

Software Lock-up Detection

The software lock-up detect feature monitors processes on the CPU to limit situations where the device stops functioning because of a software process issue. Monitored issues include

- software that enters a dead-lock state
- a software process that enters an infinite loop

The software lock-up detect feature monitors processes to ensure that the software functions within expected time limit.

The CPU logs detail about suspended tasks in the log file.

Port Lock feature

You can use the Port Lock feature to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify locked ports until the ports are first unlocked.

Jumbo Frames

Jumbo packets and large packets are particularly useful in server and storage over Ethernet applications. If the payload to header relation increases in a packet, the bandwidth can be used more efficiently. You can increase the Ethernet frame size

to the maximum frame size your switch supports to transmit large amounts of data efficiently and minimize the task load on a server CPU.

Tagged VLAN Support

A port with VLAN tagging activated can send tagged frames. If you plan to use Jumbo frames in a VLAN, ensure that you configure the ports in the VLAN to accept Jumbo frames and that the server or hosts in the VLAN do not send frames that exceed the maximum transmission unit (MTU) size for your switch. For more information about how to configure VLANs, see [VLAN Configuration](#) on page 3981.

Jumbo Frames with Extreme Integrated Application Hosting (IAH)

The ixgbevf Ethernet device driver within the TPVM does not correctly handle the interface MTU setting. Specifically, if you configure the interface in SR-IOV mode, packets larger than the MTU size are allowed.

To avoid this problem, configure the desired MTU size on both the relevant front-panel port and IAH port from the NOS CLI.

Multi-speed Ports

If a port supports multiple speeds, the software configures the speed automatically based on the optic type it detects in the port; you do not need to configure the port speed. For multi-speed copper ports, Auto-Negotiation detects the speed.



Note

Some VIMs must operate with all ports at the same speed, or with a group of ports at the same speed, while others can operate with ports at different speeds. For more information, see *Fabric Engine Release Notes*. The **sys vim-speed** command is supported only on VIMs that must operate with all ports at the same speed. An error message displays if you run the command on an unsupported VIM.

Port Type Compatibility

The following table summarizes the designed speeds and additional supported speeds for common Ethernet port types. It highlights which ports support multiple speeds and channelization to help guide decisions based on hardware capabilities.

Port type	Designed speed	Additional supported speed
SFP	1 Gbps	100 Mbps
SFP+	10 Gbps	<ul style="list-style-type: none"> 1 Gbps 100 Mbps
SFP28	25 Gbps	<ul style="list-style-type: none"> 10 Gbps 1 Gbps
QSFP+	40 Gbps	4x10 Gbps with channelization

Port type	Designed speed	Additional supported speed
QSFP28 SFP56-DD	100 Gbps	<ul style="list-style-type: none"> 40 Gbps 4x25 Gbps with channelization 4x10 Gbps with channelization <p>Note: A 100 Gbps DAC in a 100 Gbps port can negotiate down to 40 Gbps depending on the hardware and peer connection.</p>
QSFP56-DD	400 Gbps	<ul style="list-style-type: none"> 4x100 Gbps with channelization 100 Gbps 40 Gbps <ul style="list-style-type: none"> 4x25 Gbps with channelization 4x10 Gbps with channelization

Port Form Factor and Compatibility

The following list details port form factor similarities and guidance to determine multi-speed or channelization support:

- SFP28, SFP+, and SFP ports share the same physical size.
- QSFP28, QSFP+, SFP56-DD, and QSFP56-DD ports share the same physical size.

To determine if a port supports multiple speeds or channelization, refer to the applicable hardware documentation.

7830 Series Port Speed Restrictions

In the 7830-32CE-8DE port range 1/1-1/32, a group of four consecutive ports must operate at the same speed from either of the following groups:

- 10G or 40G or 4x10G
- 25G or 100G or 4x25G

For example, using the first speed group, ports 1/1-1/4 can all operate as 10G or 40G but not a mix of 10G and 40G. Because of channelization restrictions, ports 1/1 and 1/3 can be channelized as 4x10G. For more information, see [Channelization](#) on page 704.

In the 7830-VIM-24CE, a group of four consecutive ports must operate at the same speed from either of the following groups:

- 10G
- 25G or 100G

For testing purposes only, you can mix transceivers with different capabilities in the same port group. When you install transceivers with different speeds in the same port group, group speed depends on the auto-speed configuration. For more information, see [Auto-speed](#) on page 677.

25/10G Dual Speed Optics Behavior

The following lists the default behavior for 25/10G dual speed optics when inserted in a port:

Table 67: Default behavior for 25/10G dual speed optics

Port type	Default speed	Auto-negotiation	FEC
SPF28	25G	Disabled	Enabled
SFP+	10G	Disabled	Disabled
QSFP28 with QSA	25G	Disabled	Enabled
QSFP+ with QSA	10G	Disabled	Disabled

10/1G Dual Speed Optics Behavior

The following lists the default behavior for 10/1G dual optics when inserted in a port:

Table 68: Default behavior for 10/1G dual optics

Port type	Default speed	Auto-negotiation	FEC
SPF28	10G	Enabled	Disabled
SFP+	10G	Enabled	Disabled
SFP	1G	Enabled	Disabled
QSFP28 with QSA	10G	Enabled	Disabled
QSFP+ with QSA	10G	Enable	

Auto-speed



Note

Auto-speed applies to 7830 Series only.

For testing purposes only, you can mix transceivers with different capabilities in the same port group. When you install transceivers with different speeds in the same port group, group speed depends on the auto-speed configuration.

The following is an example output of a log message that displays the port and the group it belongs to when you insert a transceiver in the 7830 Series:

```
GlobalRouter HAL WARNING Mismatched pluggable inserted in port 1/7. Group (1/5-1/8) speed
is set to 100Gb
GlobalRouter HAL INFO Group 2 (1/5-1/8) speed set to 100Gb as first detected pluggable is
in port 1/5
GlobalRouter HAL WARNING Mismatched pluggable inserted in port 1/7. Group (1/5-1/8) speed
is set to 100Gb
```

```
GlobalRouter HAL WARNING Mismatched pluggable inserted in port 1/8. Group (1/5-1/8) speed
is set to 100Gb
GlobalRouter HAL INFO Group 3 (1/9-1/12) speed set to 40Gb as first detected pluggable is
in port 1/10
GlobalRouter HAL INFO Group 4 (1/13-1/16) speed set to 100Gb as first detected pluggable
is in port 1/15
```

The following table outlines the behavior of group port speed configuration.

Auto-speed status	Insert time behavior	Boot time behavior	Use Case
On (default)	Group speed is determined by the type of the first inserted transceiver.	Group speed is determined by the type of the first detected transceiver (lowest port).	Simplifies operation by automatically configuring the group speed based on transceiver.
Off	You must manually configure the speed if you insert a transceiver with a speed different than 100G.	Ports operate based on the manually set speed.	Preserves manually configured speeds across reboots, even if you swap transceivers.



Note

The default is 100 Gbps in the port range 1/1-1/32. For information about group port speeds supported on 7830 Series, see [Multi-speed Ports](#) on page 675.

To configure auto-speed on 7830 Series, see [Disable Auto-Speed](#) on page 731.

Auto-Negotiation

Table 69: Auto-Negotiation product support

Feature	Product	Release introduced
Auto-Negotiation	4220 Series	Fabric Engine 9.2 Not supported on 4220-8X
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4 SFP-DD ports are only available if you disable Advanced Feature Bandwidth Reservation.
	5520 Series	VOSS 8.2.5 Not supported on 5520-24X or 5520-VIM-4X QSFP28 ports are only available if you disable Advanced Feature Bandwidth Reservation. 5520-48SE and 5520-VIM-4XE at 1 Gbps only 5520-VIM-4YE at 25 Gbps only
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1 Not supported on: <ul style="list-style-type: none"> • VIM5-2Q • VIM5-4X VSP4900-12MXU-12XE - all fixed ports, with ports 13 to 24 at 1 Gbps only VSP4900-24S and VSP4900-48P - all fixed ports VSP4900-24XE - all fixed ports at 1 Gbps only VIM5-4XE at 1 Gbps only VIM5-4YE at 25 Gbps only
	VSP 7400 Series	VOSS 8.0

The Auto-Negotiation feature enables the device to switch between the various operational modes in an ordered fashion and lets you select a specific operational mode. The Auto-Negotiation feature also provides a parallel detection (called

autosensing) function to recognize compatible devices, even if they do not support Auto-Negotiation and helps the device sense the link speed only; not the duplex mode.

You can use the **show interfaces gigabitEthernet 11-config** command to see the Auto-Negotiation operational state on a port. The operational state uses the configuration and transceiver type present in the port. If you enable Auto-Negotiation for the port but the transceiver type does not support Auto-Negotiation, the operational state is disabled (false).



Important

The software requires the same Auto-Negotiation configuration on link partners to avoid incorrect declaration of link status. Mismatched configuration can cause the links to stay down as well as unpredictable behavior. Ensure the Auto-Negotiation configuration between local ports and their remote link partners match before upgrading software releases.

Default Auto-Negotiation Behavior

The default Auto-Negotiation behavior depends on the switch model, port type, and transceiver type. The following tables provide the default combinations.



Note

Although the switch disables Auto-Negotiation by default if you insert a 1 Gbps SFP transceiver in an SFP+ or SFP28 port, depending on the platform, you can enable 1 Gbps Auto-Negotiation. For more information, see [1Gb SFP Auto-Negotiation in SFP+ and SFP28 Ports](#) on page 697.

Table 70: 4220 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
10MB/100MB/ 1000BASE-T Copper	4220-12P-4X 4220-12T-4X 4220-24P-4X 4220-24T-4X 4220-4MW-8P-4X 4220-4MW-20P-4X		Enabled	
10G Ports SFP+	4220-12P-4X 4220-12T-4X 4220-24P-4X 4220-24T-4X 4220-4MW-8P-4X 4220-4MW-20P-4X	1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP		
		10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		10G DAC		
		10G SFP+		
1G/2.5G/5G BASE- T Copper	4220-4MW-8P-4X 4220-4MW-20P-4X		Enabled. The operational state is always enabled regardless of configuration.	

Table 71: 5320 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
10/100/1000Base-T	5320-16P-2MXT-2X 5320-24P-8XE 5320-24T-8XE 5320-24T-4X-XT 5320-24T-24S-4XE-XT 5320-48P-8XE 5320-48T-8XE 5320-16P-4XE		Enabled	
SFP	5320-24T-24S-4XE-XT	100M-FX		
		1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP		
10G Ports SFP+	5320-16P-2MXT-2X 5320-24T-4X-XT 5320-24T-24S-4XE-XT	1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP		
		10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		10G DAC		
		10G SFP+		

Table 71: 5320 Series default behavior (continued)

Port type	Model	Transceiver type	Auto-Negotiation	FEC
First 4 XE SFP+ Note: By default, the first 3 XE ports on these 24- and 48-port models are not available for front panel use because they are reserved for Advanced Feature Bandwidth Reservation.	5320-24P-8XE 5320-24T-8XE 5320-48P-8XE 5320-48T-8XE 5320-16P-4XE	100M	No 100M support	
		1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP	Not Supported	
		10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		10G DAC	Not Supported	
		10G SFP+		
Last 4 XE SFP+	5320-24P-8XE 5320-24T-8XE 5320-48P-8XE 5320-48T-8XE	100M	100FX part number 10063 only	
		1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP	Disabled	
		10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		10G DAC	Not Supported	
		10G SFP+		

Table 72: 5420 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
10/100/1000Base-T	5420F-8W-16P-4XE 5420F-16MW-32P-4XE 5420F-16W-32P-4XE 5420F-24P-4XE 5420F-24T-4XE 5420F-48P-4XE 5420F-48P-4XL 5420F-48T-4XE 5420M-16MW-32P-4YE 5420M-24T-4YE 5420M-24W-4YE 5420M-24W-24S-4YE 5420M-48T-4YE 5420M-48W-4YE		Enabled	
100M/1/2.5Gbps Copper	5420F-16MW-32P-4XE 5420M-16MW-32P-4YE		Enabled	
SFP	5420F-24S-4XE	100BASE-FX		
		1GBASE-T	Enabled	
		1G SFP	Enabled	
SFP+	5420F-8W-16P-4XE 5420F-16MW-32P-4XE 5420F-16W-32P-4XE 5420F-24P-4XE 5420F-24S-4XE 5420F-24T-4XE 5420F-48P-4XE 5420F-48P-4XL 5420F-48T-4XE	1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP	Disabled	
		10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G SFP+		

Table 72: 5420 Series default behavior (continued)

Port type	Model	Transceiver type	Auto-Negotiation	FEC
SFP28	5420M-16MW-32P-4YE 5420M-24T-4YE 5420M-24W-4YE 5420M-48T-4YE 5420M-48W-4YE	1GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		1G SFP	Disabled	
		10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G SFP+		
		25G DAC	Enabled	FEC is negotiated requesting CL108
		25G AOC		FEC auto:CL108
		25G SFP28		FEC auto:CL108

Table 72: 5420 Series default behavior (continued)

Port type	Model	Transceiver type	Auto-Negotiation	FEC
SFP-DD	5420F-24T-4XE	1GBASE-T	Enabled.	
	5420F-24P-4XE		The operational state is always enabled regardless of configuration.	
	5420F-48T-4XE			
	5420F-48P-4XE			
	5420F-48P-4XL			
	5420F-8W-16P-4XE			
	5420F-16W-32P-4XE	1G SFP	Disabled	
	5420F-24S-4XE			
	5420M-16MW-32P-4YE	10GBASE-T	Enabled.	
	5420M-24W-4YE		The operational state is always enabled regardless of configuration.	
	5420M-24W-24S-4YE			
	5420M-48W-4YE			
	5420M-24T-4YE			
	5420M-48T-4YE			
	5420M-16MW-32P-4YE	10G DAC	Disabled	
		10G SFP+		

Table 73: 5520 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
1GBASE-T	5520-24T 5520-24W 5520-48T 5520-48W		Enabled	
100M/1/2.5/5 Gbps Copper	5520-12MW-36W		Enabled	
SFP	5520-48SE	100BASE-FX		
		1GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		1G SFP	Enabled	

Table 73: 5520 Series default behavior (continued)

Port type	Model	Transceiver type	Auto-Negotiation	FEC
SFP+	5520-24X 5520-VIM-4X and 5520-VIM-4XE	1GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		1G SFP	Disabled	
		10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G SFP+		
SFP28	5520-VIM-4YE	10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G SFP+		
		25G DAC	Enabled	FEC is negotiated requesting CL108.
		25G AOC		FEC auto: CL108
		25G SFP28		FEC auto: CL108

Table 74: 5720 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
100 M/1/2.5/5 Gbps Copper	5720-24MW 5720-48MW		Enabled	
100 M/1/2.5/5/10 Gbps Copper	5720-24MXW 5720-48MXW		Enabled	

Table 74: 5720 Series default behavior (continued)

Port type	Model	Transceiver type	Auto-Negotiation	FEC
QSFP28	5720-24MW 5720-48MW 5720-24MXW 5720-48MXW	40G DAC	Enabled	
		40G AOC		
		40G QSFP+		
		100G DAC	Enabled	FEC is negotiated requesting CL91
		100G AOC		FEC auto: CL91
		100G QSFP28		FEC auto: CL91
SFP28	5720-VIM-6YE	1GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		1G SFP	Disabled	
		10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G AOC		
		10G SFP+		
		25G DAC	Enabled	FEC is negotiated requesting CL108
		25G AOC		FEC auto:CL108
		25 SFP28		FEC auto:CL108

Table 74: 5720 Series default behavior (continued)

Port type		Model	Transceiver type	Auto-Negotiation	FEC
QSFP28	with QSFP adapter	5720-VIM-2CE	10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
			10G DAC	Disabled	
			10G AOC		
			10G SFP+		
			25G DAC	Enabled	FEC is negotiated requesting CL108
			25G AOC		FEC auto:CL108
			25G SFP28		FEC auto:CL108
			40G DAC	Enabled	
			40G AOC		
			40G QSFP+		
			100G DAC	Enabled	FEC is negotiated requesting CL91
			100G AOC		FEC auto:CL91
			100G QSFP28		FEC auto:CL91

Table 75: 7520 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
10GBase-T	7520-48XT-6C		Enabled 1G, 10G	
SFP28	7520-48Y-8C 7520-48YE-8CE	1GBASE-T	Enabled The operational state is	

Table 75: 7520 Series default behavior (continued)

Port type		Model	Transceiver type	Auto-Negotiation	FEC
			Note: 7520-48Y-8C only	always enabled regardless of configuration.	
			1G SFP Note: 7520-48Y-8C only	Not Supported	
			10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
			10G DAC	Not Supported	
			10G AOC		
			25G DAC	Enabled	Auto-Negotiation
			25G AOC		CL108
QSFP28		7520-48XT-6C 7520-48Y-8C 7520-48YE-8CE	40G DAC	Enabled	
			40G AOC		
			100G DAC	Enabled	Auto-Negotiation
			100G AOC		CL91
QSFP28	Channelized	7520-48XT-6C 7520-48Y-8C 7520-48YE-8CE	40G DAC: 4x10G	Enabled	
			40G AOC: 4x10G		
			100G DAC: 4x25	Enabled	Auto-Negotiation
			100G AOC: 4x25		CL108

Table 75: 7520 Series default behavior (continued)

Port type		Model	Transceiver type	Auto-Negotiation	FEC
QSFP28	with QSFP adapter	7520-48XT-6C 7520-48Y-8C 7520-48YE-8CE	1G BASE-T	Enabled The operational state is always enabled regardless of configuration.	
			Note: 7520-48XT-6C and 7520-48Y-8C only		
			1G SFP	Not Supported	
			Note: 7520-48XT-6C and 7520-48Y-8C only		
			10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
			10G DAC	Not Supported	
			10G AOC		
			25G DAC	Enabled	Auto-Negotiation
			25G AOC		CL108

Table 76: 7720 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
QSFP28	7720-32C	40G DAC	Enabled	
		40G AOC		
		100G DAC	Enabled	Auto-Negotiation
		100G AOC		CL91

Table 76: 7720 Series default behavior (continued)

Port type		Model	Transceiver type	Auto-Negotiation	FEC
QSFP28	Channelized	7720-32C	40G DAC: 4x10G	Enabled	
			40G AOC: 4x10G		
			100G DAC: 4x25	Enabled	Auto-Negotiation
			100G AOC: 4x25		CL108
QSFP28	with QSFP adapter	7720-32C	1G BASE-T	Enabled The operational state is always enabled regardless of configuration.	
			1G SFP	Not Supported	
			10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
			10G DAC	Not Supported	
			10G AOC		
			25G DAC	Enabled	Auto-Negotiation
			25G AOC		CL108

Table 77: 7830 Series default behavior

Port type	Model	Transceiver type	Auto-Negotiation	FEC
SFP28	7830-VIM-24YE	10GBASE-T	Enabled The operational	

Table 77: 7830 Series default behavior (continued)

Port type	Model	Transceiver type	Auto-Negotiation	FEC
			state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G AOC		
		10G SFP+		
		25G DAC	Enabled	FEC is negotiated requesting CL108
		25G AOC		fec auto:CL108
		25G SFP28		fec auto:CL108
SFP56-DD	7830-VIM-24CE	10GBASE-T	Enabled The operational state is always enabled regardless of configuration.	
		10G DAC	Disabled	
		10G AOC		
		10G SFP+		
		25G DAC	Enabled	FEC is negotiated requesting CL108
		25G AOC		fec auto:CL108
		25G SFP28		fec auto:CL108
		100G SFP56-DD		fec auto:CL91

Table 77: 7830 Series default behavior (continued)

Port type		Model	Transceiver type	Auto-Negotiation	FEC
QSFP28	with QSFP adapter	7830-VIM-16CE 7830-32CE-8 DE	10GBASE-T	Enabled. The operational state is always enabled regardless of configuration.	
			10G DAC	Disabled	
			10G AOC		
			10G SFP+		
			25G DAC	Enabled	FEC is negotiated requesting CL108
			25G AOC		fec auto:CL108
			25G SFP28		fec auto:CL108
	Channelized	7830-VIM-16CE 7830-32CE-8 DE	40G DAC:4x10G		
			40G AOC:4x10G		
			100G DAC:4x25G	Enabled	FEC is negotiated requesting CL108
			100G AOC: 4x25G		fec auto:CL108
		7830-VIM-16CE 7830-32CE-8 DE	40G DAC	Enabled	
			40G AOC		
			40G QSFP+		
			100G DAC	Enabled	FEC is negotiated requesting CL91
			100G AOC		fec auto:CL91
			100G QSFP28		fec auto:CL91

Table 77: 7830 Series default behavior (continued)

Port type		Model	Transceiver type	Auto-Negotiation	FEC
QSFP56-DD	with QSFP adapter	7830-32CE-8 DE 7830-VIM-8DE	10GBASE-T		
			10G DAC		
			10G AOC		
			10G SFP+		
			25G DAC		
			25G AOC		
			25G SFP28		
	Channelized	7830-32CE-8 DE 7830-VIM-8DE	40G DAC:4x10G		
			40G AOC:4x10G		
			100G DAC:4x25G	Enabled	FEC is negotiated requesting CL108
			100G AOC: 4x25G		fec auto:CL108
			400G: 4x100G		fec auto:CL91
		7830-32CE-8 DE 7830-VIM-8DE	40G DAC	Enabled	
			40G AOC		
			40G QSFP+		
			100G DAC	Enabled	FEC is negotiated requesting CL91
			100G AOC		fec auto:CL91
			100G QSFP28		fec auto:CL91
			400G DAC		fec auto: CL119
			400G AOC		fec auto: CL119
			400G QSFP-DD		fec auto: CL119

10/100/1000 Mbps Port Considerations

Auto-Negotiation lets devices share a link, and automatically configures both devices so that they take maximum advantage of their abilities. Auto-Negotiation uses a modified 10BASE-T link integrity test pulse sequence to determine device ability.

Configure Auto-Negotiation as shown in the following table, where A and B are two Ethernet devices.

**Note**

Although the switch disables Auto-Negotiation by default if you insert a 1 Gbps SFP transceiver in an SFP+ or SFP28 port, depending on the platform, you can enable 1 Gbps Auto-Negotiation. For more information, see [Default Auto-Negotiation Behavior](#) on page 680 and [1Gb SFP Auto-Negotiation in SFP+ and SFP28 Ports](#) on page 697.

Table 78: Auto-Negotiation configuration on 10/100/1000BASE-TX ports

Port on A	Port on B	Remarks	Best practice
Auto-Negotiation enabled	Auto-Negotiation enabled	Ports negotiate on highest supported mode on both sides.	Use this configuration if both ports support Auto-Negotiation mode.
Full-duplex	Full-duplex	Both sides require the same mode.	Use this configuration if you require full-duplex, but the configuration does not support Auto-Negotiation.

Auto-Negotiation cannot detect the identities of neighbors and cannot shut down misconnected ports. Upper-layer protocols perform these functions.



Note

10 GigabitEthernet (GbE) fiber-based I/O module ports can operate at either 1 Gigabit per second (Gbps) or 10 Gbps, depending upon the capabilities of the optical transceiver that you install.

This situation presents an ambiguity with respect to the Auto-Negotiation configuration of the port, while 1 GbE ports require Auto-Negotiation; Auto-Negotiation is not defined and is non-existent for 10 GbE ports.

For a 10-GbE fiber-based I/O module, you can swap between 1 GbE and 10 GbE operation by simply swapping transceivers. To help with the swap, you can configure Auto-Negotiation when you install a 10 GbE transceiver, even though Auto-Negotiation is not defined for 10 GbE.

You can do this in anticipation of a port changeover from 10 GbE to 1 GbE. In this manner, you can pre-configure a port in 1 GbE mode while the 10 GbE transceiver is still installed. The port is ready to go upon the changeover to the 1 GbE transceiver.

You can use a saved configuration file with Auto-Negotiation enabled, to boot a system with either 10 GbE or 1 GbE transceivers installed. If you install a 1 GbE transceiver, the system applies Auto-Negotiation. If you install a 10 GbE transceiver, the system does not remove the Auto-Negotiation settings from the configuration, but the system simply ignores the configuration because Auto-Negotiation settings are irrelevant to a 10 GbE transceiver. The system preserves the saved configuration for Auto-Negotiation when re-saved no matter which speed of transceiver you install.

NEW! 1Gb SFP Auto-Negotiation in SFP+ and SFP28 Ports

The switch disables Auto-Negotiation by default if you insert a 1 Gbps SFP transceiver in an SFP+ or SFP28 port. For more information on default Auto-Negotiation behavior, see [Default Auto-Negotiation Behavior](#) on page 680.

Depending on the platform, you can enable 1 Gbps Auto-Negotiation on SFP+ and SFP28 ports. The following table identifies platforms that do not support 1 Gbps Auto-Negotiation on these ports.

Table 79: No 1Gb SFP Auto-Negotiation on SFP+ and SFP28

Platform	Not supported on ports
4220-4MW-8P-4X	13-16
4220-4MW-20P-4X	25-28
4220-8X	1-8
4220-12P-4X	13-16
4220-12T-4X	13-16
4220-24P-4X	25-28

Table 79: No 1Gb SFP Auto-Negotiation on SFP+ and SFP28 (continued)

Platform	Not supported on ports
4220-24T-4X	25-28
5320-16P-2MXT-2X	-
5320-16P-4XE	17-20
5320-16P-4XE-DC	17-20
5320-24P-8XE	25-28
5320-24T-4X-XT	25-28
5320-24T-8XE	25-28
5320-24T-24S-4XE-XT	49-52
5320-48P-8XE	49-52
5320-48T-8XE	49-52
7520-48XT-6C	1/51/1 1/49-1/50,1/52-1/54 have no 1Gb support
7520-48Y-8C	1-48,1/51/1 1/49-1/50,1/52-1/56 have no 1Gb support
7720-32C	1-32

25 GbE Port Considerations

25 GbE ports typically support 25 Gbps, 10 Gbps, and 1 Gbps operational speeds. Auto-Negotiation support varies depending on the pluggable type and speed.

For more information, see [Default Auto-Negotiation Behavior](#) on page 680.

Forward Error Correction (FEC) is a negotiated port attribute for 25 GbE connections that support Auto-Negotiation. For more information, see [Forward Error Correction](#) on page 701.

40 GbE Port Considerations

Auto-Negotiation must be enabled in 40 GbE ports when using 40GbCR4 (copper Direct Attached Cables - DACs) pluggable modules as Clause 73 of the 40 GbE standard lists it as mandatory. Though the links may come up in 40 GbE ports even without Auto-Negotiation, the best practice is to always enable Auto-Negotiation. Otherwise, there might be link instability or FCS errors.

100 GbE Port Considerations

Ensure that you enable Auto-Negotiation for ports with 100GbCR4 modules plugged in.

Although Auto-Negotiation is mandatory as per the 100GbCR4 standard, and this is the default software configuration, you can disable Auto-Negotiation to connect with older systems that do not support it. The system does not support FEC on 100GbCR4 links with Auto-Negotiation disabled.

For more information about FEC, see [Forward Error Correction](#) on page 701.

400 GbE Port Considerations

400 GbE is supported on 7830 Series only.

Forward Error Correction (FEC) is automatically configured for 400G QSFP56-DD, 400G DACs, and 400G AOCs upon insertion. The 400G QSFP56-DD module exclusively supports CL119 FEC, and its FEC settings are fixed and cannot be changed.

For more information about FEC, see [Forward Error Correction](#) on page 701.

Auto-Negotiation Advertisements

Auto-Negotiation advertisements use Custom Auto-Negotiation Advertisement (CANA) to control the speed and duplex settings that the interface modules advertise during Auto-Negotiation sessions between Ethernet devices. Modules can only establish links using these advertised settings, rather than at the highest common supported operating mode and data rate.

Use CANA to provide smooth migration from 10 Mbps to 10000 Mbps on host and server connections. Using Auto-Negotiation only, the switch always uses the fastest possible data rates. In limited-uplink-bandwidth scenarios, CANA provides control over negotiated access speeds, and improves control over traffic load patterns.

Use the **auto-negotiation-advertisements** command to configure CANA.

To use CANA, you must enable Auto-Negotiation.



Important

If a port belongs to a MultiLink Trunking (MLT) group and you configure CANA on the port (that is, you configure an advertisement other than the default), you must apply the same configuration to all other ports of the MLT group if they support CANA.

The following platforms support full duplex and half duplex modes for CANA:

Platform	Full duplex	Half duplex
4220 Series	Yes	Supported on 1Gbps copper ports at 10Mbps and 100Mbps.
5320 Series	Yes	Yes Note: Half duplex is supported at 10Mbps and 100Mbps speeds.

Platform	Full duplex	Half duplex
5420 Series	Supported on <ul style="list-style-type: none"> • 5420F-8W-16P-4XE • 5420F-24P-4XE • 5420F-24S-4XE • 5420F-24T-4XE • 5420F-48P-4XE • 5420F-48P-4XL • 5420F-48T-4XE • 5420M-24T-4YE • 5420M-24W-4YE • 5420M-24W-24S-4YE • 5420M-48T-4YE • 5420M-48W-4YE • only on ports 1/17-1/48 on 5420F-16W-32P-4XE, 5420F-16MW-32P-4XE, 5420M-16MW-32P-4YE 	Supported on <ul style="list-style-type: none"> • 5420F-8W-16P-4XE • 5420F-24P-4XE • 5420F-24T-4XE • 5420F-48P-4XE • 5420F-48P-4XL • 5420F-48T-4XE • 5420M-24T-4YE • 5420M-24W-4YE • 5420M-48T-4YE • 5420M-48W-4YE • only on ports 1/17-1/48 on 5420F-16W-32P-4XE, 5420F-16MW-32P-4XE, 5420M-16MW-32P-4YE • 5420M-24W-24S-4YE
5520 Series	Yes	Supported on 5520-24T, 5520-24W, 5520-48T, 5520-48W and ports 1/1/ - 1/36 on 5520-12MW-36W.
5720 Series	Yes	Yes
7520 Series	Yes	Yes
7720 Series	Yes	Yes
7830 Series	Yes	Yes

For more information about port capabilities, see the hardware documentation.

Forward Error Correction

Table 80: Forward Error Correction product support

Feature	Product	Release introduced
Forward Error Correction (FEC) (configurable)	4220 Series	Not Applicable
	5320 Series	Not Applicable
	5420 Series	VOSS 8.4 Supported on: <ul style="list-style-type: none"> • 5420M-16MW-32P-4YE • 5420M-24T-4YE • 5420M-24W-4YE • 5420M-48T-4YE • 5420M-48W-4YE
	5520 Series	VOSS 8.2.5 Supported on 5520-VIM-4YE
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1 Supported on VIM5-4YE at 25 Gbps
	VSP 7400 Series	VOSS 8.0

Forward Error Correction (FEC) is a method of obtaining error control in data transmission over an unreliable or noisy channel in which the source (transmitter) encodes the data in a redundant way by using an error correcting code (ECC). This redundancy enables a destination (receiver) to detect a limited number of errors and correct them without requiring a re-transmission.

FEC is useful where re-transmitting data is either expensive or impossible, for example, when transmitting to multiple receivers in multicast. However, although FEC provides more error control, it introduces a latency in data transmission.

FEC Configuration Options

You typically configure FEC on a port. The supported options are:

- cl91 (Clause 91 RS-FEC)
- cl108 (Clause 108 RS-FEC)
- auto

FEC is not supported on:

- Out-of-band (OOB) management ports.

- 100 GbE ports that are changed to 40 GbE ports by dynamically swapping 100 Gb modules with 40 Gb modules. FEC does not support the 40 Gbps speed.

**Important**

- On ports that support FEC configuration, ensure that you configure the same option at both end-points. Otherwise, the link does not come up.
- You must enable FEC to achieve proper functionality when using interconnects such as the 25Gb SR, 25 Gb SR-lite, 25 Gb ESR optics or the 25 Gb AOC and 25 Gb DAC.
- FEC is not required on 100 Gb long-range optics because these optics do error checking internally.

Clause 74 Firecode R-FEC

This option supports only the 25 Gbps speed and is used in applications that require reduced latency.

Clause 91 RS-FEC

This option supports both the 25 Gbps and 100 Gbps speeds. You can configure this option on ports with either the 100GbSR4 or 100GbCR4 modules plugged in, or on 100 GbE channelized ports operating at 25Gbps speed.

**Note**

Ensure that you enable Auto-Negotiation for ports with the 100GbCR4 modules plugged in; it is mandatory.

Clause 108 RS-FEC

This option also supports both the 25 Gbps and 100 Gbps speeds. It is similar to Clause 91 but provides extra latency.

This option automatically configures FEC based on port speed and pluggable module type.

Auto

- For 25 Gbps speeds, FEC CL108 is enabled for all transceiver types.
- For 100 Gbps speeds:
 - FEC is disabled for 100GbE LR4 and ER4 transceivers.
 - FEC CL91 is enabled for all other transceiver types (for example, 100GbE SR4, CR4, AOC, CWDM4, SWDM4).
- For 400 Gbps speeds, FEC CL119 is enabled for the QSFP56-DD transceiver, 400 Gbps DAC, 400 Gbps AOC.

FEC and Auto-Negotiation

FEC is a negotiated port attribute for 25 G and 100 G connections that support Auto-Negotiation. If you enable Auto-Negotiation on a port for a supported transceiver type, the switch uses the configured FEC value in the negotiation advertisement. Peers can

advertise different values, which means the resulting FEC operational state can be different than the one advertised.

The following table lists the 25 Gb end-point advertisements and the resulting FEC operational state:

Table 81: 25 Gb end-point advertisements

Peer A	Peer B	Result
CL108	CL108	CL108
CL74	CL74	CL74
No FEC	No FEC	No FEC
No FEC	CL108	CL108
No FEC	CL74	CL74
CL74	CL108	CL108

The following table lists the 100 Gb end-point advertisements and the resulting FEC operational state:

Table 82: 100 Gb end-point advertisements

Peer A	Peer B	Result
CL91	CL91	CL91
No FEC	No FEC	CL91 Note: Even when both peers advertise no FEC, negotiation results in clause 91 FEC per IEEE standard mandatory setting.
No FEC	CL91	CL91

You can use the **show interfaces gigabitEthernet config** command to see the FEC operational state for a port.

For additional details about support, see [Default Auto-Negotiation Behavior](#) on page 680.

Channelization

Table 83: Channelization product support

Feature	Product	Release introduced
Port Channelization	4220 Series	Not Applicable
	5320 Series	Not Applicable
	5420 Series	Not Applicable
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10 Port 1/51 only
	7720 Series	Fabric Engine 8.10 All ports except 1/8
	7830 Series	Fabric Engine 9.3 Supported on: <ul style="list-style-type: none"> ports 1/1-1/32 range: all odd ports ports 1/33-1/40: all ports 7830-VIM-16CE: all ports 7830-VIM-8DE: all ports
	VSP 4900 Series	VOSS 8.1 (4x10G only)
Port Auto-Channelization	VSP 7400 Series	VOSS 8.0 Supported on VSP 7432CQ
	4220 Series	Not Applicable
	5320 Series	Not Applicable
	5420 Series	Not Applicable
	5520 Series	Fabric Engine 8.9
	5720 Series	Fabric Engine 8.9
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.9 (4x10G only)
	VSP 7400 Series	VOSS 8.9 Supported on VSP 7432CQ

Use the channelization feature to configure a single port to operate as four individual ports. Channelization can apply to the following port speeds.

Table 84: Channelized operation with applicable port speeds

Form-factor	Channelized operation	Notes
40 Gbps (QSFP+)	four 10 Gbps ports	
100 Gbps (QSFP28)	four 25 Gbps ports	<p>If the hardware supports it, you can insert a 40 Gbps QSFP+ transceiver in a 100 Gbps port, and use the 100 Gbps as a 40 Gbps port. If you enable channelization on a 100 Gbps port and the switch detects a 40 Gbps QSFP+ transceiver in the port, the port operates as four individual 10 Gbps ports. If the switch detects a 100 Gbps QSFP28 transceiver and you enable channelization, the port operates as four 25 Gbps ports.</p> <p>To know which ports support channelization, or if you can use a 100 Gbps port as a 40 Gbps port and support the channelization of that port, see the applicable hardware documentation.</p>
400 Gbps (QSFP56-DD)	four 100 Gbps ports	<p>On the 7830-32CE-8DE port range 1/1-1/32, you can channelize only odd numbered ports. When you channelize an odd numbered port, the next even numbered port is not available for data traffic. For example, if ports 1/1 and 1/3 are channelized then ports 1/2 and 1/4 are unavailable for traffic.</p> <p>If the switch detects a 100 Gbps QSFP28 transceiver and you enable channelization, the port operates as four 10 Gbps or 25 Gbps ports.</p> <p>If the switch detects a 40 Gbps QSFP+ transceiver in a 400 Gbps port, the port operates as a 40 Gbps port. If you enable channelization, the port operates as four individual 10 Gbps ports.</p>

You can use breakout direct attach cables (DAC) or transceivers with fiber breakout cables to connect the channelized ports to other servers, storage, and switches.

For the number of ports on the switch that support channelization, see the applicable hardware documentation.

If the product supports channelization and you enable or disable channelization on a port, the port QoS configuration resets to default values. For information about configuring QoS values, see [Quality of Service](#) on page 2983.

**Note**

When you use channelized ports in an Split MultiLink Trunking (SMLT) configuration, the system does not display the channelized ports properly when you show MLT information for the remote port member if the remote switch runs a release that does not support channelization.

When a port is channelized, use only break out cables (copper or active optical DAC) in it. Using other cables in either a channelized port or a non-channelized port results in mismatched link status between link partners, which can lead to network issues.

Use the following CLI show commands to view the port channelization status from the GigabitEthernet Interface Configuration mode:

- **show interfaces gigabitEthernet channelize** [{slot/port[-slot/port]}
- **show interfaces gigabitEthernet channelize detail** [{slot/port/sub-port[-slot/port/sub-port] [, ...]}

Automatic Channelization

By default, ports are not channelized, which means that a port operates as one single port at the fully supported speed. The switch automatically channelizes a port if the port operates in Auto-sense mode and you insert one of the following:

- QSFP+ to SFP+ adapter
- QSFP28 to SFP28 adapter
- QSFP+ to 2xSFP+ passive or active breakout cable
- QSFP28 to 4xSFP28 passive or active breakout cable

You can also enable or disable channelization on a port manually. To disable automatic channelization on a port, you must disable Auto-sense on the port.

**Note**

Automatic channelization does not affect 40 Gbps and 100 Gbps SR parts; you can enable or disable channelization of ports with these parts regardless of the Auto-sense port configuration.

If you remove the existing part from an auto-channelized port, the port does not automatically disable channelization until you insert a new part that does not support channelization. For more information, see [Automatic Channelization with Auto-sense](#) on page 79.

SFP-DD Ports**Note**

Product Notice: SFP-DD ports are specific to 5420 Series.

Advanced Feature Bandwidth is enabled by default on the SFP-DD ports. When Advanced Feature Bandwidth is disabled, SFP-DD ports can operate as 1 Gbps ports with an SFP and as 10 Gbps ports with an SFP+. When the SFP-DD split DAC is used, SFP-DD ports can accept two SFP+ with 10 Gbps port speed.



Note

SFP-DD split DAC applies to certain 5420 Series models.

For more information about SFP-DD ports, see [5420 Series](#) on page 4220.

Feature Interaction with Channelization

Software features operate on channelized ports. When an interface is dechannelized, the interface cleans up all the channels.

If a feature operates on channel 1/1/1 and 1/1/2, and the circuit is dechannelized, the 1/1/1 configuration is saved and the commands are configured on 1/1. The configuration on 1/1/2 is deleted.

IEEE 802.3X Pause Frame Transmit

Table 85: IEEE 802.3X Pause Frame Transmit product support

Feature	Product	Release introduced
IEEE 802.3X Pause frame transmit	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The switch uses MAC pause frames to provide congestion relief on full-duplex interfaces.

Overview

When congestion occurs on a port, the system can send or receive pause frames, also known as flow control, to temporarily pause the packet flow. The system uses flow control if the rate at which one or more ports receives or sends packets is greater than the rate the switch can process or accept the packets.

The switch can generate pause frames to tell the sending device to stop sending additional packets for a specified time period. After the time period expires, the

sending device can resume sending packets. During the specified time period, if the switch determines the congestion is reduced, it can send pause frames to the sending device to instruct it to begin sending packets immediately.

**Note**

The 7830 Series can process received pause frames but it cannot generate or send pause frames. When a port receives a pause frame, the switch reduces the traffic rate on that port but it does not send pause frames towards the ports that cause the over subscription.

Flow control mode and pause frames

If you enable flow control mode, the switch drops packets on ingress when congestion occurs. If the switch is not in flow control mode, it drops packets at egress when congestion occurs.

Configure an interface to send pause frames when congestion occurs to alleviate packet drops due to flow control mode.

Auto-Negotiation

Interfaces that support auto-negotiation advertise and exchange their flow control capability to agree on a pause frame configuration. IEEE 802.3 annex 28b defines the auto-negotiation ability fields and the pause resolution. The switch advertises only two capabilities. The following table shows the software bit settings based on the flow control configuration.

**Note**

Not all interfaces support Auto-Negotiation. For more information, see your hardware documentation.

Table 86: Advertised abilities

Interface configuration	Pause	ASM	Capability advertised
Flow control enabled	1	0	Symmetric pause
Flow control disabled	1	1	Both Symmetric pause and asymmetric pause

The following tables identifies the pause resolution.

Table 87: Pause resolution

Local device pause	Local device ASM	Peer device pause	Peer device ASM	Local device resolution	Peer device resolution
0	0	Do not care	Do not care	Disable pause transmit and receive.	Disable pause transmit and receive.
0	1	0	Do not care	Disable pause transmit and receive.	Disable pause transmit and receive.
0	1	1	0	Disable pause transmit and receive.	Disable pause transmit and receive.
0	1	1	1	Enable pause transmit. Disable pause receive.	Disable pause transmit. Enable pause receive.
1	0	0	Do not care	Disable pause transmit and receive.	Disable pause transmit and receive.
1	Do not care	1	Do not care	Enable pause transmit and receive.	Enable pause transmit and receive.
1	1	0	0	Disable pause transmit and receive.	Disable pause transmit and receive.
1	1	0	1	Disable pause transmit. Enable pause receive.	Enable pause transmit. Disable pause receive.

The following list identifies the type of interfaces that support auto-negotiated flow control:

- 10 Mbps/100 Mbps/1 Gbps copper
- 100 Mbps/1 Gbps/10 Gbps copper
- 1 Gbps fiber (in both SFP and SFP+ ports)

Auto MDIX

Automatic medium-dependent interface crossover (Auto-MDIX) automatically detects the need for a straight-through or crossover cable connection and configures the connection appropriately. This removes the need for crossover cables to interconnect switches and ensures either type of cable can be used. The speed and duplex setting of an interface must be set to Auto for Auto-MDIX to operate correctly.

Auto MDIX is supported on all platforms with fixed copper ports. All fixed copper ports are supported.



Note

The management port on the 5520 Series switch does not support Auto-MDIX (the automatic detection of transmit and received twisted pairs).

As a best practice, enable the default auto-negotiation setting on the management port.

Because the management port does not support Auto-MDIX, when auto-negotiation is disabled, a crossover cable might be necessary to have the port link up and pass traffic.

If the peer device supports Auto-MDIX, then either a straight through or crossover will work. The issue occurs only if both ends of the connection do not support Auto-MDIX.

NEW PTPv2 Transparent Clock with VLANs

Table 88: Precision Time Protocol version 2 product support

Feature	Product	Release introduced
Transparent Clock with VLANs	4220 Series	Not Supported
	5320 Series	Fabric Engine 9.4 Supported on 5320-16P-2MXT-2X only
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.4
	5720 Series	Fabric Engine 9.4
	7520 Series	Fabric Engine 9.4
	7720 Series	Fabric Engine 9.4
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.4
	VSP 7400 Series	VOSS 9.4

Precision Time Protocol Overview



Note

The IEEE 1588 Working Group updated its terminology, replacing master and slave with timeTransmitter and timeReceiver.

IEEE 1588v2 Precision Time Protocol (PTPv2) synchronizes clocks across Ethernet networks with nanosecond-level accuracy. PTP distributes frequency, phase, and time-of-day information using timestamped event messages so devices across the network can align their clocks precisely. The protocol uses these timestamped packets to maintain accurate timing across different network topologies and switching environments.

A PTP network uses a hierarchical timing structure. A grandmaster clock supplies the reference time for all other devices in the network. timeTransmitter clocks synchronize to the grandmaster by exchanging timestamped messages and can also act as timeTransmitter clocks for downstream devices in multi-tier networks. A PTP path can include boundary clocks, transparent clocks, and non-PTP-aware network switches that forward PTP event messages without performing timing correction.

Accurate synchronization depends on precise delay measurement between the timeTransmitter clock and the timeReceiver clock. Devices calculate this path delay by exchanging timestamped messages and compensating for transmission and processing latencies. Because link delay changes slowly, devices repeat these measurements regularly. After determining link delay, the grandmaster sends synchronization messages to the timeReceiver clocks.

**Note**

This feature requires a Premier or a subscription license.

End-to-End Mode

In End-to-End (E2E) mode, the switch calculates delay directly between the timeTransmitter and timeReceiver, which allows PTP to operate over networks with Transparent clocks, non-PTP aware network switches, or a combination of both. The timeReceiver sends Delay Request messages to the timeTransmitter. The timeTransmitter returns Delay Response messages, and the timeReceiver uses these timestamps to calculate the total path delay and adjust its clock within a VLAN.

Transparent Clock Operation

A Transparent Clock improves synchronization accuracy by updating the Correction Field in PTP event messages to reflect the time the switch spends processing each packet. The switch measures residence time—the interval between packet ingress and egress—and adds that value to the Correction Field.

A Transparent Clock updates the following message types:

- Sync
- Delay Request
- Delay Response

PTP Transparent Clock with VLANs

When you enable PTPv2 Transparent Clock with VLANs support, the switch measures residence time for PTP event messages that enter and exit PTP-capable ports within the same VLAN. The switch updates the Correction Field at each hop, and slave devices use the cumulative Correction Field value to correct timing offsets accurately.

The Transparent Clock does not participate in timeTransmitter and timeReceiver selection. Instead, it improves timing precision by adding accurate per-hop delay information to PTP event messages. This behavior preserves precise timing between the grandmaster and downstream timeReceiver devices, even when the path includes multiple PTP-capable switches.

Operational Considerations

Consider the following when you use the operational considerations when you deploy or manage PTP in your network.

- On 5320 Series and 5420 Series, MACsec and PTP are mutually exclusive on ports that support both features. Enable only one feature on each port.
- PTP requires a Premier or a subscription license.
- When the switch boots in Zero-Touch Deployment (ZTD) mode with a valid license, it automatically enables PTP.
- When the switch boots in ZTD mode with no valid license, it blocks PTP and generates an error message and a license-violation log.

NEW! Hot Swap a VIM on 7830 Series

You can hot swap a Versatile Interface Module (VIM) on the switch using one of the following modes:



Note

As a best practice, use Controlled mode when you replace a VIM. Use Uncontrolled mode only if CLI and EDM are not available.

- Controlled mode
 1. As a best practice, use the **slot shutdown** command to shut down the port in a VIM before you power off the VIM slot.
 2. Power off the VIM slot.
 3. Insert the replacement VIM.
 4. Power on the VIM slot.
 5. Confirm the switch detects the new VIM.
- Uncontrolled mode
 1. As a best practice, remove the transceiver.
 2. Remove the VIM.
 3. Insert the replacement VIM.
 4. Insert the transceiver only after the VIM initializes.

Node Alias

Table 89: Node Alias product support

Feature	Product	Release introduced
Node Alias	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.1
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1
DHCP Fingerprinting	4220 Series	Fabric Engine 9.1
	5320 Series	Fabric Engine 9.1
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1

Node Alias is a powerful troubleshooting tool. With Node Alias functionality, you can discover information about the end systems in the network on a per-port basis. DHCP Fingerprinting is integrated into Node Alias functionality and can provide information about the device that sent a DHCP request.

As the packets from the end systems transit through the Node Alias-enabled port, the switch stores key attributes in a client information database that you can query to retrieve specific system information, including protocol and DHCP client information.

The client database stores the following client, protocol, and DHCP Fingerprint information for each client:

- Source MAC address
- I-SID
- VLAN
- Incoming interface
- Protocols the client uses

- IPv4 Address
- IPv6 Address
- IsActive flag
- First learned time
- Last learned time
- Host name
- OS type
- Device type
- OS version
- Protocol information type
- Protocol information value
- Option 12 information
- Option 53 information
- Option 55 information
- Option 60 information

Operational Considerations

Consider the following when you use Node Alias:

- By default, the switch enables Node Alias globally, as well as on Auto-sense enabled ports. If you disable Node Alias globally, port level functionality is also disabled. If the ports are combined to form an MLT (LAG), you must enable or disable Node Alias on individual ports.
- You cannot enable Node Alias on SPBM NNI ports.
- You cannot enable Node Alias on Extreme Integrated Application Hosting ports.
- The protocol information database stores detailed information for IPv4, IPv6, BGP, VRRP, and UDP only.
- DHCP Fingerprinting is integrated into Node Alias functionality; therefore, you must enable Node Alias to perform DHCP Fingerprinting.
- The maximum number of Node Alias clients is on per switch, not per port.
- The number of clients stored in the Node Alias client database cannot be lower than the number of existing clients.
- If you do not explicitly clear entries when the client database is full, inactive entries automatically clear first followed by the first oldest entry.
- When you disable Node Alias on a port, clients learned on that port are deleted unless the port is a member of an MLT with at least one port with Node Alias enabled. If the MLT has no ports with Node Alias enabled, the entries are deleted.

Chassis Operations Configuration Using the CLI

This section provides the details to configure basic hardware and system settings.

Configure Jumbo Frames

About This Task

Configure jumbo frames to increase the Ethernet frame size to the maximum frame size your switch supports.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure jumbo frames:

```
sys mtu <sys-mtu>
```

Example

```
Switch:1> enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1#(config)#sys mtu 1523
```

Variable Definitions

The following table defines parameters for the **sys mtu** command.

Variable	Value
<sys-mtu>	Specifies the Jumbo frame size for the data path. Different hardware platforms support different MTU sizes. Use the CLI Help to see the available range for the switch. The default is 1950 bytes.

Configure Port Lock

About This Task

Configure port lock to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify a locked port until you unlock the port.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable port lock globally:

```
portlock enable
```

3. Log on to GigabitEthernet Interface Configuration mode:

```
interface gigabitethernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

4. Lock a port:

```
lock port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all]
[all]} enable
```

Example

Unlock port 1/14:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface GigabitEthernet 1/1
Switch:1(config-if)#no lock port 1/14 enable
```

Variable Definitions

The following table defines parameters for the **interface gigabitethernet** and **lock port** commands.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <p>For the <code>lock port</code> command, use the <code>no</code> form of this command to unlock a port:</p> <pre>no lock port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}</pre>

Configure SONMP

About This Task

Configure the SynOptics Network Management Protocol (SONMP) to allow a network management station (NMS) formulate a map that shows the interconnections between Layer 2 devices in a network. The default status is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable SONMP:

```
no autotopology
```

3. Enable SONMP:

```
autotopology
```

View the Topology Message Status

About This Task

View topology message status to view the interconnections between Layer 2 devices in a network.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the contents of the topology table:

```
show autotopology nmm-table
```

Unless the switch is physically connected to other devices in the network, this topology will be blank.

Example



Note

In the following example, the column “ChassisType” uses a generic name. When you use the **show autotopology nmm-table**, your switch displays the actual chassis type.

```
Switch:1>show autotopology nmm-table
```

Topology Table								
Local Port	IpAddress	SegmentId	MacAddress	ChassisType	BT	LS	CS	Rem Port
0/0	192.0.2.81	0x000000	0030ab707a00	ChassisType 1	12	Yes	HtBt	0/0
1/1	192.0.2.81	0x000000	0050ea268800	ChassisType 2	12	Yes	HtBt	1/12
1/24	192.0.2.81	0x000000	070ab307aa00	ChassisType 3	12	Yes	HtBt	1/1
2/1	192.0.2.81	0x000000	0030ab57ab00	ChassisType 4	12	Yes	HtBt	1/23
2/2	192.0.2.81	0x000000	0030ab307af0	ChassisType 5	12	Yes	HtBt	1/12



Note

When a peer switch is running an older software version that does not include support for SONMP hello messages with channelization information, it can only show the slot/port. It cannot show the sub-port.

Configure Ethernet Ports with Auto-Negotiation

Configure Ethernet ports so they operate optimally for your network conditions.

About This Task

When you use 1 Gigabit Ethernet SFP transceivers, the software disables Auto-Negotiation on the port. If you use 1 Gbps SFP transceivers, the remote end must also have Auto-Negotiation disabled.

All ports that belong to the same MLT or Link Aggregation Control Protocol (LACP) group must use the same port speed. In the case of MLTs, the software does not enforce this.

The software requires the same Auto-Negotiation settings on link partners to avoid incorrect declaration of link status. Mismatched settings can cause the links to stay down. Ensure the Auto-Negotiation settings between local ports and their remote link partners match before you upgrade the software.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...] [slot/all] [all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable Auto-Negotiation:

```
auto-negotiate [port {slot/port [-slot/port] [,...] [slot/all] [all]}
enable
```

3. Verify the configuration:

```
show interfaces gigabitEthernet l1-config {slot/port [-slot/port] [,...]
[slot/all] [all]}
```

Example

```
Switch:>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/8
Switch:1(config-if)#auto-negotiate enable
Switch:1(config-if)#show interfaces gigabitEthernet l1-config 1/8
```

```
=====
                        Port Config L1
=====
```

PORT NUM	AUTO NEG.	OPERATE AUTO-NEG	CUSTOM ADVERTISEMENTS	AUTO NEGOTIATION	CANA ORIGIN	ADMIN DPLX SPD	OPERATE DPLX SPD	ADMIN TX-FLW-CTRL	OPERATE TX-FLW-CTRL
1/8	true	true	Not Configured		RADIUS	full 10000	0	enable	enable

Variable Definitions

The following table defines parameters for the **auto-negotiate** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	Specifies the port or ports that you want to configure.
<code>enable</code>	Enables auto-negotiation for the port or other ports of the module. The default auto-negotiation behavior depends on the switch model and transceiver type.

Configure Auto-Negotiation Advertisements

Configure local port Auto-Negotiation advertisements to specify the speed and duplex mode for traffic between local ports and remote link partners. Supported speeds and duplex modes vary, depending on your hardware.

Before You Begin

You must enable Auto-Negotiation before you perform this procedure.

About This Task

Configure local port Auto-Negotiation advertisements.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Auto-Negotiation advertisements on one or more ports:

```
auto-negotiation-advertisements {25000-full|10000-full|2500-full|5000-
full|1000-full|100-full|100-half|10-full|10-half}
```

or

```
auto-negotiation-advertisements port {{slot/port[-slot/port][,...]}
[slot/all][all]} {25000-full|10000-full|2500-full|5000-full|1000-full|
100-full|100-half|10-full|10-half}
```

3. Verify the configuration:

```
show interfaces gigabitEthernet 11-config {slot/port[-slot/port][,...]
[slot/all][all]}
```

Variable Definitions

The following table defines parameters for the **auto-negotiation-advertisements** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}	Specifies the port or ports that you want to configure.
25000-full	Advertises 25 Gbps full-duplex.
10000-full	Advertises 10 Gbps full-duplex.
5000-full	Advertises 5 Gbps full-duplex.
2500-full	Advertises 2.5 Gbps full-duplex.
1000-full	Advertises 1 Gbps full-duplex.
100-full	Advertises 100 Mbps full-duplex.
100-half	Advertises 100 Mbps half-duplex.
10-full	Advertises 10 Mbps full-duplex.
10-half	Advertises 10 Mbps half-duplex.
none	Configures the Auto-Negotiate value to none.

Configure IEEE 802.3X Pause Frame Transmit

About This Task

Configure IEEE 802.3X Pause frame transmit to eliminate or minimize packet loss.

By default, flow control mode is disabled. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames.

**Note**

The 7830 Series can process received pause frames but it cannot generate or send pause frames. When a port receives a pause frame, the switch reduces the traffic rate on that port but it does not send pause frames towards the ports that cause the over subscription.

By default, an interface does not send pause frames.



Note

If you enable MACsec on an interface and you send small packet size traffic near line rate, the **In FlowCtrl** frame might increment in the output of the **show interfaces gigabitEthernet statistics** command because of the processing overhead caused by adding the MACsec header of 32 bytes. This is part of the expected over-subscription footprint.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable flow control mode:

```
boot config flags flow-control-mode
```

3. Save the configuration.

4. Return to Privileged EXEC mode:

```
end
```

5. Reboot the chassis.

```
boot
```

6. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

7. Configure the interface to generate pause frames:

```
tx-flow-control [enable]
```

8. (Optional) Configure other interfaces to generate pause frames:

```
tx-flow-control port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} enable
```

9. Verify the boot flag configuration:

```
show boot config flags
```

10. Verify the interface configuration:

```
show interfaces gigabitEthernet ll-config {slot/port[/sub-port] [-slot/
port[/sub-port]][,...][slot/all][all]}
```

11. View the pause-frame packet count:

```
show interfaces gigabitEthernet statistics {slot/port[/sub-port] [-
slot/port[/sub-port]][,...][slot/all][all]}
```

Example

Enable flow control on the system and configure slot 1, port 10 to send pause frames. Verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags flow-control-mode
Warning: Please save the configuration and reboot the switch
        for this configuration to take effect.
Switch:1<config>#save config
CP-1: Save config to file /intflash/config.cfg successful.
CP-1: Save license to file /intflash/license.xml successful.
Switch:1<config>#disable
Switch:1#boot
Are you sure you want to re-boot the switch (y/n) ?y
```



Note

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false

Switch:1(config-if)#show interfaces gigabitEthernet 1/10-config 1/10
=====
Port Config L1
=====
=====
```

PORT NUM	AUTO NEG.	OPERATE AUTO-NEG	CUSTOM ADVERTISEMENTS	AUTO NEGOTIATION	CANA ORIGIN	ADMIN DPLX	OPERATE SPD	ADMIN DPLX	OPERATE SPD	ADMIN TX-FLW-CTRL	OPERATE TX-FLW-CTRL
1/10	true	true	Not Configured		RADIUS	full	10000	0		enable	enable

View the pause-frame packet count for slot 1, port 10.

```
Switch:1(config-if)#show interfaces gigabitEthernet statistics 1/10
=====
Port Stats Interface
=====
```

PORT NUM	IN OCTETS	OUT OCTETS	IN PACKET	OUT PACKET
1/1	29964704384	22788614528	234106526	178034166

```
=====
```

PORT OUTLOSS NUM PACKETS	IN FLOWCTRL	OUT FLOWCTRL	IN PFC	OUT PFC
1/1	0	11014	0	0

```
=====
```

Variable Definitions

The following table defines parameters for the **tx-flow-control** command.

Variable	Value
<i>enable</i>	Configures the interface to send pause frames. By default, flow control is disabled.
<i>port {slot/ port[/sub-port] [-slot/port[/ sub-port]] [,...][slot/ all][all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Use the data in the following table to use the **show interfaces gigabitEthernet 11-config** and **show interfaces gigabitEthernet statistics** commands.

Variable	Value
<i>{slot/port[/ sub-port] [- slot/port[/sub- port]][,...] [slot/all] [all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Enable Channelization

Enable channelization on a port to configure it to operate as four channels, or ports.



Note

The switch automatically channelizes a port if the port operates in Auto-sense mode and you insert a QSFP+ to SFP+ or QSFP28 to SFP28 adapter, or a breakout cable.

About This Task



Important

Enabling or disabling channelization resets the port QoS configuration to default values.



Note

Only odd numbered ports can be channelized on the 7830 Series. When an odd numbered port is channelized in the 1/1-1/32 range, the next even numbered port is not available for data traffic.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable channelization on a port:

```
channelize [port {slot/port[-slot/port] [,...]}] enable
```

3. Display the status of the ports:

```
show interfaces gigabitEthernet channelize [{slot/port[-slot/port]}  
[,...]]
```

To display the details of the sub-ports, use:

```
show interfaces gigabitEthernet channelize detail [{slot/port/sub-  
port[-slot/port/sub-port] [,...]]
```

Examples

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)# interface gigabitEthernet 2/1  
Switch:1(config-if)# channelize enable  
Enabling channelization on port 2/1. Subport 2/1/1 will inherit port 2/1 configuration.
```

```
Subports 2,3,4 will use default config. QSFP will be reset as removal and re-insert.  
NOTE: Modify QOS configurations on all subports as required.  
Do you wish to continue (y/n) ? y
```

Display the port status:

```
Switch:1(config)# show interfaces gigabitEthernet channelize 2/1-2/2
```

```
=====
                        Port Channelization
=====
-----
PORT          ADMIN MODE    CHANNEL TYPE
-----
2/1           true             40G
2/2           true             40G
```

The following is an example of how to disable channelization on a port:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitEthernet 2/1/1
Switch:1(config-if)# no channelize enable
```

Variable Definitions

The following table defines parameters for the **channelization** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<div>Identifies the slot and port in one of the following formats:</div> <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) <div>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</div>

Configure FEC on a Port

About This Task

Use this procedure to configure Forward Error Correction (FEC) on supported ports.

 **Note**
For

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. (Optional) Specify the port or ports to configure for FEC:

```
fec port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all]
[all]}
```

3. Configure FEC on a port:

```
fec {auto | cl108 | cl174 | cl191}
```

4. Verify the configuration:

```
show interfaces gigabitEthernet config {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]}
```

Examples

Configure Clause 108 FEC on a 25 Gbps port 1/1:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#fec cl108
```

Verify the configuration when a 25 Gbps optic is present:

```
Switch:1(config-if)#show interfaces gigabitEthernet config 1/1
=====
Port Config
=====
PORT          DIFF-SERV  QOS  MLT  VENDOR
NUM    TYPE      EN    TYPE  LVL  ID   NAME
-----
1/1      25GbCX      true  core   1    0   Extreme
=====
Port Config
=====
PORT  ADMIN  OPERATE AUTO  ACCESS-SERV  RMON  FLEX-UNI  ADMIN  APPLICABLE  OPERATE
NUM   ROUTING ROUTING RECOVER  EN                      FEC    FEC          FEC
-----
1/1   Enable  Disable Disable false                    Disable Disable  Auto  CL108    CL108
```

Verify the configuration when a 10 Gb optic is present in a 25 Gb port:

```
Switch:1(config-if)#show interfaces gigabitEthernet config 1/1
=====
Port Config
=====
```

PORT NUM	TYPE	DIFF-SERV EN	TYPE	QOS LVL	MLT ID	VENDOR NAME				
1/1	10GbSR	true	core	1	0	Extreme				
PORT NUM	ADMIN ROUTING	OPERATE ROUTING	AUTO RECOVER	ACCESS-SERV EN	RMON	FLEX-UNI FEC	ADMIN FEC	APPLICABLE FEC	OPERATE FEC	
1/1	Enable	Disable	Disable	false	Disable	Disable	Auto	Not Applicable	Off	

Variable Definitions

The following table defines parameters for the **fec** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>{auto cl108 cl174 cl91}</code>	<p>Configures one of the following options for FEC on the port:</p> <ul style="list-style-type: none"> auto Clause 91 Clause 108 Clause 74 <p>Note: On a 100 GbE port, only the Clause 91 and Clause 108 options are supported. On 100 GbE channelized ports (operating at 25 Gbps speed), you can configure Clause 108 for extra latency or Clause 74 for reduced latency. Configuration of FEC is not supported on a management port or on 100 GbE ports operating at 40 Gbps speed.</p> <p>Important: On ports that support FEC, always configure the same option on both end-points. Otherwise, the link does not come up.</p>

Configure Serial Management Port Dropping

Configure the serial management ports to drop a connection that is interrupted for any reason. If you enable serial port dropping, the serial management ports drop the connection for the following reasons:

- modem power failure
- link disconnection
- loss of the carrier

Serial ports interrupted due to link disconnection, power failure, or other reasons force out the user and end the user session. Ending the user session ensures a maintenance port is not available with an active session that can allow unauthorized use by someone other than the authenticated user, and prevents the physical hijacking of an active session by unplugging the connected cable and plugging in another.

By default, the feature is disabled with enhanced secure mode disabled. If enhanced secure mode is enabled, the default is enabled.

For more information on enhanced secure mode, see [Enable Enhanced Secure Mode](#) on page 324.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the serial port to drop if a connection is interrupted:
`sys security-console`

Enable the Locator LED

About This Task

Perform this procedure to turn the system Locator LED on to provide a visual identification of a specific switch.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the system Locator LED:
`sys locator-led`
3. Display the system Locator LED status:
`show sys locator-led`

Enable or Disable the USB Port

Perform this procedure to control USB access. For security reasons, you may want to disable this port to prevent individuals from using it. By default, the port is automatically mounted when a USB device is inserted.

Before You Begin

- The switch must be in Enhanced Secure mode.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Disable the USB port:

```
sys usb disable
```
3. Enable a previously disabled USB port:

```
no sys usb disable
```

Configure Port Speed

About This Task

Manually configure the port speed.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the speed for one or more ports:

```
speed {10|100|1000|10000|2500|25000|5000}
```

or

```
speed port {slot/port[-slot/port]}[,...][slot/all][all]} {10|100|1000|
10000|2500|25000|5000}
```

Variable Definitions

The following table defines parameters for the **speed** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	Specifies the port or ports that you want to configure.
<code>10</code>	Configures the port speed to 10 Mbps.
<code>100</code>	Configures the port speed to 100 Mbps.
<code>1000</code>	Configures the port speed to 1 Gbps.
<code>10000</code>	Configures the port speed to 10 Gbps.
<code>2500</code>	Configures the port speed to 2.5 Gbps.
<code>25000</code>	Configures the port speed to 25 Gbps.
<code>5000</code>	Configures the port speed to 5 Gbps.

Configure Port Speeds for All VIM Ports



Note

This procedure only applies to 5520 Series, 5720 Series, and 7830 Series.

Configure all of the ports on an installed Versatile Interface Module (VIM) to operate at the same speed.



Note

Some VIMs must operate with all ports at the same speed, or with a group of ports at the same speed, while others can operate with ports at different speeds. For more information, see *Fabric Engine Release Notes*. The **sys vim-speed** command is supported only on VIMs that must operate with all ports at the same speed. An error message displays if you run the command on an unsupported VIM.

Before You Begin

Install the VIM before performing this procedure.

About This Task



Note

As a best practice, do not mix 10Gbps and 25Gbps VIMs with the auto parameter.

Use this procedure to configure the speed of all ports in a multi-port VIM to operate at either 1 Gbps, 10 Gbps, 25 Gbps, or auto.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```


2. Configure the speed for all of the VIM ports:

```
sys vim-speed {1000 | 10000 | 25000 | auto}
```

3. Configure all VIM ports into two speed setting groups:

```
sys vim-speed group <1-2> {1000 | 10000 | 25000}
```

Variable Definitions

The following table defines parameters for the **sys vim-speed** command.

Variable	Value
<code>1000 10000 25000 auto</code> Note: Exception: 1 Gbps port speed applies to 5720-VIM-6YE only.	Configures all ports in a multi-port VIM to operate at either 1 Gbps, 10 Gbps, 25 Gbps, or auto. The command auto vim-speed is enabled by default.
<code>group <1-2></code> Note: Exception: only applies to 5720-VIM-6YE.	Configures all ports into two speed setting groups. <ul style="list-style-type: none"> • group 1: ports 2/1, 2/2, 2/3 • group 2: ports 2/4, 2/5, 2/6 Ports in groups must be configured to operate at the same speed. For example, ports 2/1, 2/2, 2/3 can be configured for 10 Gbps, and ports 2/4, 2/5, 2/6 can be configured for 25 Gbps. If 5720-VIM-6YE uses the LRM/MACsec Adapter, you must configure the groups according to the SFP or SFP+ inserted in the Adapter.

Disable Auto-Speed

About This Task

Perform the following procedure to disable auto-speed on one or more ports. This procedure is for testing purposes on the 7830 Series only.



Note

Auto-speed is enabled by default. If you change auto-speed configuration on one port, you change the configuration on all four ports in the group. For example, disabling auto-speed on port 1/2 also disables it on ports 1/1, 1/3, and 1/4. For information about group port speeds on 7830 Series, see [Multi-speed Ports](#) on page 675

For more information about auto-speed, see [Auto-speed](#) on page 677.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable auto-speed:

```
no auto-speed enable [port {slot/port[/sub-port]}[-slot/port[/sub-
port]][,...][slot/all][all]] enable]
```

Variable Definitions

The following table defines parameters for the **auto-speed** command.

Variable	Value
<i>enable</i>	Configures auto-speed for a port or ports.
<i>port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

NEW Enable PTPv2

About This Task

Perform the following procedure to enable PTPv2 Transparent Clock with VLANs support on a switch with a valid license.



Note Exception: This procedure does not apply to 4220 Series and 7830 Series.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable PTPv2:

```
ptp enable
```

3. Verify the configuration:

```
show ptp
```

NEW! Swap a VIM

You can hot swap a Versatile Interface Module (VIM) on the 7830 Series using Controlled mode or Uncontrolled mode.

About This Task

Consider the following when you swap VIMs:

- Mandatory wait times between VIM swaps are as follows:
 - 15 seconds after you remove a VIM
 - 60 seconds after you insert a VIM
- When you use uncontrolled mode:
 - Remove the transceiver before you remove the VIM, and insert the replacement VIM before you install the transceiver.
 - Wait 60 seconds after you insert the VIM before you install the transceiver.
- When you swap between the same VIM type, configuration on the ports is preserved then reapplied. When you swap between different VIM types, configuration on the ports revert to default configuration.
- MACSec configuration on channelized ports is lost when you swap between the same VIM type.

Perform the following steps to swap a VIM.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Shut down all ports on the VIM slot:

```
slot shutdown {slot[-slot][,...]}
```

3. Power off the VIM:

```
no sys power slot {slot[-slot][,...]}
```

4. Remove the VIM.

5. Insert the replacement VIM.

6. Power on the VIM:

```
sys power slot {slot[-slot][,...]}
```

7. Enable the ports on VIM slot:

```
no slot shutdown {slot[-slot][,...]}
```

8. Verify the status of the VIM:

```
show sys-info
```

NEW! Variable Definitions

The following table defines parameters for the **slot shutdown** command.

Variable	Value
<code>{slot[-slot][,...]}</code>	Specifies the valid slot. The valid range is 2-3 on 7830 Series.

The following table defines parameters for the **sys power slot** command.

Variable	Value
<code>{slot[-slot][,...]}</code>	Specifies the valid slot. The valid range is 2-3 on 7830 Series.

Node Alias Configuration using CLI

This section provides procedures to configure Node Alias using Command Line Interface (CLI).

Configure Node Alias

About This Task

By default, the switch enables Node Alias globally, as well as on Auto-sense enabled ports.



Note

You cannot enable Node Alias on SPBM NNI ports.

If you disable Node Alias globally, port level functionality is also disabled.



Note

The steps in this procedure assume Node Alias is disabled globally as well as on ports.

Perform the following procedure to configure Node Alias.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable Node Alias globally:

```
nodealias enable
```

3. Configure the number of clients stored in the Node Alias client database on the switch:

```
nodealias max-clients <max-clients>
```

4. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

5. Enable Node Alias on a port:

```
nodealias enable
```

6. Verify the configuration:

```
show nodealias
```

Example

Display basic information for Node Alias clients:

```
Switch:1>show nodealias
```

Node Alias - General Info

```
NodeAlias is globally enabled:      TRUE
Maximum Number of Clients:          8192
Port List NodeAlias Enabled:        1/14,1/17,2/1
Number of learned clients:          5
Number of active clients:           5
Number of protocol related entries: 6
```

Variable Definitions

The following table defines parameters for the **nodealias** command.

<i>enable</i>	Enables Node Alias globally.
<i>max-clients</i> <max-clients>	Specifies the number of clients the switch can store in the client database. Use the CLI Help to see the available range for the switch. The default is 8,192.

*Clear Node Alias Client Information***About This Task**

Perform the following procedure to clear Node Alias client information from the database. You can clear information for all clients or a single client globally as well as on a specific port or MLT.

**Note**

If ports combine to form an MLT (LAG), you must clear Node Alias client information on individual ports.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Clear Node Alias client information:
clear nodealias [client mac <x00:0x00:0x00:0x00:0x00:0x00> i-sid <0-16777215> vlan <0-4094>] [mlt <1-512>] [port {slot/port[/sub-port]}]

Variable Definitions

The following table defines parameters for the **clear nodealias** command.

Variable	Value
<i>isid</i> <0-16777215>	Specifies the I-SID.
<i>mac</i> <x00:0x00:0x00:0x00:0x00:0x00>	Specifies the MAC address in hexadecimal format. The MAC address parameter does not accept MAC addresses beginning with 01:00:5e (01:00:5e:00:00:00 to 01:00:5e:ff:ff:ff inclusive).
<i>mlt</i> <1-512>	Specifies the MLT ID.
<i>port</i> {slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>vlan</i> <0-4094>	Specifies the VLAN ID.

*Display Node Alias Client Information***About This Task**

Perform the following procedure to display Node Alias client information.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display Node Alias client information:

```
show nodealias client-info [detail] [hostname] [i-sid <0-16777215>]
[mac <0x00:0x00:0x00:0x00:0x00:0x00>] [mlt <1-512>] [port {slot/
port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] [vlan
<0-4094>]
```

Example

Display Node Alias client information:

```
Switch:1#show nodealias client-info
```

Node Alias - Client Database							
MAC-ADDRESS	I-SID	VLAN	INTERFACE	PROTOCOLS	ACTIVE	LAST IP ADDR LAST IPv6 ADDR	FIRST LEARNED TIME LAST LEARNED TIME
00:00:11:11:16:01	0	16	1/1	4	true	192.0.2.15 0:0:0:0:0:0:0:0	2024/06/06 12:25:30 2024/06/06 12:26:12
00:00:11:11:26:02	0	16	1/2	4	true	192.0.2.16 0:0:0:0:0:0:0:0	2024/06/06 12:25:31 2024/06/06 12:25:42
00:00:11:66:00:01	2	2	1/3	4BOR	false	198.51.100.2 0:0:0:0:0:0:0:0	2024/06/06 12:27:25 2024/06/06 12:39:08
00:00:11:66:55:01	4	4	1/3	4BOR	false	198.51.100.3 0:0:0:0:0:0:0:0	2024/06/06 12:27:30 2024/06/06 12:39:08
00:00:11:77:00:02	1002	1002	MLT-1	4BOR	true	203.0.113.23 0:0:0:0:0:0:0:0	2024/06/06 12:27:26 2024/06/06 12:39:08
00:00:11:77:00:03	1002	1002	1/2	4R	true	203.0.113.24 0:0:0:0:0:0:0:0	2024/06/06 12:31:51 2024/06/06 12:32:23
6 displayed out of 6 Total Num of Clients							
Protocols Legend: Client Capabilities 4 = IPv4, 6 = IPv6, B = BGP, V = VRRP, U = UDP, L = LLDP, D = DHCP, b = BPDU, n = NON_NT_BPDU O = OSPF, R = RIP, l = LLMNR, S = SSDP, M = MDNS							

Display detailed Node Alias client information:

```
Switch:1#show nodealias client-info detail
```

Node Alias - Client Database									
MAC-ADDRESS	I-SID	VLAN	INTERFACE	PROTOCOLS	ACTIVE	LAST IP ADDR LAST IPv6 ADDR	FIRST LEARNED TIME LAST LEARNED TIME	HOSTNAME	OS TYPE OS VERSION DEVICE TYPE
00:00:50:b8:ec:61	15999999	4048	1/1	61	false	0.0.0.0 fe80:0:0:0:bde3:2f0e:7b83:80a	09/25/24 15:07:24 09/25/24 18:06:35		
40:88:2f:dc:68:81	15999999	4048	1/1	4UD	true	10.3.33.208 0:0:0:0:0:0:0:0	09/25/24 14:40:56 09/25/24 19:03:24	A1-5420	
40:88:2f:el:c4:81	15999999	4048	1/1	4UD	true	0.0.0.0 0:0:0:0:0:0:0:0	09/25/24 14:41:09 09/25/24 19:03:31	FabricEngine40882fe~	
40:88:2f:e3:a0:81	15999999	4048	1/1	4UD	true	0.0.0.0 0:0:0:0:0:0:0:0	09/25/24 14:40:18 09/25/24 19:03:28	FabricEngine40882fe~	
4 displayed out of 426 Total Num of Clients									
Protocols Legend: Client Capabilities 4 = IPv4, 6 = IPv6, B = BGP, V = VRRP, U = UDP, L = LLDP, D = DHCP, b = BPDU, n = NON_NT_BPDU O = OSPF, R = RIP, l = LLMNR, S = SSDP, M = MDNS									
DeviceType Legend: B= Bridge, D= DOCSIS, O= Other, R= Repeater, S= Station, T= Telephone, W= WLAN, r= Router									

Variable Definitions

The following table defines parameters for the **show nodealias client-info** command.

Variable	Value
<i>detail</i>	Displays detailed information.
<i>hostname</i>	Displays the hostname.
<i>i-sid</i> <0-16777215>	Displays details for the specified I-SID.
<i>mac</i> <0x00:0x00:0x00:0x00:0x00:0x00>	Displays details for the specified MAC address.
<i>mlt</i> <1-512>	Displays details for the specified MLT.
<i>port</i> { <i>slot/port</i> [/ <i>sub-port</i>] [- <i>slot/port</i> [/ <i>sub-port</i>]][,...][<i>slot/all</i>][<i>all</i>]}	Displays details for the specified port.
<i>vlan</i> <0-4059>	Displays details for the specified VLAN.

Display Node Alias Protocol Information

About This Task

Perform the following procedure to display Node Alias protocol information.



Note

Detailed information displays for IPv4, IPv6, BGP, VRRP, and broadcast UDP protocols only.

Procedure

- To enter User EXEC mode, log on to the switch.
- Display Node Alias protocol information:

```
show nodealias protocol-info [i-sid <0-16777215>] [mac
<0x00:0x00:0x00:0x00:0x00:0x00>] [mlt <1-512>] [port {slot/port[/sub-
port] [-slot/port[/sub-port]][,...][slot/all][all]]] [protocol {bgp|
ipv4|ipv6|udp|vrrp}] [vlan <0-4094>]
```

Example

Display Node Alias protocol information:

```
Switch:1>show nodealias protocol-info
=====
Node Alias - Protocol Info Database
=====
```

MAC-ADDRESS	I-SID	VLAN	INTERFACE	PROTOCOL	PROTO INFO	LAST LEARNED TIME
00:00:11:11:16:01	0	16	1/1	IPv4	16.1.1.1	2024/06/06 12:26:12
00:00:11:66:00:01	2	2	1/3	IPv4	66.1.0.1	2024/06/06 12:39:08
00:00:11:66:00:01	2	2	1/3	BGP	20	2024/06/06 12:36:49
00:00:11:66:11:01	3	3	1/3	BGP	20	2024/06/06 12:36:49
00:00:11:66:16:01	0	16	1/1	IPv4	16.1.1.1	2024/06/06 12:34:10
00:00:11:66:55:01	4	4	1/3	BGP	20	2024/06/06 12:36:49
00:00:11:77:00:02	1002	1002	MLT-1	IPv4	77.2.0.2	2024/06/06 12:39:08
00:00:11:77:00:02	1002	1002	MLT-1	BGP	20	2024/06/06 12:36:49
00:00:11:77:00:03	1002	1002	1/2	IPv4	77.2.0.3	2024/06/06 12:32:23

00:00:11:77:00:03	0	2002	1/3	IPv4	77.3.0.3	2024/06/06 12:35:08
00:00:11:77:11:02	1003	1003	MLT-1	IPv4	77.2.1.2	2024/06/06 12:39:08
00:00:11:77:11:02	1003	1003	MLT-1	BGP	20	2024/06/06 12:36:49
00:00:11:77:11:03	1003	1003	1/2	IPv4	77.2.1.3	2024/06/06 12:32:28
00:00:11:77:11:03	0	2003	1/3	IPv4	77.3.1.3	2024/06/06 12:35:09
90:e2:ba:21:7d:10	0	16	1/1	IPv6	1016:0:0:0:0:120d:1	2024/06/06 12:23:08
90:e2:ba:21:7d:11	0	26	1/2	IPv6	2026:0:0:0:0:120d:1	2024/06/06 12:23:19

16 displayed out of 16 Total Num of Protocol-Info entries.

Variable Definitions

The following table defines parameters for the **show nodealias protocol-info** command.

Variable	Value
<i>i-sid</i> <0-16777215>	Displays details for the specified I-SID.
<i>mac</i> <0x00:0x00:0x00:0x00:0x00:0x00>	Displays details for the specified MAC address.
<i>mlt</i> <1-512>	Displays details for the specified MLT.
<i>port</i> { <i>slot/port</i> [/ <i>sub-port</i>] [- <i>slot/port</i> [/ <i>sub-port</i>]][,...][<i>slot/all</i>][<i>all</i>]}	Displays details for the specified port.
<i>protocol</i> { <i>bgp</i> <i>ipv4</i> <i>ipv6</i> <i>udp</i> <i>vrrp</i> }	Displays details for the specified protocol.
<i>vlan</i> <0-4094>	Displays details for the specified VLAN ID.

Display DHCP Fingerprint Information

About This Task

Perform the following procedure to display DHCP fingerprint information stored for DHCP clients.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display Node Alias client information:

show nodealias dhcp-fingerprint [i-sid <0-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>] [mlt <1-512>] [port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}] [vlan <0-4094>]

Example

Display DHCP fingerprint information:

Switch:1>show nodealias dhcp-fingerprint					
=====					
Node Alias - DHCP Fingerprint Info					
=====					
MAC-ADDRESS	I-SID	VLAN	INTERFACE	CODE	VALUE

00:00:00:00:00:15	0	20	cpp	12	host1
				53	2.2.2.2
				55	17,17,17,16

```
-----
                                60      extreme networks
-----
00:00:00:00:00:21    0      20    cpp    12      abracadabra
                                53      0.0.0.0
                                55      3
                                60      extreme networks
-----

2 displayed.
```

Variable Definitions

The following table defines parameters for the `show nodealias dhcp-fingerprint` command.

Variable	Value
<code>i-sid <0-16777215></code>	Displays details for the specified I-SID.
<code>mac <0x00:0x00:0x00:0x00:0x00:0x00></code>	Displays details for the specified MAC address.
<code>mlt <1-512></code>	Displays details for the specified MLT.
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/ all][all]}</code>	Displays details for the specified port.
<code>vlan <0-4094></code>	Displays details for the specified VLAN ID.

Display Ports Speeds for All VIM Ports



Note
This procedure only applies to 5520 Series and 5720 Series.

Display the configured speed on all VIM ports.



Note
Some VIMs must operate with all ports at the same speed, while others can operate with ports at different speeds. For more information, see *Fabric Engine Release Notes*. The `show sys vim-speed` command is supported only on VIMs that must operate with all ports at the same speed. An error message displays if you run the command with an unsupported VIM installed.

Before You Begin

Install the VIM before performing this procedure.

About This Task

Use this procedure to display the configured speed of all ports in a multi-port VIM.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display the speed for all of the VIM ports:

```
show sys vim-speed
```

View the Management Port Statistics

Use this procedure to view the management port statistics.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the management port statistics:
show interfaces mgmtethernet statistics

Example

View management port statistics:

```
Switch:1>enable
Switch:1#show interfaces mgmtethernet statistics
=====
Port Stats Interface
=====
```

PORT NUM	IN OCTETS	OUT OCTETS	IN PACKET	OUT PACKET
mgmt	7222116	44282	81789	586

```
=====
```

PORT NUM	IN FLOWCTRL	OUT FLOWCTRL	IN PFC	OUT PFC	OUTLOSS PACKETS
mgmt	0	0	0	0	0

```
=====
```

Display Detailed Statistics for Ports

Display detailed statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View statistics for specific ports:
show interfaces GigabitEthernet statistics verbose {slot/port[/sub-port]} [-slot/port[/sub-port]] [,...]

Example

View statistics for various ports:

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet statistics verbose
```

Please widen the terminal for optimal viewing of data.

```
=====
```

```

Port Stats Interface Extended
=====
PORT_NUM IN_UNICST OUT_UNICST IN_MULTICST OUT_MULTICST IN_BRDCST OUT_BRDCST IN_LSM OUT_LSM
-----
1/1      0          0          0          0          0          0          0          0
1/2      11421       10          0          255132      154         412502      0          0
1/3      0           0          0          0          0          0          0          0
1/4      0           0          0          0          0          0          0          0
1/5      0           0          0          0          0          0          0          0
1/6      0           0          0          0          0          0          0          0
1/7      0           0          0          0          0          0          0          0
1/8      0           0          0          0          0          0          0          0
1/9      0           0          0          0          0          0          0          0
1/10     0           0          0          0          0          0          0          0
1/11     0           0          0          0          0          0          0          0
1/12     49811       4388       13200      730098      0           0           0          0

--More-- (q = quit)

```

Variable Definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics verbose** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

NEW! Configure the Minimum Fan Speed



Note

This procedure only applies to 5720 Series.

About This Task

Perform this procedure to adjust the thermal operations of the switch. For example, if you want installed transceivers to operate at a lower temperature, increase the fan speed.

The switch monitors temperature sensors every 15 seconds. If the temperature exceeds the threshold, the fan speed increases.

If you change the minimum fan speed value, the switch can take up to 15 seconds to read the temperature sensors and then change the fan speed. If you lower the minimum fan speed by more than 20%, the switch decreases the speed by 20% every 15 seconds.



Important

Increasing the fan speed can also increase the fan noise.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Change the fan speed:

```
sys fan set-min-speed <20-100>
```
3. Verify the configuration:

```
show sys-info fan
```

Example

The following example configures the fan speed to 20%.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#sys fan set-min-speed 20
Switch:1(config)#show sys-info fan
```

Tray	Unit	Oper Speed	Oper Speed (Rpm)	Airflow Type	Status
1	1	LOW	4960	F2B	OK
1	2	LOW	4320	F2B	OK
2	1	LOW	5120	F2B	OK
2	2	LOW	4320	F2B	OK
3	1	LOW	4960	F2B	OK
3	2	LOW	4320	F2B	OK
4	1	LOW	4960	F2B	OK
4	2	LOW	4480	F2B	OK

```
Fan minimum speed: 20
```

NEW! Variable Definitions

The following table defines parameters for the **sys fan set-min-speed** command.

Variable	Value
<20-100>	Configures the minimum fan speed as a percentage between the minimum (20%) and maximum (100%) speed imposed by the platform design. The default is 20%.

View Fan Information

About This Task

View fan information to monitor the alarm status of the cooling ports in the chassis.



Note

This procedure does not apply to 5320-16P-2MXT-2X.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display fan information:

show sys-info fan

Example

View fan information:

```
Switch:1>show sys-info fan
*****
Tray      Unit      Oper Speed      Oper Speed (Rpm)  Airflow Type      Status
1         1         LOW             5782              F2B               OK
1         2         LOW             3366              F2B               OK
```

View Baseboard Management Controller Information

About This Task

View Baseboard Management Controller (BMC) status and device information.



Note

This procedure applies to 7520 Series, 7720 Series, and 7830 Series.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display BMC information:

show sys-info bmc

Example

View BMC information:

```
Switch:1>show sys-info bmc

Current BMC Info :
-----
Status                : OPERATIONAL
Device ID              : 32
Device Revision        : 1
Firmware Revision      : 3.04
IPMI Version           : 2.0
Manufacturer ID        : 7244
```

Display Power Status

About This Task

Perform the following procedure to display information about the installed power supplies.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display power supply information:
 - `show sys-info power`
 - `show sys power power-supply`

Example

Display power supply information:



Note
Switches do not report information for Voltage in, Voltage out, Current in, Current out, Power in, and Power out for fixed power supply units.

```
Switch:1#show sys-info power
```

Power Info :								
Power Supply	Type	Input Voltage	Serial Num	Part Num	Oper Status	Max Power	Fan FlowType	Hardware Revision
PS#1	AC	110/220	1234A-12345	10951	UP	715	front-back	14
Power Supply		Voltage in[V]	Voltage out[V]	Current in[A]	Current out[A]	Power in[W]	Power out[W]	
PS#1		235.75	54.81	0.24	0.55	44.50	30.00	

```
Switch:1>show sys power power-supply
```

Power Supply Information								
Power Supply	Type	Input Voltage	Serial Num	Part Num	Oper Status	Max Power	Fan FlowType	Hardware Revision
PS#1	AC	unknown	--	--	DOWN	0	front-back	--
PS#2	AC	110/220	2041U-30076	801092-00-01	UP	150	front-back	01
Power Supply		Voltage in[V]	Voltage out[V]	Current in[A]	Current out [A]	Power in[W]	Power out[W]	
PS#1		--	--	--	--	--	--	
PS#2		240.25	12.08	0.21	1.53	23.41	18.72	

Chassis operations configuration using EDM

This section provides the details to configure basic hardware and system settings using Enterprise Device Manager (EDM).

Edit System Information

About This Task

Edit system identification information, configuration file information, and perform system actions.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System** tab.
4. Enter or edit the information as required.
5. Select **Apply**.

System Field Descriptions

Use the data in the following table to use the **System** tab.

Name	Description
sysDescr	Shows the system assigned name and the software version.
sysUpTime	Shows the elapsed time since the system last started.
sysContact	Configures the contact information.
sysName	Configures the name of this device.
sysLocation	Configures the physical location of this device.
VirtualIpAddr	Configures the virtual IP address that the CPU advertises and stores in the switch configuration file.
VirtualNetMask	Configures the net mask of the virtual management IP address.
VirtualIpv6Addr	Specifies the virtual IPv6 address.
VirtualIpv6PrefixLength	Specifies the length of the virtual IPv6 address prefix (in bits).
LastChange	Displays the time since the last configuration change.
LastVlanChange	Displays the time since the last VLAN change.
LastStatisticsReset	Displays the time since the statistics counters were last reset.
LastRunTimeConfigSave	Displays the last run-time configuration saved.
DefaultRuntimeConfigFileName	Displays the default Run-time configuration file directory name.

Name	Description
ConfigFileName	Specifies the name of a new configuration file.
ActionGroup1	<p>Performs one of the following actions:</p> <ul style="list-style-type: none"> • resetCounters— Resets all statistic counters. • saveRuntimeConfig— Saves the current run-time configuration. • loadLicense— Loads a software license file to enable features. • revokeLicensePremier— Revokes the Premier license and generates a revocation certificate in XML format. • revokeLicenseMacsec— Revokes the MACsec license and generates a revocation certificate in XML format. • revokeLicense10G4P— Revokes the 4-port 10 Gbps Port License on 5320 Series and generates a revocation certificate in XML format. • revokeLicense10G8P— Revokes the 8-port 10 Gbps Port License on 5320 Series and generates a revocation certificate in XML format.
LicenseFileName	Specifies the name of the license file in the /intflash directory.
ActionGroup2	<p>Specifies the following action:</p> <ul style="list-style-type: none"> • resetIstStatCounters—Resets the IST statistic counters
ActionGroup3	<p>Can be the following action:</p> <ul style="list-style-type: none"> • flushIpRouteTbl—flushes IP routes from the routing table
ActionGroup4	<p>Can be the following action:</p> <ul style="list-style-type: none"> • softReset—resets the device without running power-on tests • resetConsole—resets the switch console • softResetCoreDump—reset with coredump
Result	Displays a message after you select Apply .
LocatorLED	Configures the system Locator LED on or off. The default is off.

Edit Chassis Information

About This Task

Edit the chassis information to make changes to chassis-wide settings.

Procedure

1. In the Device Physical View tab, select the device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **Chassis** tab.
5. Edit the necessary options.
6. Select **Apply**.

Chassis Field Descriptions

Use the data in the following table to use the **Chassis** tab.

Name	Description
Type	Specifies the chassis type.
ModelName	Specifies the chassis model name.
BrandName	Specifies the chassis brand name.
PartNumber	Specifies the device part number.
SerialNumber	Specifies a unique chassis serial number.
HardwareRevision	Specifies the current hardware revision of the device chassis.
NumSlots	Specifies the number of slots available in the chassis.
NumPorts	Specifies the number of ports currently installed in the chassis.
BaseMacAddr	Specifies the starting point of the block of MAC addresses used by the switch for logical and physical interfaces.
MacAddrCapacity	Specifies the number of routable MAC addresses based on the BaseMacAddr.
AutoRecoverDelay	Specifies the time interval, in seconds, after which auto-recovery runs on ports to clear actions taken by CP Limit or link flap. The default is 30.
MTUSize	Configures the maximum transmission unit size. The default is 1950 bytes.
MgidUsageVlanCurrent	Number of MGIDs for VLANs currently in use.
MgidUsageVlanRemaining	Number of remaining MGIDs for VLANs.
MgidUsageMulticastCurrent	Number of MGIDs for multicast currently in use.
MgidUsageMulticastRemaining	Number of remaining MGIDs for multicast.
DdmMonitor	Enables or disables the monitoring of the DDM. When enabled, the user gets the internal performance condition (temperature, voltage, bias, Tx power and Rx power) of the SFP/XFP. The default is disable.

Name	Description
DdmMonitorInterval	Configures the DDM monitor interval in the range of 5 to 60 in seconds. If any alarm occurs, the user gets the log message before the specific interval configured by the user. The default value is 5 seconds.
DdmTrapSend	Enables or disables the sending of trap messages. When enabled, the trap message is sent to the Device manager, any time the alarm occurs. The default is enable.
DdmAlarmPortdown	Sets the port down when an alarm occurs. When enabled, the port goes down when any alarm occurs. The default is disable.

View Physical Entities

Perform this procedure to view information about the functional components of the switch.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Entity**.

Physical Entities Field Descriptions

Use the data in the following table to use the **Physical Entities** tab.

Name	Description
Index	Indicates the index of the entry.
Descr	Indicates the name of the manufacturer for the physical entity.
VendorType	Indicates the vendor-specific hardware type for the physical entity. Because there is no vendor-specifier registration for this device, the value is 0.
ContainedIn	Indicates the index value for the physical entity which contains this physical entity. A value of zero indicates that this physical entity is not contained in any other physical entity.
Class	Indicates the general hardware type of the physical entity. The value is configured to the standard enumeration value that indicates the general class of the physical entity.
ParentRelPos	Indicates the relative position of the child component among the sibling components.

Name	Description
Name	<p>Indicates the name of the component, as assigned by the local device, and that is suitable to use in commands you enter on the console of the device. Depending on the physical component naming syntax of the device, the name can be a text name such as console, or a component number such as port or module number.</p> <p>If there is no local name, there is no value.</p>
HardwareRev	<p>Indicates the vendor-specific hardware revision string for the physical entity.</p> <p>If no specific hardware revision string is associated with the physical component, or if this information is unknown, then this object contains a zero-length string, or there is no value.</p> <p>If there is no information available, there is no value.</p>
FirmwareRev	<p>Indicates the vendor-specific firmware revision string for the physical entity.</p> <p>If no specific firmware programs are associated with the physical component, or if this information is unknown, then this object contains a zero-length string, or there is no value.</p> <p>If there is no information available, there is no value.</p>
SoftwareRev	<p>Indicates the vendor-specific software revision string for the physical entity.</p> <p>If no specific software programs are associated with the physical component, or if this information is unknown, then this object contains a zero-length string, or there is no value.</p> <p>If there is no information available, there is no value.</p>
SerialNum	<p>Indicates the vendor-specific serial number string for the physical entity. The value is the serial number string printed on the component, if present.</p> <p>If there is no information available, there is no value.</p>

Name	Description
MfgName	<p>Indicates the name of the manufacturer of the physical component. The value is the manufacturer name string printed on the component, if present.</p> <p>If the manufacturer name string associated with the physical component is unknown, then this object contains a zero-length string.</p> <p>If there is no information available, there is no value.</p>
ModelName	<p>Indicates the vendor-specific model name identifier string associated with the physical component. The value is the part number which is printed on the component.</p> <p>If the model name string associated with the physical component is unknown, then this object contains a zero-length string.</p>
Alias	<p>Indicates an alias name for the physical entity that is specified by a network manager, and provides a nonvolatile handle for the physical entity.</p> <p>The software supports read-only and provides values for the port interface only.</p>
AssetID	<p>Indicates a user-assigned asset tracking identifier for the physical entity. This value is specified by a network manager, and provides nonvolatile storage of this information.</p> <p>Because this object is not supported, there is no value.</p>
IsFRU	<p>Indicates whether or not the physical entity is considered a field replaceable unit.</p> <ul style="list-style-type: none"> • If the value is <code>true (1)</code>, then the component is a field replaceable unit. • If the value is <code>false (2)</code>, then the component is permanently contained within a field replaceable unit.
MfgDate	<p>Indicates the manufacturing date of the managed entity. If the manufacturing date is unknown, then the value is '0000000000000000' H.</p>
Uris	<p>Indicates additional identification information about the physical entity.</p> <p>Uris is not supported, therefore there is no value.</p>

View Entity Aliases

About This Task

Perform this procedure to view the entity aliases on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Entity**.
3. Select the **Alias** tab.

Alias Field Descriptions

Use the data in the following table to use the **Alias** tab.

Name	Description
Index	The index of the entry
LogicalIndexOrZero	The index of the entry. The value of this object identifies the logical entity that defines the naming scope for the associated instance of the Mapping Identifier object. This is always 0.
MappingIdentifier	The value of this object identifies a particular conceptual row associated with the indicated Physical Index and Logical Index pair. Because only physical ports are modeled in this table, only entries that represent interfaces or ports are allowed. If an ifEntry exists on behalf of a particular physical port, then this object should identify the associated ifEntry. This is the OID of ifIndex.Port.

View Entity Child Indexes

About This Task

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Entity**.
3. Select the **Child Index** tab.

Child Index *Field Descriptions*

Use the data in the following table to use the **Child Index** tab.

Name	Description
Index	Indicates the index of the entry.
ChildIndex	The index of the entry. The value of Physical Index for the contained physical entity.

Configure System Flags

About This Task

Configure the system flags to enable or disable flags for specific configuration settings.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System Flags** tab.
4. Select the system flags you want to activate.
5. Clear the system flags you want to deactivate.
6. Select **Apply**.

**Important**

After you change certain configuration parameters, you must save the changes to the configuration file.

System Flags *Field Descriptions*

Use the data in the following table to use the **System Flags** tab.

Name	Description
EnableAccessPolicy	Activates access policies. The default is disabled.
ForceTrapSender	Configures circuitless IP as a trap originator. The default is disabled.
ForceIpHdrSender	If you enable Force IP Header Sender, the system matches the IP header source address with SNMP header sender networks. The default is disabled.
AuthSuccessTrapEnable	Enables the system to send the authentication success trap, rcnAuthenticationSuccess. The default is disabled.
MrouteStrLimit	Enable or disable Mroute stream limit in system. The default is disabled.
DataPathFaultShutdownEnable	Enable or disable data path fault shutdown. The default is enabled.

Name	Description
PingTracerouteContextType	Configures the default context for executing ping commands and traceroute commands. The default is grt.
UdpSrcByVirtualIpEnable	Enables or disables virtual IP as the User Datagram Protocol (UDP) source. The default is disabled.
ForceTopologyIpFlagEnable	Activates or disables the flag that configures the CLIP ID as the topology IP. Values are true or false. The default is disabled.
CircuitlessIpId	Uses the CLIP ID as the topology IP.
PtpEnable Note: Exception: does not apply to 4220 Series and 7830 Series.	Enables or disables PTPv2 on a switch with a valid license. The default is enabled if a the switch is in Zero Touch Deployment mode.

Configure Channelization

Use this procedure to enable or disable channelization on a port. Channelization configures the port to operate as four channels, or ports.



Note

The switch automatically channelizes a port if the port operates in Auto-sense mode and you insert a QSFP+ to SFP+ or QSFP28 to SFP28 adapter, or a breakout cable.

About This Task



Important

Enabling or disabling channelization resets the port QoS configuration to default values.

Procedure

1. In the Device Physical View tab, select a port that supports channelization.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Channelization** tab.
5. To enable channelization on the port, select **enable**.
6. Select **Apply**. Alternatively, you can right-click the port on the Device Physical View tab, and then select **Channelization Enable**.
7. To disable channelization on a port, select the first sub-port for the corresponding port, slot/port/1.
8. In the navigation pane, expand **Configuration > Edit > Port**.
9. Select **General**.
10. Select the **Channelization** tab.

11. To disable channelization on the port, select **disable**. This action will disable the four sub-ports.
12. Select **Apply**. Alternatively, you can right-click the port on the Device Physical View tab, and then select **Channelization Disable**.

Channelization Field Descriptions

Use the data in the following table to use the **Channelization** tab.

Name	Description
Channelization	Specifies whether channelization is enabled or disabled on the selected port. The default is disable .

Configure Basic Port Parameters

Configure options for port operations.

About This Task

If you select more than one port, the format of the tab changes to a table-based tab.

When you use 1 Gigabit Ethernet SFP transceivers, the software disables Auto-Negotiation on the port. If you use 1 Gbps SFP transceivers, the remote end must also have Auto-Negotiation disabled.

Procedure

1. In the Device Physical View tab, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. Configure the fields as required.

10/100BASE-TX ports do not consistently auto-negotiate with older 10/100BASE-TX equipment. You can sometimes upgrade the older devices with new firmware or driver revisions. If an upgrade does not enable auto-negotiation to correctly identify the link speed and duplex settings, you can manually configure the settings for the link in question.

Check the Extreme Networks website for the latest compatibility information.

6. Select **Apply**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Index	Displays the index of the port, written in the slot/port[/sub-port] format.
Name	Configures the name of the port.

Name	Description
Descr	Displays the description of the port. A textual string containing information about the interface.
Type	Displays the type of connector plugged in the port.
Mtu	Displays the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram which can be sent or received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.
PhysAddress	Displays the physical address of the port. The address of the interface at the protocol layer immediately below the network layer in the protocol stack. For interfaces which do not have such an address, for example, a serial line, this object should contain an octet string of zero length.
VendorDescr	Displays the vendor of the connector plugged in the port.
DisplayFormat	Identifies the slot and port numbers (slot/port). If the port is channelized, the format also includes the sub-port in the format slot/port/sub-port
AdminStatus	Configures the port as enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
OperStatus	Displays the current status of the port. The status includes enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
ShutdownReason	Indicates the reason for a port state change.
LastChange	Displays the timestamp of the last change.
LinkTrap	Enable or disable link trapping.
AutoNegotiate	Enables or disables Auto-Negotiation for this port. The default Auto-Negotiation behavior depends on the switch model and transceiver type.

Name	Description
AutoNegAd	<p>Specifies the port speed and duplex abilities to advertise during link negotiation.</p> <p>Supported speeds and duplex modes vary, depending on your hardware.</p> <p>The abilities specified in this object are only used when auto-negotiation is enabled on the port. If all bits in this object are disabled, and auto-negotiation is enabled on the port, then the physical link process on the port is disabled (if hardware supports this ability).</p> <p>Any change to this configuration restarts the auto-negotiation process, which has the same effect as physically unplugging and reattaching the cable attached to the port.</p> <p>If you select default, all capabilities supported by the hardware are advertised.</p>
AdminDuplex	Configures the administrative duplex setting for the port.
OperDuplex	Indicates the operational duplex setting for the port.
AdminSpeed	Configures the administrative speed for the port.
OperSpeed	Indicates the operational speed for the port.
QoSLevel	Selects the Quality of Service (QoS) level for this port. The default is level1.
DiffServ	Enables the Differentiated Service feature for this port. The default is disabled.
Layer3Trust	Configures if the system should trust Layer 3 packets coming from access links or core links only. The default is core.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled (selected) or disabled (cleared) on the port. The default is disabled (clear).
MLtId	Shows the MLT ID associated with this port. The default is 0.
Locked	Shows if the port is locked. The default is unlocked.
UnknownMacDiscard	Discards packets that have an unknown source MAC address, and prevents other ports from sending packets with that same MAC address as the destination MAC address. The default is disabled.

Name	Description
DirectBroadcastEnable	<p>Specifies that an Isolated Routing Port (IRP) can forward directed broadcast traffic. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcast on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling this function protects a host from possible denial of service (DoS) attacks.</p> <p>With the feature enabled, the Control Processor (CP) does not receive a copy of the directed broadcast. As a result, the system does not respond to a subnet broadcast ping sent from a remote subnet.</p> <p>The default is disabled.</p>
OperRouting	Shows the routing status of the port.
HighSecureEnable	Enables or disables the high secure feature for this port.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
FlexUniEnable	Enables Flex UNI on the port. The default is disabled.
EgressRateLimitState	Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the port. The default is disabled.
EgressRateLimit	<p>Specifies the egress rate limit in Kbps. You cannot configure the egress shaper rate to exceed the port capability.</p> <p>If you configure this value to 0, shaping is disabled on the port.</p>
TxFlowControl	<p>Configures if the port sends pause frames. By default, an interface does not send pause frames.</p> <p>You must also enable the flow control feature globally before an interface can send pause frames.</p>
TxFlowControlOperState	Shows the operational state of flow control.
BpduGuardTimerCount	Shows the time, starting at 0, since the port became disabled. When the BpduGuardTimerCount reaches the BpduGuardTimeout value, the port is enabled. Displays in 1/100 seconds.

Name	Description
BpduGuardTimeout	Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.
BpduGuardAdminEnabled	Enables BPDU Guard on the port. The default is disabled.
ForwardErrorCorrection	Configures one of the following options for Forward Error Correction (FEC) on the port: <ul style="list-style-type: none"> • CL 91 • CL 108 • CL 74 • disable • auto The disable option disables this configuration on the port.
ForwardErrorCorrectionApplicability	Displays whether FEC is applicable on the interface.
OperAutoNegotiate	Shows the operational state of Auto-Negotiation.
OperForwardErrorCorrection	Shows the negotiated operational FEC clause. If the value is off, the port supports FEC and is up but not configured for FEC. If the value is notApplicable, the port does not support FEC. If the value is unknown, the port supports FEC but is down.
Action	Performs one of the following actions on the port <ul style="list-style-type: none"> • none - none of the following actions • flushMacFdb - flush the MAC forwarding table • flushArp - flush the ARP table • flushIp - flush the IP route table • flushAll - flush all tables • triggerRipUpdate — manually triggers a RIP update The default is none.
Result	Displays the result of the selected action. The default is none.
AutoSense	Enables or disables Auto-sense on the specific port. The default value is disabled for existing configurations but enabled for new Zero Touch Fabric Configuration deployments.

Name	Description
AutoSenseKeepAutoConfig	Retains the Auto-sense configuration if you disable Auto-sense on the port. The dynamic configuration becomes a manual configuration and is visible in the show running-config output.
CustomAutoNegAdOrigin	Specifies the origin of Custom Auto Negotiation Advertisements (CANA) configuration on the port. The supported values are: <ul style="list-style-type: none"> config - Set by the user. radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
BpduGuardOrigin	Specifies the origin of BPDU Guard configuration on the port. The supported values are: <ul style="list-style-type: none"> config - Set by the user. radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
AutoSenseState	Displays the Auto-sense port state.
LinkDebounce	Specifies the extended debounce timer on the port. The value 0 milliseconds disables debounce time. The default value is 1000. Note: This field does not apply to Auto-sense ports. To configure Link Debounce on Auto-sense ports, see Configure Link Debounce for Auto-sense Ports on page 65.
AutoSenseDataIsid	Specifies the Auto-sense data I-SID per port.
AutoSenseWaitInterval	Specifies the Auto-sense Wait Interval for the port. The default is 35.
NodeAliasEnable	Enables or disables Node Alias on this specific port. The default is enabled on Auto-sense ports.

Name	Description
NodeAliasClear	Clears Node Alias on this specific port.
IpfixEnable	Enables or disables IPFIX on the port. The default is true, which means enabled. If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.

Configure Basic Parameters on an Extreme Integrated Application Hosting Port



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to configure basic parameters on Extreme Integrated Application Hosting (IAH) ports, for example, auto negotiation, QoS level, and remote monitoring.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **Interface** tab.
4. In **Name**, type a name for the IAH port.
5. Configure the parameters as required.
6. Select **Apply**.

Interface Field Descriptions

Use data in the following table to use the **Interface** tab.

Name	Description
Index	Specifies the index of the Extreme Integrated Application Hosting (IAH) port, written in the slot/port[/sub-port] format.
Name	Specifies the name of the IAH port.
Descr	Specifies the information about the interface.
Type	Specifies the type of connector plugged in the IAH port.

Name	Description
Mtu	Specifies the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram which can be sent or received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.
PhysAddress	Specifies the physical address of the IAH port. The address of the interface at the protocol layer immediately below the network layer in the protocol stack. For interfaces which do not have such an address (like a serial line), this object should contain an octet string of zero length.
VendorDescr	Specifies the vendor of the connector plugged in the IAH port.
DisplayFormat	Specifies the slot and port numbers (slot/port).
AdminStatus	Specifies the operational status of the IAH port. The testing state indicates that no operational packets can be passed.
OperStatus	Specifies the current status of the IAH port. The testing state indicates that no operational packets can be passed.
LicenseControlStatus	Specifies the IAH port license status.
ShutdownReason	Specifies the reason for the IAH port state change.
LastChange	Specifies the timestamp of the last change.
LinkTrap	Enables or disables link trapping. The default is enabled.
AutoNegotiate	Enables or disables auto-negotiation for the IAH port. The default is true (enabled).

Name	Description
AutoNegAd	<p>Specifies the port speed and duplex abilities to be advertised during link negotiation.</p> <p>The abilities specified in this object are only used when auto-negotiation is enabled on the IAH port. If all bits in this object are disabled, and auto-negotiation is enabled on the IAH port, then the physical link process on the IAH port will be disabled (if hardware supports this ability).</p> <p>Any change in the value of this bit map will force the switch to restart the auto-negotiation process.</p> <p>The capabilities being advertised are either all the capabilities supported by the hardware or the user-configured capabilities, which is a subset of all the capability supported by hardware.</p> <p>By default, all capabilities supported by the hardware are enabled.</p>
AdminDuplex	Specifies the administrative duplex setting for the IAH port.
OperDuplex	Specifies the operational duplex setting for the IAH port.
AdminSpeed	Specifies the administrative speed for the IAH port.
OperSpeed	Specifies the operational speed for the IAH port.
QoSLevel	Specifies the Quality of Service (QoS) level for the IAH port. The default is level1.
DiffServ	Enables the Differentiated Service feature for the IAH port. The default is enabled.
Layer3Trust	Specifies if the system should trust Layer 3 packets coming from access links or core links only. The default is core.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled or disabled. The default is disabled.
MLTid	Specifies the MLT ID associated with the IAH port. The default is 0.
Locked	Specifies if the IAH port is locked. The default is false.

Name	Description
UnknownMacDiscard	Enables the functionality to discard packets with an unknown source MAC address, and prevents the other IAH port from sending packets with the same MAC address as the destination MAC address. The default is disabled.
DirectBroadcastEnable	<p>Specifies if the IAH port forwards direct broadcast traffic.</p> <p>A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcast on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling this function protects a host from possible denial of service (DoS) attacks.</p> <p>With the feature enabled, the Control Processor (CP) does not receive a copy of the directed broadcast. As a result, the system does not respond to a subnet broadcast ping sent from a remote subnet.</p> <p>The default is disabled.</p>
OperRouting	Specifies the routing status of the IAH port. The default is disabled.
HighSecureEnable	Enables or disables the high secure feature for the IAH port. The default is disabled.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the IAH port. The default is disabled.
FlexUniEnable	Enables or disables Flex UNI on the IAH port. The default is disabled.
EgressRateLimitState	Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the IAH port. The default is disabled.
EgressRateLimit	<p>Specifies the egress rate limit in Kbps. Different hardware platforms provide different port speeds.</p> <p>The default is 0.</p>
TxFlowControl	<p>Specifies if the IAH port is sending pause frames. The default is disabled.</p> <p>Note: You must enable the flow control feature globally.</p>

Name	Description
TxFlowControlOperState	Specifies the operational state of flow control.
BpduGuardTimerCount	Specifies the duration since when the IAH port is disabled. When the BpduGuardTimerCount reaches the BpduGuardTimeout value, the IAH port is enabled.
BpduGuardTimeout	Specifies the time (in seconds) for the IAH port-state recovery. After the IAH port is disabled by the BPDU guard, the IAH port remains in the disabled state until this timer expires. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.
BpduGuardAdminEnabled	Enables or disables BPDU Guard on the IAH port. The default is disabled.
ForwardErrorCorrection	Configures one of the following options for Forward Error Correction (FEC) on the IAH port: <ul style="list-style-type: none"> • CL 91 • CL 108 • CL 74 • disable • auto The disable option disables this configuration on the port.
ForwardErrorCorrectionApplicability	Displays whether FEC is applicable on the interface.
OperAutoNegotiate	Shows the operational state of Auto-Negotiation.
OperForwardErrorCorrection	Shows the negotiated operational FEC clause. If the value is off, the port supports FEC and is up but not configured for FEC. If the value is notApplicable, the port does not support FEC. If the value is unknown, the port supports FEC but is down.

Name	Description
Action	<p>Specifies the following actions on the IAH port:</p> <ul style="list-style-type: none"> • none - no action. • flushMacFdb - flush the MAC forwarding table. • flushArp - flush the ARP table. • flushIp - flush the IP route table. • flushAll - flush all tables. • triggerRipUpdate - manually triggers a RIP update. • clearLoopDetectAlarm - clears the loop detection alarm on the IAH port. <p>The default is none.</p>
Result	<p>Specifies the result of the selected action.</p> <p>The default is none.</p>
AutoSense	<p>Enables or disables Auto-sense on the specific port. The default value is disabled for existing configurations but enabled for new Zero Touch Fabric Configuration deployments.</p>
AutoSenseKeepAutoConfig	<p>Retains the Auto-sense configuration if you disable Auto-sense on the port. The dynamic configuration becomes a manual configuration and is visible in the show running-config output.</p>
AutoSenseState	<p>Displays the Auto-sense port state.</p>
CustomAutoNegAdOrigin	<p>Specifies the origin of Custom Auto Negotiation Advertisements (CANA) configuration on the port. The supported values are:</p> <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
BpduGuardOrigin	<p>Specifies the origin of BPDU Guard configuration on the port. The supported values are:</p> <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.

Configure IEEE 802.3X Pause Frame Transmit

Configure IEEE 802.3X Pause frame transmit to eliminate or minimize packet loss.

About This Task

By default, flow control mode is disabled. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames.



Note

The 7830 Series can process received pause frames but it cannot generate or send pause frames. When a port receives a pause frame, the switch reduces the traffic rate on that port but it does not send pause frames towards the ports that cause the over subscription.

By default, an interface does not send pause frames.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **Boot Config** tab.
4. For EnableFlowControlMode, select **enable**.
5. Select **Apply**.
6. Save the switch configuration.
7. Reboot the chassis, and log in again.
8. In the Device Physical View, select a port or ports.
9. In the navigation pane, expand **Configuration > Edit > Port**.
10. Select **General**.
11. Select the **Interface** tab.
12. For TxFlowControl, select **enable** to enable the interface to generate pause frames.
13. Select **Apply**.

View the Boot Configuration

About This Task

View the boot configuration to determine the software version, as well as view the source from which the switch last started.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **Boot Config** tab.

Boot Config *Field Descriptions*

Use the data in the following table to use the **Boot Config** tab.

Name	Description
SwVersion	Specifies the software version that currently runs on the switch.
LastRuntimeConfigSource	Specifies the last source for the run-time image.
PrimaryConfigSource	Specifies the primary configuration source.
PrimaryBackupConfigSource	Specifies the backup configuration source to use if the primary does not exist.

Name	Description
EnableFactoryDefaultsMode	<p>Specifies whether the switch uses the factory default settings after you reboot the switch. The options are:</p> <ul style="list-style-type: none"> • false: The node does not use factory default settings. • fabric: This mode is not supported. • noFabric: The node uses the factory default mode settings. <p>Removes the configuration files, primary and secondary config file names, all user account passwords, digital certificates, IKE/OSPF/IS-IS keys, and SNMP communities. All ports are disabled and assigned to VLAN 1. License files are not removed.</p> <ul style="list-style-type: none"> • config-only <p>Boots the switch with a blank configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities. All ports are disabled and assigned to VLAN 1.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p> <ul style="list-style-type: none"> • reset-all-files: <p>Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset. If required, you must revoke the license file. The 30-day factory license does not apply to 4220 Series.</p> <p>Important:</p> <p>Only CLI provides two options to achieve the same behavior; there is no option other than reset-all-files in EDM.</p> <ul style="list-style-type: none"> • zero-touch:

Name	Description
	<p>Boots the switch with a default configuration that enables Auto-sense. This parameter resets secure files but keeps the security mode and performs a ZTP+ configuration after reboot. License files are not removed.</p> <ul style="list-style-type: none"> • zero-touch-config-only <p>Boots the switch with a blank configuration that enables Auto-sense and performs a ZTP+ configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities.</p> <p>Note: Feature support can vary across platforms. For more information, see <i>Fabric Engine and VOSS Feature Matrix</i>.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p> <p>The default value is false. This flag is automatically reset to the default value after the switch restarts.</p> <p>If you select an option but then want to select a different option and have not yet rebooted the switch, you must disable the first option before you select the second option.</p>
EnableDebugMode	<p>Enabling the debugmode allows a user to enable TRACE on any port by prompting the selection on the console during boot up. This allows the user start trace for debugging earlier on specified port. It only works on console connection. The default value is disabled.</p> <p>Important: Do not change this parameter.</p>

Name	Description
EnableRebootOnError	Activates or disables automatic reboot on a fatal error. The default value is activated. Important: Do not change this parameter.
EnableTelnetServer	Activates or disables the Telnet server service. The default value is disabled.
EnableFtpServer	Activates or disables the FTP server on the switch. The default value is disabled. To enable FTP, ensure that the TFTPD flag is disabled.
EnableTftpServer	Activates or disables Trivial File Transfer Protocol server service. The default value is disabled.
EnableSshServer	Activates or disables the SSH server service. The default value is disabled.
EnableSpbmConfigMode	Enables you to configure SPB and IS-IS, but you cannot configure PIM and IGMP either globally or on an interface. The boot flag is enabled by default.
EnableIpv6Mode Note: Exception: does not apply to 4220 Series	Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits. This flag is disabled by default.
EnableEnhancedsecureMode	Enables or disables the enhanced secure mode. Select either jitc or non-jitc to enable the enhanced secure mode in one of these sub-modes. The default is disabled. Note: As a best practice, enable the enhanced secure mode in the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.
EnableUrpMode Note: Exception: does not apply to 4220 Series	Enables Unicast Reverse Path Forwarding (uRPF) globally. You must enable uRPF globally before you configure it on a port or VLAN. The default is disabled.

Name	Description
EnableVxlanGwFullInterworkingMode Note: Exception: only supported on 7520 Series and 7720 Series.	Enables VXLAN Gateway in Full Interworking Mode, which supports SPB, SMLT, and vIST. By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes by enabling this boot configuration flag. In Base Interworking Mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.
EnableFlowControlMode	Enables or disables flow control globally. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames. The default is disabled.

Name	Description
<p>AdvancedFeatureBwReservation</p> <p>Note: Exception: high is supported on 4220 Series, 5720 Series, 7520 Series, and 7720 Series . Exception: low is not supported on 5720 Series. Exception: vim is only supported on 5520 Series and 5720 Series.</p>	<p>Enables the switch to support advanced features by reserving ports as loopback ports. When disabled, you can use all ports on the switch, but advanced features do not work.</p> <p>When you select low or high, the corresponding bandwidth of the reserved Ethernet ports is allocated to the loopback required for advanced feature functionality.</p> <p>When you select vim, the bandwidth is allocated from the unused vim slot capacity; you cannot select this parameter if the VIM is present.</p> <p>The default varies depending on the platform:</p> <ul style="list-style-type: none"> • The default for 4220 Series is enabled with high level. • The default for 5320 Series, 5420 Series, 7520 Series, and 7720 Series is enabled with low level. • The default for 5520 Series is enabled with vim level if Versatile Interface Module (VIM) is not installed, else low level is enabled. • The default for 5720 Series is enabled with vim level if Versatile Interface Module (VIM) is not installed, else high level is enabled. <p>If you change this parameter, you must restart the switch.</p>
<p>EnableDvrLeafMode</p> <p>Note: Exception: does not apply to 4220 Series</p>	<p>Enables the switch to be configured as a DvR Leaf.</p> <p>When enabled, you cannot configure the switch to operate as a DvR Controller.</p>
<p>EnablevrfScaling</p> <p>Note: Exception: does not apply to 4220 Series</p>	<p>Changes the maximum number of VRFs and Layer 3 VSNs that the switch supports. If you select this check box, the maximum number increases. The default is disabled.</p> <p>Important: If you select both this check box and the EnableSpbmConfigMode check box, the switch reduces the number of configurable VLANs. For more information about maximum scaling numbers, see <i>Fabric Engine Release Notes</i>.</p>

Name	Description
EnableSyslogRfc5424Format	Enables or disables the RFC 5424 syslog format. The default is enabled. If the pre-existing configuration file is for a release prior to this enhancement, then the flag is disabled automatically.
NniMstp Note: Exception: does not apply to 4220 Series	Enables MSTP, and allows non SPBM B-VLAN configuration on SPBM network-to-network interface (NNI) ports. The default is disabled. Note: Spanning Tree is disabled on all SPBM NNIs. You cannot add an SPBM NNI port or MLT port to any non SPBM B-VLAN.
EnableIpv6EgressFilterMode Note: Exception: does not apply to 4220 Series	Enables IPv6 egress filters. The default is disabled. If you change this parameter, you must restart the switch. For 5320 Series, 5420 Series, and 5720 Series platforms, EnableIpv6EgressFilterMode and EnableMacsec are mutually exclusive.
EnableMacsec Note: Exception: only required for MACsec supporting models of 5320 Series and 5420 Series.	Enables Media Access Control Security (MACsec) mode globally. To enable MACsec mode, you must configure the boot flag. EnableIpv6EgressFilterMode and EnableMacsec are mutually exclusive.
EnableSpbmNodeScaling Note: Exception: only applies to 5320 Series (not 5320-16P-2MXT-2X or 5320-24T-4X-XT) and 5420 Series.	Enables the switch to increase the number of supported SPB nodes per area. By default, the switch supports up to 350 SPB nodes per area. The default is disabled. If you change this parameter, you must restart the switch.
Slot	Specifies the slot number.
TftpHash	Enables TFTP hashing.
TftpRetransmit	Set TFTP retransmit timeout counter.
TftpTimeout	Set TFTP timeout counter.
User	Configure host user.
Password	Configure host password.

Configure Boot Flags

About This Task

Change the boot configuration to determine the services available after the system starts.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Boot Config** tab.
3. Select the services you want to enable.
4. Select **Apply**.

Boot Config Field Descriptions

Use the data in the following table to use the **Boot Config** tab.

Name	Description
SwVersion	Specifies the software version that currently runs on the switch.
LastRuntimeConfigSource	Specifies the last source for the run-time image.
PrimaryConfigSource	Specifies the primary configuration source.
PrimaryBackupConfigSource	Specifies the backup configuration source to use if the primary does not exist.

Name	Description
EnableFactoryDefaultsMode	<p>Specifies whether the switch uses the factory default settings after you reboot the switch. The options are:</p> <ul style="list-style-type: none"> • false: The node does not use factory default settings. • fabric: This mode is not supported. • noFabric: The node uses the factory default mode settings. <p>Removes the configuration files, primary and secondary config file names, all user account passwords, digital certificates, IKE/OSPF/IS-IS keys, and SNMP communities. All ports are disabled and assigned to VLAN 1. License files are not removed.</p> <ul style="list-style-type: none"> • config-only <p>Boots the switch with a blank configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities. All ports are disabled and assigned to VLAN 1.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p> <ul style="list-style-type: none"> • reset-all-files: <p>Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset. If required, you must revoke the license file. The 30-day factory license does not apply to 4220 Series.</p> <p>Important:</p> <p>Only CLI provides two options to achieve the same behavior; there is no option other than reset-all-files in EDM.</p> <ul style="list-style-type: none"> • zero-touch:

Name	Description
	<p>Boots the switch with a default configuration that enables Auto-sense. This parameter resets secure files but keeps the security mode and performs a ZTP+ configuration after reboot. License files are not removed.</p> <ul style="list-style-type: none"> • zero-touch-config-only <p>Boots the switch with a blank configuration that enables Auto-sense and performs a ZTP+ configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities.</p> <p>Note: Feature support can vary across platforms. For more information, see <i>Fabric Engine and VOSS Feature Matrix</i>.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p> <p>The default value is false. This flag is automatically reset to the default value after the switch restarts.</p> <p>If you select an option but then want to select a different option and have not yet rebooted the switch, you must disable the first option before you select the second option.</p>
EnableDebugMode	<p>Enabling the debugmode allows a user to enable TRACE on any port by prompting the selection on the console during boot up. This allows the user start trace for debugging earlier on specified port. It only works on console connection. The default value is disabled.</p> <p>Important: Do not change this parameter.</p>

Name	Description
EnableRebootOnError	Activates or disables automatic reboot on a fatal error. The default value is activated. Important: Do not change this parameter.
EnableTelnetServer	Activates or disables the Telnet server service. The default value is disabled.
EnableFtpServer	Activates or disables the FTP server on the switch. The default value is disabled. To enable FTP, ensure that the TFTPD flag is disabled.
EnableTftpServer	Activates or disables Trivial File Transfer Protocol server service. The default value is disabled.
EnableSshServer	Activates or disables the SSH server service. The default value is disabled.
EnableSpbmConfigMode	Enables you to configure SPB and IS-IS, but you cannot configure PIM and IGMP either globally or on an interface. The boot flag is enabled by default.
EnableIpv6Mode Note: Exception: does not apply to 4220 Series	Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits. This flag is disabled by default.
EnableEnhancedsecureMode	Enables or disables the enhanced secure mode. Select either jitc or non-jitc to enable the enhanced secure mode in one of these sub-modes. The default is disabled. Note: As a best practice, enable the enhanced secure mode in the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.
EnableUrpMode Note: Exception: does not apply to 4220 Series	Enables Unicast Reverse Path Forwarding (uRPF) globally. You must enable uRPF globally before you configure it on a port or VLAN. The default is disabled.

Name	Description
EnableVxlanGwFullInterworkingMode Note: Exception: only supported on 7520 Series and 7720 Series.	Enables VXLAN Gateway in Full Interworking Mode, which supports SPB, SMLT, and vIST. By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes by enabling this boot configuration flag. In Base Interworking Mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.
EnableFlowControlMode	Enables or disables flow control globally. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames. The default is disabled.

Name	Description
<p>AdvancedFeatureBwReservation</p> <p>Note: Exception: high is supported on 4220 Series, 5720 Series, 7520 Series, and 7720 Series . Exception: low is not supported on 5720 Series. Exception: vim is only supported on 5520 Series and 5720 Series.</p>	<p>Enables the switch to support advanced features by reserving ports as loopback ports. When disabled, you can use all ports on the switch, but advanced features do not work.</p> <p>When you select low or high, the corresponding bandwidth of the reserved Ethernet ports is allocated to the loopback required for advanced feature functionality.</p> <p>When you select vim, the bandwidth is allocated from the unused vim slot capacity; you cannot select this parameter if the VIM is present.</p> <p>The default varies depending on the platform:</p> <ul style="list-style-type: none"> • The default for 4220 Series is enabled with high level. • The default for 5320 Series, 5420 Series, 7520 Series, and 7720 Series is enabled with low level. • The default for 5520 Series is enabled with vim level if Versatile Interface Module (VIM) is not installed, else low level is enabled. • The default for 5720 Series is enabled with vim level if Versatile Interface Module (VIM) is not installed, else high level is enabled. <p>If you change this parameter, you must restart the switch.</p>
<p>EnableDvrLeafMode</p> <p>Note: Exception: does not apply to 4220 Series</p>	<p>Enables the switch to be configured as a DvR Leaf.</p> <p>When enabled, you cannot configure the switch to operate as a DvR Controller.</p>
<p>EnablevrfScaling</p> <p>Note: Exception: does not apply to 4220 Series</p>	<p>Changes the maximum number of VRFs and Layer 3 VSNs that the switch supports. If you select this check box, the maximum number increases. The default is disabled.</p> <p>Important: If you select both this check box and the EnableSpbmConfigMode check box, the switch reduces the number of configurable VLANs. For more information about maximum scaling numbers, see <i>Fabric Engine Release Notes</i>.</p>

Name	Description
EnableSyslogRfc5424Format	Enables or disables the RFC 5424 syslog format. The default is enabled. If the pre-existing configuration file is for a release prior to this enhancement, then the flag is disabled automatically.
NniMstp Note: Exception: does not apply to 4220 Series	Enables MSTP, and allows non SPBM B-VLAN configuration on SPBM network-to-network interface (NNI) ports. The default is disabled. Note: Spanning Tree is disabled on all SPBM NNIs. You cannot add an SPBM NNI port or MLT port to any non SPBM B-VLAN.
EnableIpv6EgressFilterMode Note: Exception: does not apply to 4220 Series	Enables IPv6 egress filters. The default is disabled. If you change this parameter, you must restart the switch. For 5320 Series, 5420 Series, and 5720 Series platforms, EnableIpv6EgressFilterMode and EnableMacsec are mutually exclusive.
EnableMacsec Note: Exception: only required for MACsec supporting models of 5320 Series and 5420 Series.	Enables Media Access Control Security (MACsec) mode globally. To enable MACsec mode, you must configure the boot flag. EnableIpv6EgressFilterMode and EnableMacsec are mutually exclusive.
EnableSpbmNodeScaling Note: Exception: only applies to 5320 Series (not 5320-16P-2MXT-2X or 5320-24T-4X-XT) and 5420 Series.	Enables the switch to increase the number of supported SPB nodes per area. By default, the switch supports up to 350 SPB nodes per area. The default is disabled. If you change this parameter, you must restart the switch.
Slot	Specifies the slot number.
TftpHash	Enables TFTP hashing.
TftpRetransmit	Set TFTP retransmit timeout counter.
TftpTimeout	Set TFTP timeout counter.
User	Configure host user.
Password	Configure host password.

Reserve Bandwidth for Advanced Features

Use this procedure if you want the switch to support advanced features. When you enable the boot flag, you need to save and reboot with the new configuration.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Boot Config** tab.
3. In **AdvancedFeatureBWReservation**, select **high**, **low**, or **vim** to enable the boot flag.



Note

high applies to 4220 Series, 5720 Series, 7520 Series, and 7720 Series.

low does not apply to 5720 Series.

vim does not apply to 4220 Series, 5320 Series, 5420 Series, 7520 Series, or 7720 Series.

4. Select **Apply**.

A message displays to remind you that the configuration cannot include reserved ports, and that you must save the configuration and reboot the switch for changes to take effect.

5. Select **Yes** to continue or select **No** to cancel the change because the configuration includes reserved ports.

If you selected No, you can modify your switch configuration to remove the reserved ports and then return to this tab to change the **AdvancedFeatureBWReservation** configuration.

6. Save the configuration, and then reboot the switch.

Configure Jumbo Frames

About This Task

Configure jumbo frames to increase the Ethernet frame size to the maximum frame size your switch supports. The default is MTU size is 1950.

Procedure

1. On the Device Physical View, select the device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **Chassis** tab.
5. In **MTUSize**, select the desired frame size.
6. Select **Apply**.

Configure CP Limit

Configure CP Limit functionality to protect the switch from becoming congested by an excess of data flowing through one or more ports.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **CP Limit** tab.
5. Select **AutoRecoverPort**.
6. Select **Apply**.

CP Limit Field Descriptions

Use the data in the following table to use the **CP Limit** tab.

Name	Description
AutoRecoverPort	Activates or disables auto recovery of the port from action taken by CP Limit or link flap features. The default value is disabled.

Configure CP Limit on an Extreme Integrated Application Hosting Port



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to configure CP Limit functionality to protect the switch from becoming congested by excess data flow through Extreme Integrated Application Hosting (IAH) ports.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **CP Limit** tab.
4. Select **AutoRecoverPort**.
5. Select **Apply**.

CP Limit Field Descriptions

Use data in the following table to use the **CP Limit** tab.

Name	Description
AutoRecoverPort	Enables or disables auto recovery of the Extreme Integrated Application Hosting port from action taken by CP Limit or the link flap features. The default is disabled.

Edit the Management Port Parameters

Configure Layer 1 and Layer 2 networking parameters for the management port.



Note

This procedure only applies to hardware with a dedicated physical management interface.

The management port on the switch is a 10/100/1000 Mb/s Ethernet port that you can use for an out-of-band management connection to the switch.



Note

The 7830 Series has two management ports. RJ-45 (10M/1G/10G) and SFP+ (1G/10G) interfaces operate as primary and backup connections, respectively, with only one port active at a time.

About This Task

Static routes for the management OOB must use the Segmented Management Instance. For more information, see [Segmented Management Instance Configuration using EDM](#) on page 156.

If you use EDM to configure the static routes of the management port, you do not receive a warning if you configure a non-natural mask. After you save the changes, the system deletes those static routes after the next restart, possibly causing the loss of IP connectivity to the management port.

If you are uncertain whether the mask you configure is non-natural, use the CLI to configure static routes.

Procedure

1. In the Device Physical View tab, select the management port.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Mgmt Port**.
4. Select the **General tab**.
5. Modify the appropriate settings.
6. Select **Apply**.

General Field Descriptions

Use the data in the following table to use the **General** tab.

Name	Description
Index	Specifies the slot and port number of the management port.
AdminStatus	Configures the administrative status of the device as up (ready to pass packets) or down. The testing state indicates that no operational packets can be passed.
OperStatus	Specifies the operational status of the device.

Name	Description
LicenseControlStatus	Shows the license status of the port: <ul style="list-style-type: none"> Locked means the port requires a Port License but one is not present on the switch. Unlocked means the port requires a Port License and one is present on the switch. notApplicable means the port does not require a Port License.
Mtu	Shows the configuration for the maximum transmission unit, the size of the largest packet which can be sent or received on the interface, specified in octets. For interfaces that transmit network datagrams, this is the size of the largest network datagram that can be sent on the interface.
LinkTrap	Enables or disables traps for the link status.
PhysAddress	Shows the MAC address.
AutoNegotiate	Enables or disables Auto-Negotiation for the management port. The default is enabled.
AdminDuplex	Specifies the administrative duplex mode for the management port. The default is full.
OperDuplex	Specifies the operational duplex configuration for this port.
AdminSpeed	Specifies the administrative speed for this port. The default is 100 Mb/s.
OperSpeed	Shows the current operating data rate of the port.

View SFP Management Port Details

About This Task



Note

This procedure applies to 7830 Series only.

The SFP management port on the switch is a 1G/10G Ethernet port that you can use for an out-of-band management connection to the switch. The SFP management port operates as a backup interface and becomes active only when the primary interface is inactive.



Note

The 7830 Series has two management ports. RJ-45 (10M/1G/10G) and SFP+ (1G/10G) interfaces operate as primary and backup connections, respectively, with only one port active at a time. Only specific SFP modules are supported for the management port on the 7830 Series. For optics compatibility, see the [Extreme Optics](#) website.

Procedure

1. On the Device Physical View, select the SFP+ port.

2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Mgmt SFP Port**.
4. Select **General**.

Mgmt SFP Port Field Descriptions

Use the data in the following table to use the **Mgmt SFP Port** tab.

Name	Description
ConnectorType	Specifies the type of connector.
SupportsDDM	Specifies if the transceiver supports DDM.
DdmStatusMask	Specifies the DDM status. A value other than ddm-ok represents a specific error.
CLEI	Specifies the Telcordia register assignment CLEI code.
VendorName	Specifies the name of the manufacturer.
VendorPartNumber	Specifies the part number for the transceiver. Select the part number to access the technical specifications for the part.
VendorRevNumber	Specifies the manufacturer revision level for the transceiver.
VendorSN	Specifies the manufacturer serial number for the transceiver.
VendorDateCode	Specifies the manufacturer date code for the transceiver.
Wavelength	Specifies the wavelength in nm. This field is valid for optical transceivers only.
Calibration	Specifies if the calibration is internal or external.
PowerMeasure	Specifies Rx power measurement as average or OMA.
Aux1Monitoring	Specifies if auxiliary monitoring is implemented for the transceiver.
Aux2Monitoring	Specifies if auxiliary monitoring is implemented for the transceiver.
TemperatureLowAlarmThreshold	Specifies the low alarm threshold in degrees Celsius.
TemperatureLowWarningThreshold	Specifies the low warning threshold in degrees Celsius.
Temperature	Specifies the current temperature in degrees Celsius of the transceiver.
TemperatureHighWarningThreshold	Specifies the high warning threshold in degrees Celsius.
TemperatureHighAlarmThreshold	Specifies the high alarm threshold in degrees Celsius.

Name	Description
TemperatureStatus	Specifies if any temperature thresholds were exceeded.
VoltageLowAlarmThreshold	Specifies the low alarm threshold in volts.
VoltageLowWarningThreshold	Specifies the low warning threshold in volts.
Voltage	Specifies the current voltage in volts.
VoltageHighWarningThreshold	Specifies the high warning threshold in volts.
VoltageHighAlarmThreshold	Specifies the high alarm threshold in volts.
VoltageStatus	Specifies if any voltage thresholds were exceeded.
BiasLowAlarmThreshold	Specifies the bias current low alarm threshold in mA.
BiasLowWarningThreshold	Specifies the bias current low warning threshold in mA.
Bias	Specifies the laser bias current in mA.
BiasHighWarningThreshold	Specifies the bias current high warning threshold in mA.
BiasHighAlarmThreshold	Specifies the bias current high alarm threshold in mA.
BiasStatus	Specifies if any bias thresholds were exceeded.
TxPowerLowAlarmThreshold	Specifies the low alarm threshold in dBm for the Tx power.
TxPowerLowWarningThreshold	Specifies the low warning threshold in dBm for the Tx power.
TxPower	Specifies the current Tx power in dBm.
TxPowerHighWarningThreshold	Specifies the high warning threshold in dBm for the Tx power.
TxPowerHighAlarmThreshold	Specifies the high alarm threshold in dBm for the Tx power.
TxPowerStatus	Specifies if any Tx power thresholds were exceeded.
RxPowerLowAlarmThreshold	Specifies the low alarm threshold in dBm for the Rx power.
RxPowerLowWarningThreshold	Specifies the low warning threshold in dBm for the Rx power.
RxPower	Specifies the current Rx power in dBm.
RxPowerHighWarningThreshold	Specifies the high warning threshold in dBm for the Rx power.
RxPowerHighAlarmThreshold	Specifies the high alarm threshold in dBm for the Rx power.
RxPowerStatus	Specifies if any Rx power thresholds were exceeded.

Name	Description
Aux1LowAlarmThreshold	Specifies the low alarm threshold auxiliary 1 reading.
Aux1LowWarningThreshold	Specifies the low warning threshold auxiliary 1 reading.
Aux1	Specifies the current auxiliary 1 reading.
Aux1HighWarningThreshold	Specifies the high warning threshold auxiliary 1 reading.
Aux1HighAlarmThreshold	Specifies the high alarm threshold auxiliary 1 reading.
Aux1Status	Specifies if any auxiliary 1 thresholds were exceeded.
Aux2LowAlarmThreshold	Specifies the low alarm threshold auxiliary 2 reading.
Aux2LowWarningThreshold	Specifies the low warning threshold auxiliary 2 reading.
Aux2	Specifies the current auxiliary 2 reading.
Aux2HighWarningThreshold	Specifies the high warning threshold auxiliary 2 reading.
Aux2HighAlarmThreshold	Specifies the high alarm threshold auxiliary 2 reading.
Aux2Status	Specifies if any auxiliary 2 thresholds were exceeded.



Note

1. Threshold and actual values for TxBias, TxPower, and RxPower are provided for all four channels in QSFP+ and QSFP28 optical transceivers.
2. Auxiliary monitoring does not apply to QSFP+s or QSFP28s.

Automatically Reactivate a Port Shutdown by SLPP

About This Task

Use the following procedure to automatically reactivate the port that is shut down by the SLPP.

Procedure

1. In the Device Physical View tab, select the port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **CP Limit** tab.

5. Select **AutoRecoverPort** to activate auto recovery of the port from the action taken by SLPP shutdown features.
The default value is disabled.
6. Select **Apply**.

Edit Serial Port Parameters

About This Task

Perform this procedure to specify serial port communication settings. The serial port on the device is the console port.

Procedure

1. In the Device Physical View tab, select the console port on the device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Serial Port**.
4. Edit the port parameters as required.
5. Select **Apply**.

Serial Port Field Descriptions

Use the data in the following table to use the **Serial Port** tab.

Name	Description
IfIndex	Identifies the port as a serial port.
BaudRate	Specifies the baud rate of this port. The default value is 115200.
DataBits	Specifies the number of data bits, for each byte of data, this port sends and receives. The default is eight.

Enable Port Lock

About This Task

Use the port lock feature to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify locked ports until you first unlock the port.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Port Lock** tab.
4. To enable port lock, select **Enable**.
5. Select **Apply**.

Port Lock Field Descriptions

Use the data in the following table to use the **Port Lock** tab.

Name	Description
Enable	Activates the port lock feature. Clear this check box to unlock ports. The default is disabled.
LockedPorts	Lists the locked ports. Select the ellipsis (...) to select the ports you want to lock or unlock.

Lock a Port

Before You Begin

- You must enable port lock before you can lock or unlock a port.

About This Task

Use the port lock feature to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify locked ports until you first unlock the port.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **Port Lock** tab.
4. For **LockedPorts**, select the ellipsis (...).
5. Select the desired port or ports.
6. Select **Ok**.
7. In the **Port Lock** tab, select **Apply**.

Port Lock Field Descriptions

Use the data in the following table to use the **Port Lock** tab.

Name	Description
Enable	Activates the port lock feature. Clear this check box to unlock ports. The default is disabled.
LockedPorts	Lists the locked ports. Select the ellipsis (...) to select the ports you want to lock or unlock.

View Power Information

About This Task

View power information to see the amount of power available and used by the chassis and all components.

Procedure

1. On the Device Physical View, select the Device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **Power Info** tab.

Power Info Field Descriptions

Use the data in the following table to use the **Power Info** tab.

Name	Description
TotalPower	Shows the total power for the chassis.
RedundantPower	Shows the redundant power for the chassis.

View Power Status

About This Task

Perform the following procedure to view the power consumption of the modules in the chassis.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **Power Consumption** tab.

Power Consumption Field Descriptions

Use the data in the following table to use the **Power Consumption** tab.

Name	Description
Index	Specifies an index value that identifies the component.
PowerStatus	Specifies the power status: on or off.
ConsumedPower	Specifies the actual consumed power for the slot. This value is 0 if the power status is off.
SlotDescription	Specifies the slot number.
CardDescription	Specifies the type of module in the slot.

NEW! Configure the Minimum Fan Speed



Note

This procedure only applies to 5720 Series.

About This Task

Perform this procedure to adjust the thermal operations of the switch. For example, if you want installed transceivers to operate at a lower temperature, increase the fan speed.

The switch monitors temperature sensors every 15 seconds. If the temperature exceeds the threshold, the fan speed increases.

If you change the minimum fan speed value, the switch can take up to 15 seconds to read the temperature sensors and then change the fan speed. If you lower the minimum fan speed by more than 20%, the switch decreases the speed by 20% every 15 seconds.



Important

Increasing the fan speed can also increase the fan noise.

Procedure

1. In the navigation pane, expand **Configuration > Edit**
2. Select **Chassis**.
3. Select the **Fan Control** tab.
4. In **FanMinSpeed**, type the fan speed as a percentage.
5. Select **Apply**.

NEW! Fan Control Field Descriptions

Use the data in the following table to use the **Fan Control** tab.

Name	Description
FanMinSpeed	Configures the minimum fan speed as a percentage between the minimum (20%) and maximum (100%) speed imposed by the platform design. The default is 20%.

View Fan Tray Information

View fan tray information to see manufacturing information about the fans.



Note

This procedure does not apply to 5320-16P-2MXT-2X.

Procedure

1. On the Device Physical View, select the Device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **Fan Tray Info** tab.

Fan Tray Info Field Descriptions

Use the data in the following table to use the **Fan Tray Info** tab.

Name	Description
Description	Shows a description of the fan tray.
FlowType	Specifies whether the air flow is front-to-back or back-to-front.

View Fan Information

View fan information to monitor the alarm status of the cooling ports in the chassis.



Note

This procedure does not apply to 5320-16P-2MXT-2X.

About This Task

For platforms that support both back-to-front and front-to-back airflow, the airflow direction must be the same for both the power supply fans and the chassis fan.

Procedure

1. On the Device Physical View, select the Device.
2. In the navigation pane, expand **Configuration > Edit**.
3. Select **Chassis**.
4. Select the **Fan Info** tab.

Fan Info Field Descriptions

Use the data in the following tables to use the **Fan Info** tab.

Name	Description
Description	Specifies a description of the fan location.
OperStatus	Specifies the operation status of the fan.
OperSpeed	Specifies the actual fan speed.
OperSpeedRPM	Specifies the current operational speed of the fan in RPM.

View USB Port Information

About This Task

Perform this procedure to view information about the USB port on the switch.

Procedure

1. In the Device Physical View, select the USB port.
2. In the navigation pane, expand **Configuration > Edit**.

3. Select **USB Port**.
4. Select the **General** tab.

General Field Descriptions

Use the data in the following table to use the **General** tab.

Name	Description
UsbStatus	Displays the current status of USB storage: either present or notPresent.
UsbDescription	Displays a description of the USB storage.

View Topology Status Information

About This Task

View topology status information, which includes MIB status information, to view the configuration status of the SynOptics Network Management Protocol (SONMP) on the system.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Topology**.
3. Select the **Topology** tab.

Topology Field Descriptions

Use the data in the following table to use the **Topology** tab.

Name	Description
IpAddr	Specifies the IP address of the device.
Status	Indicates whether topology (SONMP) is on or off for the device.
NmmLstChg	Specifies the value of sysUpTime, the last time an entry in the network management MIB (NMM) topology table was added, deleted, or modified, if the table did not change since the last cold or warm start of the agent.
NmmMaxNum	Specifies the maximum number of entries in the NMM topology table.
NmmCurNum	Specifies the current number of entries in the NMM topology table.

View the Topology Message Status

About This Task

View topology message status to view the interconnections between Layer 2 devices in a network.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Topology**.
3. Select the **Topology Table** tab.

Topology Table Field Descriptions

Use the data in the following table to use the **Topology Table** tab.

Name	Description
Slot	Specifies the slot number in the chassis that received the topology message.
Port	Specifies the port that received the topology message.
SubPort	Specifies the channel of a channelized 40 Gbps port that received the topology message.
IpAddr	Specifies the IP address of the sender of the topology message.
SegId (RemPort)	Specifies the segment identifier of the segment from which the remote agent sent the topology message. This value is extracted from the message.
MacAddr	Specifies the MAC address of the sender of the topology message.
ChassisType	Specifies the chassis type of the device that sent the topology message.
BkplType	Specifies the backplane type of the device that sent the topology message.
LocalSeg	Indicates if the sender of the topology message is on the same Ethernet segment as the reporting agent.
CurState	Specifies the current state of the sender of the topology message. The choices are <ul style="list-style-type: none"> • topChanged—Topology information recently changed. • heartbeat—Topology information is unchanged. • new—The sending agent is in a new state.

Configure Port Speed for All VIM Ports



Note

This procedure only applies to 5520 Series and 5720 Series.

Configure all of the ports on an installed Versatile Interface Module (VIM) to operate at the same speed.

**Note**

Some VIMs must operate with all ports at the same speed, or a group of ports at the same speed, while others can operate with ports at different speeds. For more information, see *Fabric Engine Release Notes*. You can configure VIM ports speed only on VIMs that must operate with all ports at the same speed.

Before You Begin

Install the VIM before performing this procedure.

About This Task**Note**

It is a best practice not to use a mix of 10Gbps and 25Gbps VIMs with the auto parameter.

Use this procedure to configure the speed of all ports in a multi-port VIM to operate at either 1 Gbps, 10 Gbps, 25 Gbps, or auto.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **VIM** tab.
4. Select **mbps1000**, **mbps10000**, or **mbps25000**.
5. Select **Apply**.

VIM Field Descriptions

Use the data in the following table to use the **VIM** tab.

Name	Description
AdminSpeed Note: Exception: 1 Gbps port speed applies to 5720-VIM-6YE only.	<ul style="list-style-type: none"> • mbps1000: Configures all ports or a group of ports in a multi-port VIM to operate at 1 Gbps. • mbps10000: Configures all ports in a multi-port VIM to operate at 10 Gbps. • mbps25000: Configures all ports in a multi-port VIM to operate at 25 Gbps.
AutoVimSpeed	Configures all ports in a multi-port VIM to operate at the detected link speed. The default is enabled.
GroupSpeed Exception: only applies to 5720-VIM-6YE.	Configures all ports into two speed setting groups. <ul style="list-style-type: none"> • group1: ports 2/1, 2/2, 2/3 • group2: ports 2/4, 2/5, 2/6 Ports in groups must be configured to operate at the same speed. For example, ports 2/1, 2/2, 2/3 can be configured for 10 Gbps, and ports 2/4, 2/5, 2/6 can be configured for 25 Gbps. If 5720-VIM-6YE uses the LRM/MACsec Adapter, you must configure the groups according to the SFP or SFP+ inserted in the Adapter.

NEW Swap a VIM

You can hot swap a Versatile Interface Module (VIM) on the 7830 Series using Controlled mode or Uncontrolled mode.

About This Task

Consider the following when you swap VIMs:

- Mandatory wait times between VIM swaps are as follows:
 - 15 seconds after you remove a VIM
 - 60 seconds after you insert a VIM
- When you use uncontrolled mode:
 - Remove the transceiver before you remove the VIM, and insert the replacement VIM before you install the transceiver.
 - Wait 60 seconds after you insert the VIM before you install the transceiver.
- When you swap between the same VIM type, configuration on the ports is preserved then reapplied. When you swap between different VIM types, configuration on the ports revert to default configuration.
- MACSec configuration on channelized ports is lost when you swap between the same VIM type.

Perform this procedure to swap a VIM.

Procedure

1. In the Device Physical View, right-click the VIM.
2. Select **Edit**.
3. In the Card tab, select **off** to power off the VIM.
4. Remove the VIM.
5. Insert the replacement VIM.
6. In the Card tab, select **on** to power on the VIM.

Card Field Descriptions

Use the data in the following table to use the **Card** tab.

Name	Description
CardType	Displays the model number of the module.
SerialNum	Displays the serial number for the installed module.
PartNumber	Displays the part number.
CardAssemblyDate	Displays the date the module was assembled.
HWVersion	Displays the hardware revision.
AdminStatus Note: Exception: only applies to 7830 Series.	Enables or disables the VIM.
OperStatus Note: Exception: only applies to 7830 Series.	Displays the operational status of the VIM.
SlotPower Note: Exception: only applies to 7830 Series.	Enables or disables power on the VIM slot.

Node Alias Configuration using EDM

This section provides procedures to configure Node Alias using Enterprise Device Manager (EDM).

Configure Node Alias Globally

About This Task

By default, the switch enables Node Alias globally, as well as on Auto-sense enabled ports.



Note

You cannot enable Node Alias on SPBM NNI ports.

If you disable Node Alias globally, port level functionality is also disabled.



Note

The steps in this procedure assume Node Alias is disabled globally as well as on ports.

Perform the following procedure to configure Node Alias.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Node Alias**.
3. Select the **GlobalConfig/Info** tab.
4. Select **GlobalEnable**.
5. In **ClientsMaxNumberEntries**, type the number of clients to be stored in the Node Alias client database on the switch.
6. (Optional) Select **ClientsClearAll** to clear all client information from the database.
7. Select **Apply**.

Global Config/Info Field Descriptions

Use the data in the following table to use the **Global Config/Info** tab.

Name	Description
ClientsMaxNumberEntries	Specifies the total number of clients entries. The default is 8,192.
ClientsClearAll	Clears all client information from the database.
StatsTotalClientEntries	Displays the total number of stored clients.
StatsTotalActiveClientEntries	Displays the total number of active stored clients.
StatsTotalProtocolInfoEntries	Displays the total number of stored protocol entries for clients.
GlobalEnable	Enables or disables Node Alias globally.

View Node Alias Client Information

About This Task

Perform the following procedure to view Node Alias client information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Node Alias**.
3. Select the **Client Info** tab.

4. (Optional) To clear a single client from the database, double-click an entry in the **EntryClear** column.
 - a. Select **true**.
 - b. Select **Apply**.

Client Info Field Descriptions

Use the data in the following table to use the **Client Info** tab.

Name	Description
MacAddress	Displays the MAC address of the client.
Isid Index	Displays details for the specified I-SID.
Vlan Index	Displays details for the specified VLAN ID.
Isid Value	Displays the I-SID value of the I-SID-to-VLAN assignment.
InterfaceIndex	Displays the receiving interface port or MLT.
Capabilities	Displays the protocols the client uses.
IsActive	Displays if the client is active.
Ipv4Addr	Displays the last IPv4 address used by this client. Note: Option 53 reads the IP address.
Ipv6Addr	Displays the last IPv6 address used by this client.
FirstLearnedTimer	Displays the timestamp when the client was first detected on the port.
LastLearnedTimer	Displays the timestamp when the client was last detected on the port.
EntryClear	Clears the client entry from the database.
Hostname	Displays the client host name.
OSType	Displays the client operating system type.
DeviceType	Displays the client device type.
OSVersion	Displays the client operating system version.

View Node Alias Protocol Information

About This Task

Perform the following procedure to view Node Alias protocol information.



Note

Detailed information displays for IPv4, IPv6, BGP, VRRP, and broadcast UDP protocols only.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Node Alias**.
3. Select the **Protocol Info** tab.

Protocol Info Field Descriptions

Use the data in the following table to use the **Protocol Info** tab.

Name	Description
SourceMacAddress	Displays the source MAC address.
Isid	Displays the I-SID.
Vlan	Displays the VLAN.
InterfaceIndex	Displays the receiving port or MLT interface.
Type	Displays the protocol type.
Value	Displays additional details regarding the used protocol, such as: <ul style="list-style-type: none">• IP address for the IPv4 protocol type• IP address for the IPv6 protocol type• AS Number for the BGP protocol type• VR ID for the VRRP protocol type• Source port for the broadcast UDP protocol type
Isid Value	Displays the I-SID value of the I-SID-to-VLAN assignment.
LastLearnedTimer	Displays the timestamp when the client was last updated on the port.

View Node Alias DHCP Fingerprint Information

About This Task

Perform the following procedure to view Node Alias DHCP fingerprint information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Node Alias**.
3. Select the **DHCP Fingerprint** tab.

DHCP Fingerprint Field Descriptions

Use the data in the following table to use the **DHCP Fingerprint** tab.

Name	Description
SourceMacAddress	Displays the source MAC address.
Isid	Displays the I-SID.
Vlan	Displays the VLAN ID.
Isid Value	Displays the I-SID value of the I-SID-to-VLAN assignment.
InterfaceIndex	Displays the receiving port or MLT interface.
Ipv4Addr	Displays the last known IPv4 address of the client. Note: Option 53 reads the IP address.
Option12	Displays the hostname of the client received from the DHCP server.
Option55	Displays the list of configuration parameters the client requests.
Option60	Displays the vendor class domain name.

View Modular SSD Information



Note

This procedure only applies to 5720 Series.

About This Task

Perform this procedure to display information about an installed Solid State Drive (SSD) on a switch.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **SSD** tab.

SSD Field Descriptions

Use the data in the following table to use the **SSD** tab.

Name	Description
ProductName	Specifies Solid State Drive (SSD) product name.
VendorName	Specifies the SSD vendor.
ManufactureDate	Specifies the date on which the SSD was manufactured.
SerialNum	Specifies the SSD serial number.

Name	Description
PartNum	Specifies the SSD part number.
DeviceVersion	Specifies the version of the SSD.
TotalSize	Specifies the total memory size of the SSD.

Graph Chassis Statistics

About This Task

Create graphs of chassis statistics to generate a visual representation of your data.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. On the Graph Chassis tab, select the tab with the data you want to graph:
 - **System**
 - **SNMP**
 - **IP**
 - **ICMP In**
 - **ICMP Out**
 - **TCP**
 - **UDP**
 - **Fabric Attach**
4. Select the statistic you want to graph.
5. Select the graph type:
 - line chart
 - area chart
 - bar chart
 - pie chart

Graph Port Statistics

About This Task

You can create a graph of the port statistics to generate a visual representation of your data.

You can also use the shortcut menu by right-clicking the selected port or ports, and then select **Graph**.

Procedure

1. In the Device Physical View, select the port or ports for which you want to create a graph.

2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. On the **Graph Port** tab, select the item you want to graph.
5. Select an icon to select the type of graph you require:
 - Line Chart
 - Area Chart
 - Bar Chart
 - Pie Chart

View Chassis System Statistics

About This Task

Use the following procedure to create graphs for chassis statistics.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. Select the **System** tab.

System Field Descriptions

The following table describes the fields on the **System** tab.

Name	Description
MemUsed	The percentage of memory space used. Only the AbsoluteValue column is valid in the System tab. All other columns display as N/A because they are percentages and not actual memory counters.
MemFree	The amount in kilobytes of free memory.
CpuUtil	Percentage of CPU utilization.

View Chassis SNMP Statistics

About This Task

View chassis SNMP statistics to monitor network performance.

Procedure

1. In the Device Physical View, select the chassis.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Chassis**.
4. Select the **SNMP** tab.

SNMP Field Descriptions

The following table describes parameters on the **SNMP** tab.

Name	Description
InPkts	The number of messages delivered to the SNMP entity from the transport service.
OutPkts	The number of SNMP messages passed from the SNMP protocol entity to the transport service.
InTotalReqVars	The number of MIB objects retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
InTotalSetVars	The number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.
InGetRequests	The number of SNMP Get-Request PDUs the SNMP protocol accepts and processes.
OutGetRequests	The number of SNMP Get-Request PDUs that are generated by the SNMP protocol entity.
InGetNexts	The number of SNMP Get-Next PDUs the SNMP protocol accepts and processes.
OutGetNexts	The number of SNMP Get-Next PDUs that are generated by the SNMP protocol entity.
InSetRequests	The number of SNMP Set-Request PDUs the SNMP protocol accepts and processes.
OutSetRequests	The number of SNMP Set-Request PDUs that are generated by the SNMP protocol entity.
InGetResponses	The number of SNMP Get-Response PDUs the SNMP protocol accepts and processes.
OutGetResponses	The number of SNMP Get-Response PDUs that are generated by the SNMP protocol entity.
InTraps	The number of SNMP Trap PDUs the SNMP protocol accepts.
OutTraps	The number of SNMP Trap PDUs the SNMP protocol generates.
OutTooBigs	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is tooBig.
OutNoSuchNames	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is noSuchName.
OutBadValues	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is badValue.
OutGenErrs	The number of SNMP PDUs the SNMP protocol generates for which the value of the error-status field is genErr.
InBadVersions	The number of SNMP messages delivered to the SNMP protocol entity for an unsupported SNMP version.

Name	Description
InBadCommunityNames	The number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to said entity.
InBadCommunityUses	The number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed by the SNMP community named in the message.
InASNParseErrs	The number of ASN.1 or BER errors the SNMP protocol encountered when decoding received SNMP messages.
InTooBig	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is tooBig.
InNoSuchNames	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is noSuchName.
InBadValues	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is badValue.
InReadOnlys	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is readOnly. It is a protocol error to generate an SNMP PDU containing the value "readOnly" in the error-status field. This object is provided to detect incorrect implementations of the SNMP.
InGenErrs	The number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is genErr.

View Chassis IP Statistics

About This Task

View chassis IP statistics to monitor network performance.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. Select the **IP** tab.

IP Field Descriptions

The following table describes parameters on the **IP** tab.

Name	Description
InReceives	The number of input datagrams received from interfaces, including those received in error.
InHdrErrors	The number of input datagrams discarded due to errors in the IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options.
InAddrErrors	The number of input datagrams discarded because the IP address in the IP header destination field was not a valid address to be received at this entity. This count includes invalid addresses and addresses of unsupported Classes (for example, Class E). For entities that are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
ForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities that do not act as IP Gateways, this counter includes only those packets that were Source-Routed by way of this entity and had successful Source-Route option processing.
InUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
InDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing but that were discarded (for example, for lack of buffer space). This counter does not include any datagrams discarded while awaiting reassembly.
InDelivers	The number of input datagrams successfully delivered to IP user-protocols (including ICMP).
OutRequests	The number of IP datagrams that local IP user-protocols (including ICMP) supplied to IP in requests for transmission. This counter does not include any datagrams counted in ipForwDatagrams.
OutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but that were discarded (for example, for lack of buffer space). This counter includes datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
OutNoRoutes	The number of IP datagrams discarded because no route was found to transmit them to their destination. This counter includes any packets counted in ipForwDatagrams that meet this no-route criterion. This counter includes any datagrams a host cannot route because all default gateways are down.

Name	Description
FragOKs	The number of IP datagrams that were successfully fragmented at this entity.
FragFails	The number of IP datagrams that were discarded because they needed to be fragmented at this entity but can not be, for example, because the Don't Fragment flags were set.
FragCreates	The number of IP datagram fragments that were generated as a result of fragmentation at this entity.
ReasmReqds	The number of IP fragments received that needed to be reassembled at this entity.
ReasmOKs	The number of IP datagrams successfully reassembled.
ReasmFails	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so on). This number is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.

View Chassis ICMP In Statistics

About This Task

View chassis ICMP In statistics to monitor network performance.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. Select the **ICMP In** tab.

ICMP In Field Descriptions

The following table describes parameters on the **ICMP In** tab.

Name	Description
SrcQuenchs	The number of ICMP Source Quench messages received.
Redirects	The number of ICMP Redirect messages received.
Echos	The number of ICMP Echo (request) messages received.
EchoReps	The number of ICMP Echo Reply messages received.
Timestamps	The number of ICMP Timestamp (request) messages received.
TimestampReps	The number of ICMP Timestamp Reply messages received.
AddrMasks	The number of ICMP Address Mask Request messages received.
AddrMaskReps	The number of ICMP Address Mask Reply messages received.

Name	Description
ParmProbs	The number of ICMP Parameter Problem messages received.
DestUnreachs	The number of ICMP Destination Unreachable messages received.
TimeExcds	The number of ICMP Time Exceeded messages received.

View Chassis ICMP Out Statistics

About This Task

View chassis ICMP Out statistics to monitor network performance.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. Select the **ICMP Out** tab.

ICMP Out Field Descriptions

The following table describes parameters on the **ICMP Out** tab.

Name	Description
SrcQuenchs	The number of ICMP Source Quench messages sent.
Redirects	The number of ICMP Redirect messages received. For a host, this object is always zero, because hosts do not send redirects.
Echos	The number of ICMP Echo (request) messages sent.
EchoReps	The number of ICMP Echo Reply messages sent.
Timestamps	The number of ICMP Timestamp (request) messages sent.
TimestampReps	The number of ICMP Timestamp Reply messages sent.
AddrMasks	The number of ICMP Address Mask Request messages sent.
AddrMaskReps	The number of ICMP Address Mask Reply messages sent.
ParmProbs	The number of ICMP Parameter Problem messages sent.
DestUnreachs	The number of ICMP Destination Unreachable messages sent.
TimeExcds	The number of ICMP Time Exceeded messages sent.

View Chassis TCP Statistics

About This Task

View TCP statistics to monitor network performance.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. Select the **TCP** tab.

TCP Field Descriptions

The following table describes parameters on the **TCP** tab.

Name	Description
ActiveOpens	The number of times TCP connections made a direct transition to the SYN-SENT state from the CLOSED state.
PassiveOpens	The number of times TCP connections made a direct transition to the SYN-RCVD state from the LISTEN state.
AttemptFails	The number of times TCP connections made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections made a direct transition to the LISTEN state from the SYN-RCVD state.
EstabResets	The number of times TCP connections made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
CurrEstab	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
InSegs	The number of segments received, including those received in error. This count includes segments received on currently established connections.
OutSegs	The number of segments sent, including those on current connections, but excluding those containing only retransmitted octets.
RetransSegs	The number of segments retransmitted that is, the number of TCP segments transmitted containing one or more previously transmitted octets.
InErrs	The number of segments received in error (for example, bad TCP checksums).
OutRsts	The number of TCP segments sent containing the RST flag.
HCInSegs	The number of segments received, including those received in error. This count includes segments received on currently established connections. This object is the 64-bit equivalent of InSegs.
HCOutSegs	The number of segments sent, including those on current connections, but excluding those containing only retransmitted octets. This object is the 64-bit equivalent of OutSegs.

View Chassis UDP Statistics

About This Task

Display User Datagram Protocol (UDP) statistics to see information about the UDP datagrams.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.
2. Select **Chassis**.
3. Select the **UDP** tab.
4. Select the information you want to graph.
5. Select the type of graph you want:

- line
- area
- bar
- pie

6. To clear counters, select **Clear Counters**.

Discontinuities in the value of these counters can occur when the management system reinitializes, and at other times as indicated by discontinuities in the value of sysUpTime.

UDP Field Descriptions

Use the data in the following table to use the **UDP** tab.

Name	Description
NoPorts	The number of received UDP datagrams with no application at the destination port. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.
InErrors	The number of received UDP datagrams that were not delivered for reasons other than the lack of an application at the destination port. Discontinuities in the value of this counter can occur at reinitialization of the management system and at other times as indicated by discontinuities in the value of sysUpTime.
InDatagrams	The number of UDP datagrams delivered to UDP users, for devices that can receive more than 1 000 000 UDP datagrams for each second. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.

Name	Description
OutDatagrams	The number of UDP datagrams sent from this entity. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.
HCInDatagrams	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT.
HCOutDatagrams	The number of UDP datagrams sent from this entity, for devices that can transmit more than 1 million UDP datagrams for each second. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by discontinuities in the value of sysUpTime.

View Port Interface Statistics

About This Task

View port interface statistics to manage network performance.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **Interface** tab.

Interface Field Descriptions

The following table describes parameters on the **Interface** tab.

Name	Description
InOctets	Specifies the number of octets received on the interface, including framing characters.
OutOctets	Specifies the number of octets transmitted from the interface, including framing characters.
InUcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that were not addressed to a multicast or broadcast address at this sublayer.
OutUcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this sublayer. The total number includes those packets discarded or not sent.

Name	Description
InMulticastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both group and functional addresses.
OutMulticastPkts	Specifies the number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this number includes both group and functional addresses.
InBroadcastPkts	Specifies the number of packets delivered by this sublayer to a higher sublayer that are addressed to a broadcast address at this sublayer.
OutBroadcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent.
InDiscards	Specifies the number of inbound packets that are discarded because of frames with errors or invalid frames or, in some cases, to fill up buffer space.
InErrors	For packet-oriented interfaces, specifies the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol.
InUnknownProtos	For packet-oriented interfaces, specifies the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received through the interface that were discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter is always 0.
HCInPfcPkts	Specifies the total number of Priority Flow Control (PFC) packets received by this interface. This number does not increment for port-level flow control.
HCOutPfcPkts	Specifies the total number of PFC packets transmitted by this interface. This number does not increment for port-level flow control.
InFlowCtrlPkts	Specifies the number of port-level flow control packets received by this interface.
OutFlowCtrlPkts	Specifies the number of port-level flow control packets transmitted by this interface.
InPfcPkts	Specifies the total number of port-level flow control packets received by this interface.

Name	Description
OutPfcPkts	Specifies the total number of port-level flow control packets transmitted by this interface.
NumStateTransition	Specifies the number of times the port went in and out of service; the number of state transitions from up to down.
RateLimitDropCounter Note: Exception: does not display on 7830 Series	Displays the number of packets dropped on the port because they exceeded the configured rate limit.

View Port Ethernet Errors Statistics

About This Task

View port Ethernet errors statistics to manage network performance.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **Ethernet Errors** tab.

Ethernet Errors Field Descriptions

The following table describes parameters on the **Ethernet Errors** tab.

Name	Description
AlignmentErrors	Specifies account of frames received on a particular interface that are not an integral number of octets in length and do not pass the FCS check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtain are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
FCSErrors	Specifies a count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. The count represented by an instance of this object increments when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.

Name	Description
InternalMacTransmitErrors	Specifies a count of frames for which transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.
InternalMacReceiveErrors	Specifies a count of frames for which reception on a particular interface fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSerrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of receive errors on a particular interface that are not otherwise counted.
CarrierSenseErrors	Specifies the number of times that the carrier sense condition is lost or not asserted when the switch attempts to transmit a frame on a particular interface. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
FrameTooLongs	Specifies a count of frames received on a particular interface that exceed the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
SQETestErrors	Specifies a count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation described in section 7.2.4.6 of the same document.
DeferredTransmissions	Specifies a count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.

Name	Description
SingleCollisionFrames	Specifies a count of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the UcastPkts, MulticastPkts, or BroadcastPkts objects and is not counted by the corresponding instance of the MultipleCollisionFrames object.
MultipleCollisionFrames	Specifies a count of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the UcastPkts, MulticastPkts, or BroadcastPkts objects and is not counted by the corresponding instance of the SingleCollisionFrames object.
LateCollisions	Specifies the number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet; 512 corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
ExcessiveCollisions	Specifies a count of frames for which transmission on a particular interface fails due to excessive collisions.
FrameTooShorts	Specifies the number of frames, encountered on this interface, that are too short.
LinkFailures	Specifies the number of link failures encountered on this interface.
PacketErrors	Specifies the number of packet errors encountered on this interface.
CarrierErrors	Specifies the number of carrier errors encountered on this interface.
LinkInactiveErrors	Specifies the number of link inactive errors encountered on this interface.



DHCP and UDP Configuration

[DHCP Fundamentals](#) on page 819
[DHCP Server](#) on page 820
[DHCP Relay](#) on page 825
[UDP Broadcast Forwarding](#) on page 831
[DHCP and UDP Configuration Using the CLI](#) on page 832
[IPv6 DHCP Relay Configuration using CLI](#) on page 864
[DHCP and UDP Configuration Using EDM](#) on page 870
[IPv6 DHCP Relay Configuration using EDM](#) on page 902

Table 90: Dynamic Host Configuration Protocol product support

Feature	Product	Release introduced
Dynamic Host Configuration Protocol (DHCP) Server	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0
	5420 Series	Fabric Engine 9.0
	5520 Series	Fabric Engine 9.0
	5720 Series	Fabric Engine 9.0
	7520 Series	Fabric Engine 9.0
	7720 Series	Fabric Engine 9.0
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.0
	VSP 7400 Series	VOSS 9.0

Table 90: Dynamic Host Configuration Protocol product support (continued)

Feature	Product	Release introduced
Option 43 (Vendor Specific) and Option 60 (Class Id)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0.2
	5420 Series	Fabric Engine 9.0.2
	5520 Series	Fabric Engine 9.0.2
	5720 Series	Fabric Engine 9.0.2
	7520 Series	Fabric Engine 9.0.2
	7720 Series	Fabric Engine 9.0.2
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.0.2
	VSP 7400 Series	VOSS 9.0.2
Dynamic Host Configuration Protocol (DHCP) Relay	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Option 82 (Relay Agent Information)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 90: Dynamic Host Configuration Protocol product support (continued)

Feature	Product	Release introduced
DHCP Smart Relay	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.10
	5420 Series	Fabric Engine 8.10
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.10
	VSP 7400 Series	VOSS 8.10

Table 91: Dynamic Host Configuration Protocol Relay for IPv6 product support

Feature	Product	Release introduced
IPv6 DHCP Relay	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

DHCP Fundamentals

The Dynamic Host Configuration Protocol (DHCP) simplifies network management by eliminating the need to have an administrator assign IP addresses manually to each client. A DHCP server assigns IP addresses to clients in the local network automatically from a preconfigured pool of addresses.

DHCP support includes the following main subfeatures:

- DHCP Server
- DHCP Relay

For information on DHCP-related logs and alarms, see *Fabric Engine Alarms and Logs Reference*.

DHCP Server

With the DHCP Server feature, the switch hosts the server thereby removing the need to deploy an external DHCP server. The hosted DHCP Server has the same basic functionality as a standalone external DHCP server. The switch listens for DHCP requests from clients on the network and assigns IP addresses dynamically from a preconfigured address pool.

**Note**

DHCP Server information in this section applies to the DHCP Server feature hosted on the switch. If you deploy an external DHCP server, see the documentation that comes with your device.

DHCP Server includes the following support options:

- DHCP Server supports IPv4 deployments. For IPv6-only deployments, you must deploy an external DHCP server on a standalone server or other network device.
- You can deploy DHCP Server with clients that are in the GRT as well as user-defined VRFs.
- You can configure DHCP Server and DHCP Relay to provide DHCP services to clients in local VLANs as well as clients that connect from a remote switch.
- You can configure DHCP Server to listen for DHCP messages or BOOTP messages, or both.
- If the DHCP Server operates in authoritative mode, the server is the sole authority on the network and assigns all IP addresses within that network. In non-authoritative mode, the server listens for DHCP messages and forwards these messages to another server for processing.

In authoritative mode, the switch rejects INT_REBOOT messages from clients with a DHCP NAK message. In non-authoritative mode, the switch sends a DHCP NAK only if a lease is present for the same client, but using a different IP address.

DHCP Server Options Configuration

You can configure DHCP Server to provide clients with standard options such as DNS, NTP, TFTP servers, boot image files, NetBIOS servers, and with customized non-standard options that the server passes to clients.

DHCP Server contains multiple global configuration parameters, which you can define within scopes such as host, subnet, vendor class, or global.

The switch process the DHCP options in the following order of precedence:

1. host—Host configuration applies only to a specific host.
2. DHCP subnet—DHCP subnet options apply to only those clients within the specified subnet.
3. vendor class—Vendor class options apply to only those clients with a string that matches a specified vendor class identifier (VCI).
4. global—Global options apply to all clients.

DHCP Server builds a cumulative configuration from all scopes but applies precedence rules in instances of conflict.

The following table illustrates the resulting configuration specifically for Host A following inheritance and precedence rules for configured DHCP Server options.

Global	Vendor class printer	Subnet 198.51.100.0/24	Host A	Resulting configuration
ip dhcp-server tftp server-ip 192.0.2.100	ip dhcp-server vendor-class tftp server-ip 192.0.2.33	ip dhcp-server tftp server-ip 198.51.100.1	host 00-00-5E-00-53-0 0 tftp server-ip 203.0.113.23	host 00-00-5E-00-53-0 0 tftp server-ip 203.0.113.23
ip dhcp-server tftp server-ip 192.0.2.100	ip dhcp-server vendor-class tftp server-ip 192.0.2.33	None	None	ip dhcp-server vendor-class tftp server-ip 192.0.2.33
ip dhcp-server tftp server-ip 192.0.2.100 ip dhcp-server domain-name saalem	None	ip dhcp-server tftp server-ip 198.51.100.1	None	ip dhcp-server tftp server-ip 198.51.100.1 ip dhcp-server domain-name saalem
ip dhcp-server tftp server-ip 192.0.2.100	None	None	None	ip dhcp-server tftp server-ip 192.0.2.100

For more information see the following sections:

- [DHCP Custom Options](#) on page 821
- [DHCP Vendor Options](#) on page 822

DHCP Custom Options

In addition to standard DHCP options, DHCP Server supports the customization of non-standard parameters that the DHCP Server provides to clients in response to DHCP requests. The following conditions apply:

- You must assign a DHCP code in the unassigned range from 224 to 254.
- You must first define the option. After the option is defined, you can assign a data value to the option globally or for a specific subnet or vendor class. For more information about the different scopes and their precedence in application, see [DHCP Server](#) on page 820.
- In the Custom Option definition, you must assign one of three data types: IPv4 address, string, or integer. After you define a data type, all data values for that option, irrespective of which scope they are assigned in, must use that data type.
- Custom Options do not support the assignment of arrays or the override of standard DHCP options.

DHCP Vendor Options

In addition to DHCP standard and custom options, you can configure up to 255 vendor-specific options. Vendor-specific options always transmit within a single DHCP Option 43.

You can use the **ip dhcp-server vendor-option-def** and **ip dhcp-server vendor-option-data** commands to configure the vendor-specific options to always transmit in DHCP Offer, as shown in the following example configuration:

```
ip dhcp-server vendor-option-def 1 type ipv4-address
ip dhcp-server vendor-option-data 1 value 192.0.2.24
```

However, different client types expect to find different information within Option 43. These clients advertise which type of device they are in DHCP Option 60 so it is more common to specify a vendor-specific option conditional upon a match of a specific client type from Option 60.

In the following configuration example, vendor-specific information transmits within Option 43 if the client reports itself as "WingAP.AP510" in Option 60.

```
ip dhcp-server vendor-option-def 1 type string name XIQ-Controllers
ip dhcp-server vendor-class 1 matches "WingAP.AP510"
ip dhcp-server vendor-class 1 vendor-option-data 1 value "192.0.2.146,192.0.2.143"
```

The vendor-option ID is encoded within Option 43. Multiple vendor-option IDs (up to 255) can be included within the same Option 43. The vendor-class ID is not encoded in any DHCP packet and serves only to link a vendor-option value with a conditional match of an Option 60 device type. The conditional sending of DHCP options can also be used with standard and custom DHCP options.

The configuration shown in the following example, transmits vendor-specific information within Option 43 for clients that report as "IOT.Bot" in Option 60, and sends DHCP standard options to configure the TFTP server IP and Domain-Name, as well as a custom-option 224, which includes an IP address for an IOT Controller:

```
ip dhcp-server vendor-option-def 1 type string name IOT-Controller
ip dhcp-server custom-option-def 224 type ipv4-address name IOT-collection-server
ip dhcp-server vendor-class 1 matches "IOT.Bot"
ip dhcp-server vendor-class 1 vendor-option-data 1 value 192.0.2.146
ip dhcp-server vendor-class 1 domain-name mydomain.com
ip dhcp-server vendor-class 1 tftp server-ip 192.0.2.24
ip dhcp-server vendor-class 1 custom-option-data 224 value 203.0.113.21
```

The following list provides valid data types for vendor-options and custom-options:

- ipv4
- string
- 32-bit integer

For more information about the different scopes and their precedence in application, see [DHCP Server](#) on page 820.

High Availability

You can deploy DHCP Server as a standalone server or in a High Availability (HA) pairing on two switches in a primary-standby configuration. In an HA pair, the primary DHCP Server manages the DHCP service for clients and shares lease information with the standby server. The standby server becomes active only if the primary server is inactive.

To deploy High Availability, perform the following actions:

- Configure DHCP Server on both peer switches. You can assign the primary and standby roles to each switch manually or the switches can assign roles automatically. By default, the switch with the lower IP address is primary.
- Configure the network so that clients can connect to both peer switches.
- Synchronize the clocks on the two HA peer switches using a protocol, such as NTP or configure the system clocks on the two HA peers manually.



Note

For DHCP server synchronization to function correctly, the maximum allowable clock skew is 60 seconds. If the skew exceeds this limit, the High Availability service on each server is terminated.

DHCP Server Deployment Types

The following examples illustrate a single-VRF and multi-VRF deployment for DHCP Server. Both examples illustrate a small campus network where DHCP Relay is on the same switch as DHCP Server. In larger networks, you can configure the relay on a neighboring switch instead of the switch that hosts DHCP Server.

DHCP Server with Single VRF - Small Campus Network

DHCP Server manages local attached VLANs and IP subnets and supports local configuration or cloud configuration of the DHCP Server. This deployment requires that you deploy DHCP Server with DHCP Relay configured on the same switch as the server.

- Configure each VLAN that will use the DHCP service with a Management CLIP IP interface.
- Configure DHCP Server on the management CLIP IP interface of two High Availability (HA) peer switches for redundancy.
- Configure two DHCP Relay forwarding paths; one path to the local management CLIP IP and one path to the management CLIP IP of the peer switch. For small networks, configure the relay on the same switch as DHCP Server.




Figure 45: DHCP Server with a Single VRF - Small Campus Network

DHCP Server with Multiple VRFs - Small Campus Network

Redundant DHCP Servers manage local attached VLANs and IP subnets within multiple VRFs. DHCP Server supports local configuration or cloud configuration.

- Configure DHCP Server on the management CLIP IP interface of two HA peer switches for redundancy.
- Configure DHCP Relay with two forwarding paths (one path to the local management CLIP IP and one path to the management CLIP IP of the peer switch). For smaller networks, configure the relay on the same switch as DHCP Server.
- Configure inter-VRF routes or IS-IS accept policies so that both switches can forward DHCP requests and redundancy communication between VRFs.




Figure 46: DHCP Server with Multiple VRF - Small Campus Network

DHCP Server Configuration Considerations

Consider the following when you configure DHCP Server:

- You must configure at least one DHCP subnet and a management CLIP IP interface before you enable DHCP Server.
- If you deploy multiple VRFs, you must share DHCP pools across all VRFs.
- You must configure local IP routing VLANs with relay configuration to be unique across all VRFs.

- If you associate the VLAN to a different VRF than the management interface, configure inter-VRF routes or IS-IS accept policies so that DHCP requests can relay between VRFs.
- You must configure DHCP Relay. Depending on the size of your network, you can configure the relay on the same switch that hosts DHCP Server or on a neighboring switch. In either case, configure the relay with forwarding paths that point to the management CLIP IP interface of the switch that hosts DHCP Server.

DHCP Relay

With DHCP Relay, the switch can forward DHCP messages from one VLAN or VRF to a DHCP server that resides in a different network thereby increasing the size of the network that the DHCP server (hosted or external) services.

DHCP Relay overcomes a limitation that exists due to the fact that, under the default configuration, the switch does not forward UDP broadcasts to different VLANs or VRFs. If a client broadcasts a DHCP Discover message to request an IP address from any connected DHCP server, the broadcast fails unless there is a DHCP server in the same VLAN. When you configure DHCP Relay, the switch can receive the broadcast in one VLAN or VRF and retransmit the message as a unicast message to the IP address of a remote DHCP server. The switch also relays the server response back to the client.

The following conditions exist for this feature:

- You can configure DHCP Relay whether you deploy the hosted DHCP Server feature or if you deploy an external DHCP server.
- Supports the relay of either DHCP or Bootstrap (BOOTP) messages (default behavior is to relay both message types).
- Supports Smart Relay, which turns DHCP Relay on for secondary addresses also.
- You can configure the relay using global DHCP settings or using port or VLAN settings.
- When you deploy DHCP Relay with the hosted DHCP Server feature, you can configure the relay on the same switch that hosts DHCP Server (local relay) or on a neighboring switch (remote relay). In either case, configure the relay forwarding path to point to the management CLIP IP interface where DHCP Server listens.



Note

In the case of asymmetric routing, when a DHCP message is received from the server, there must be a working forwarding path. You can achieve this in one of two ways:

- Configure a DHCP forwarding path using the actual IP address of that interface.
- Use a VRRP address, in which case, if not the VRRP master, you must enable Backup Master (local ARP entry present).

DHCP Option 82

The DHCP option 82 is the DHCP Relay Agent Information option. The DHCP relay agent inserts option 82 when it forwards the client-originated DHCP packets to a DHCP server. The Relay Agent Information option is organized as a single DHCP option that contains one or more sub-options that convey information known by the relay agent. The DHCP server echoes the option back to the relay agent in server-to-client replies, and the relay agent removes the option before forwarding the reply to the client.

The DHCP option 82 is added at the DHCP relay level as shown in the following image, which illustrates the architecture. This example illustrates an external DHCP server. However, the option is also supported if you are deploying a hosted DHCP Server.




Figure 47: DHCP Client-Relay-Server Architecture (with External DHCP Server)

The Relay Agent Information option (code 82) is a container for specific agent-supplied suboptions; Agent Circuit ID (code 1) and Agent Remote ID (code 2). The suboptions can represent different information relevant for the relay. The fields are encoded in the following manner, where N or n is the total number of octets in the Agent Information Field (all bytes of the suboptions):

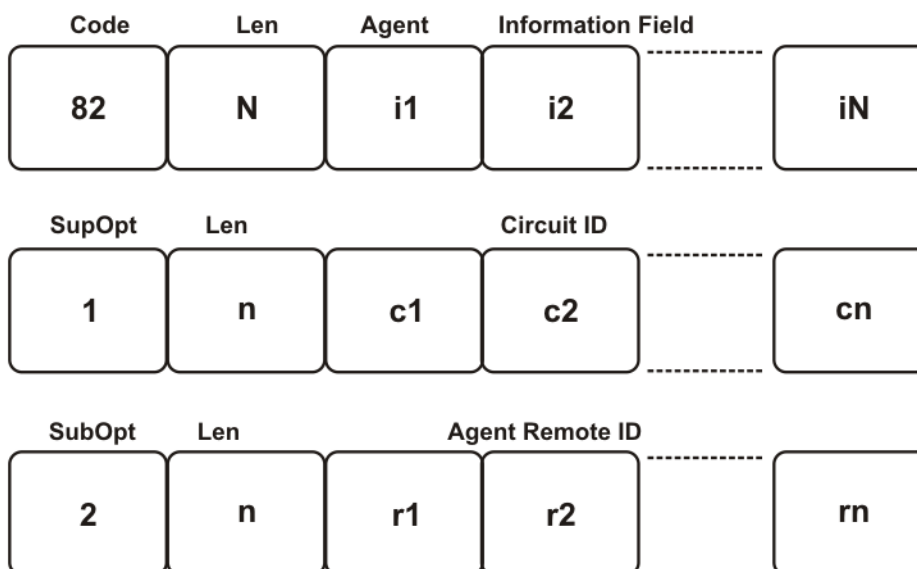


Figure 48: Format of the Relay Agent Information

Because at least one of the sub-options must be defined, the minimum Relay Agent Information length is two (2), and the length n of the suboption can be zero (0). The sub-options do not have to display in any particular order. No pad suboption is defined and the Information field is not terminated with 255 suboption.

DHCP Suboptions

The suboptions are Agent Circuit ID and Agent Remote ID.

The DHCP relay agents can add the Agent Circuit ID to terminate switched or permanent circuits. The Agent Circuit ID encodes an agent-local identifier of the circuit from which a DHCP client-to-server packet was received. Agents can use the Circuit ID to relay DHCP responses back to the proper circuit. In the switch, the Agent Circuit ID field contains the ifIndex of the interface on which the packet is received.

DHCP relay agents can add the Agent Remote ID to terminate switched or permanent circuits, and can identify the remote host end of the circuit. The switch uses the Agent Remote ID field to encode the MAC address of the interface on which the packet is received. The Agent Remote ID must be globally unique.

Agent Operations

A DHCP relay agent adds a Relay Agent Information field as the last option in the DHCP options field of any recognized BOOTP or DHCP packet forwarded from a client to a server. However, if the End Option 255 is present, then the DHCP relay agent adds a Relay Agent information field before the End Option 255 field.

Relay agents can receive a DHCP packet from an untrusted circuit with the gateway IP address (GIADDR) set to zero to indicate that the relay agent is the first-hop router from the gateway. If a Relay Agent Information option is present in the packet, the relay agent discards the packet and increments an error counter. A trusted circuit can contain a trusted downstream network element, for example, a bridge, between the relay agent and the client. The bridge can add a relay agent option but does not set the GIADDR field. In this case, the relay agent forwards the DHCP packet per normal DHCP relay agent operations, and sets the GIADDR field to the relay address. The relay agent does not add a second relay agent option.

You can distinguish between a trusted circuit and an untrusted circuit based on the type of circuit termination equipment you use. To make a circuit trusted, set the trusted flag under DHCP for each interface.

After packets append the Relay Agent Information option, the packets that exceed the MTU or the vendor size buffer of 64 bits, are forwarded without adding the Agent Information option, and an error counter is incremented.

The relay agent or the trusted downstream network element removes the Relay Agent Information option echoed by a server that is added when forwarding a server-to-client response back to the client.

The following list outlines the operations that the relay agent does not perform:

- The relay agent does not add an Option Overload option to the packet or use the file or sname fields to add the Relay Agent Information option. The agent does not parse or remove Relay Agent Information options, the system can display it in the sname or file fields of a server-to-client packet forwarded through the agent.
- The relay agent does not monitor or modify client-originated DHCP packets addressed to a server unicast address; this includes the DHCP-REQUEST sent when entering the RENEWING state.
- The relay agent does not modify DHCP packets that use the IPSEC Authentication Header or IPSEC Encapsulating Security Payload.

A DHCP relay agent can receive a client DHCP packet forwarded from a BOOTP/DHCP relay agent closer to the client. This packet has a GIADDR as non-zero, and may or may not already have a DHCP Relay Agent option in it.

Relay agents configured to add a Relay Agent option which receive a client DHCP packet with a nonzero GIADDR, discards the packet if the GIADDR spoofs a GIADDR address implemented by the local agent itself. Otherwise, the relay agent forwards any received DHCP packet with a valid non-zero GIADDR without adding any relay agent options. The GIADDR value does not change.

DHCP Relay for IPv6

The Dynamic Host Configuration Protocol (DHCP) for IPv6 (RFC 3315) enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. DHCP supports automatic allocation of reusable network addresses and of additional configuration parameters. This protocol is a stateful counterpart to stateless address autoconfiguration, and you can use it separately or concurrently with the

latter to obtain configuration parameters. For more information about stateless address autoconfiguration, see [Host autoconfiguration](#) on page 2265.

To request the assignment of one or more IPv6 addresses, a client first locates a DHCP server, and then requests the assignment of addresses and other configuration information from the server:

1. The client sends a solicit message to the All_DHCP_Relay_Agents_and_Servers (FF02::1:2) multicast address to find available DHCP servers.
2. Any server that can meet the requirements responds with an advertise message.
3. The client then chooses one of the servers and sends a request message to the server asking for confirmed assignment of addresses and other configuration information.
4. The server responds with a reply message that contains the confirmed addresses and configuration.

If a DHCP client does not need a DHCP server to assign it an IPv6 address, the client can obtain configuration information such as a list of available DNS servers or NTP servers through a single message and reply exchanged with a DHCP server.

IPv6 DHCP clients use link-local addresses to send and receive DHCP messages. To permit a DHCP client to send a message to a DHCP server that is not attached to the same link, you must configure a DHCP relay agent on the client link to relay messages between the client and server. The operation of the relay agent is transparent to the client.

A relay agent relays messages from clients and messages from other relay agents. The switch supports DHCP Relay for IPv6. Configure at least one relay agent when the client and server are in different networks.

You must configure the relay agent to use a list of destination addresses for available DHCP servers. The software does not support IPv6 multicast for site-local and global addresses.

The DHCP relay can be a Virtual Router Redundancy Protocol (VRRP) Address. The relay forwards the DHCP messages only if VRRP is in the Master state, otherwise the relay discards the messages.

**Note**

DHCP cannot work on the backup VRRP if the master fails. To achieve optimum results and to leverage redundancy, you must configure DHCP on the backup VRRP.

Clients listen for DHCP messages on UDP port 546. Servers and relay agents listen for DHCP messages on UDP port 547.

Remote ID

IPv6 DHCP Relay supports the remote ID parameter (RFC4649). After you enable remote ID on the switch, the relay agent adds information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The server can use the

supplied information in the process of assigning the addresses, delegated prefixes, and configuration parameters that the client is to receive.

The remote ID option contains two fields:

- vendor ID
- MAC address of the client

The switch uses a vendor ID of 1584.

Limitations

The following list identifies configuration limitations:

- You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.
- The maximum number of servers to which a relay can send a message from one client, is 10.
- You can configure the number of forwarding paths per system. For information on the maximum limit, see *Fabric Engine Release Notes*.
- A DHCP Relay IPv6 is established only between agents within the context of each VRF and when no cross VRF interaction is present.



Note

`All_DHCP_Servers` multicast address option is not implemented for IPv6, as there is no IPv6 MCAST support for site-local and global address.

DHCP Smart Relay for Secondary IP Interfaces

With DHCP Smart Relay functionality, the relay agent inserts both the primary and secondary IP addresses of the client-facing interface (VLAN or VRRP) as the gateway IP address in the DHCP packet. The DHCP relay agent can use secondary IP addresses if there is no reply to the DHCPOFFER message on the server.

For more information about secondary IP addresses, see [Secondary IP Interfaces](#) on page 3992.

DCHP Smart Relay functionality provides two modes:

- parallel

In parallel mode, the relay agent makes multiple copies of DHCP requests and inserts each IP address of the client-facing interface (VLAN or VRRP) as the gateway address in each copy of the request. The switch simultaneously relays multiple packets to the DHCP server, assigning each packet the correct IP address in the correct subnet.

- sequential

In sequential mode, the relay agent forwards received DHCP packets from all IP addresses on the VLAN or VRRP interface one at a time, starting with the

primary address and continuing with each secondary address. Each time a packet is received, the agent selects an address on the VLAN or VRRP interface. The relay agent forwards three client requests to the DHCP Server before repeating the process and cycling through the primary and all secondary IP addresses.

For DHCP Smart Relay to use a subset of VRRP addresses, you must configure the agent IP in the forwarding path with a VRRP address and you must configure at least the following VRRP items on the VLAN:

**Note**

These requirements are not specific to DHCP Smart Relay and also apply to DHCP Relay in general.

- VRRP version
- VRRP address
- VRRP status (enabled)

For configuration information, see the following sections:

- [Configure DHCP Relay on a Port or VLAN](#) on page 853 (CLI)
- [Configure DHCP Relay on an IP Interface](#) on page 892 (EDM)
- [Configure VRRP Using CLI](#) on page 4121
- [Configure VRRP using EDM](#) on page 4142

UDP Broadcast Forwarding

Some network applications, such as the NetBIOS name service, rely on a User Datagram Protocol (UDP) broadcast to request a service or locate a server for an application. If a host is on a network, subnet segment, or VLAN that does not include a server for the service, UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. You can resolve this problem by forwarding the broadcasts to the server through physical or virtual router interfaces.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface out to other router IP interfaces as a rebroadcast or to a configured IP address. If the address is that of a server, the packet is sent as a unicast packet to this address. If the address is that of an interface on the router, the frame is rebroadcast.

After a UDP broadcast is received on a router interface, it must meet the following criteria to be eligible for forwarding:

- It must be a MAC-level broadcast.
- It must be an IP limited broadcast.
- It must be for the specified UDP protocol.
- It must have a time-to-live (TTL) value of at least 2.

For each ingress interface and protocol, the policy specifies how the UDP broadcast is retransmitted: to a unicast host address or to a broadcast address.

DHCP and UDP Configuration Using the CLI

Use Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), to provide host configuration information to the workstations dynamically. Use the DHCP relay commands to configure DHCP relay behavior on a port or on a VLAN.

This section describes CLI commands for DHCP and User Datagram Protocol (UDP) configuration.

DHCP Server Configuration

Complete the following tasks to configure a hosted DHCP Server within your network. With this configuration, the switch hosts the DHCP server.

Before You Begin

- You must configure a management CLIP IP interface.
- Plan the DHCP Server configuration. Many of the parameters are optional and you can apply them either globally, within a subnet, or to a specific host only.

Table 92: DHCP Server CLI configuration task flow

Steps	Procedure	Description
1	Configure DHCP Server General Parameters on page 833	Optional. Configure global DHCP parameters for the DHCP Server.
2	Configure DHCP Custom Options on page 835	Optional. Configure non-standard DHCP options to provide to clients.
3	Configure Vendor Options or a Vendor Class	Optional. Configure a vendor class and vendor options.
4	Configure a DHCP Subnet on page 836	Create DHCP subnets and configure per-subnet values, including an IP address pool.
5	Configure High Availability for the DHCP Server on page 839	Optional. Configure High Availability (HA) for the DHCP server. Repeat the configuration process on the peer switch.
6	Configure a Static DHCP Server Lease on page 839	Optional. Complete this procedure to reserve DHCP configurations, such as a specific IP address, for a static host.
7	Enable DHCP Server on page 841	Enable the DHCP server.

Complete the following optional tasks to configure vendor options and a vendor class.

Table 93: Vendor options and vendor class CLI configuration task flow

	Procedure	Description
1	Define Vendor Options Definitions on page 843	Define DHCP Server vendor options definitions.
2	Configure Vendor Option Data Globally on page 844	Configure DHCP vendor option data global parameters.
3	Configure a Vendor Class on page 845	Configure a vendor class.
4	Configure Custom Option Data for a Vendor Class on page 847	Configure custom option data for the vendor class.
5	Configure Vendor Option Data for a Vendor Class on page 849	Configure a value for the vendor option for the specified vendor class.

What to do Next

[DHCP Relay Configuration](#) on page 850—You must also configure DHCP Relay with forwarding paths that point to the management CLIP IP interface where DHCP Server listens.

Configure DHCP Server General Parameters

Configure DHCP general parameters for the hosted DHCP Server to change default values or specify information to advertise on all subnets. Configure only the parameters that apply to your deployment.

About This Task

It is not mandatory to configure all parameters at the global level. You can also configure many of these parameters for specific subnets or hosts.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the DHCP Server as an authoritative DHCP server:

```
ip dhcp-server authoritative
```

By default, this option is disabled.
3. Configure the global lease expiration:

```
ip dhcp-server lease-time <60-300000000>
```
4. Configure the following optional parameters:
 - a. Configure the DNS servers globally:

```
ip dhcp-server domain-name-servers {A.B.C.D} [...]
```
 - b. Configure the NTP servers globally:

```
ip dhcp-server ntp-servers {A.B.C.D} [...]
```

c. Configure the TFTP server globally:

```
ip dhcp-server tftp server-ip {A.B.C.D}
```

```
ip dhcp-server tftp server-name WORD<0-64>
```

```
ip dhcp-server tftp boot-file-name WORD<0-128>
```

d. Configure a TFTP server list:

```
ip dhcp-server tftp-server-list {A.B.C.D} [... {A.B.C.D}]
```

e. Configure a domain name globally:

```
ip dhcp-server domain-name WORD<0-255>
```

f. Configure global NetBIOS values:

```
ip dhcp-server netbios-name-server {A.B.C.D}
```

```
ip dhcp-server netbios-node-type {0x1|0x2|0x4|0x8}
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server** command.

Variable	Definition
<i>domain-name WORD<0-255></i>	Assigns a domain name to the DHCP Server.
<i>domain-name-servers {A.B.C.D} [...]</i>	Specifies the IPv4 address of one or more DNS servers. You can assign a maximum of eight servers. With multiple servers, list each server address on the same command line with a space between each address entry.
<i>lease-time <60-300000000></i>	Assigns the global lease timer in seconds. The host lease expires when the timer expires. The default is 86400. Note: The DHCP renewal process begins when a lease reaches 50% of the remaining expiration timer. The client unicasts a renewal request to DHCP Server, which checks its database before confirming the renewal.
<i>netbios-name-server {A.B.C.D}</i>	Specifies the IPv4 address of the NetBIOS server.
<i>netbios-node-type {0x1 0x2 0x4 0x8}</i>	Specifies the node type of the NetBIOS server. Enter one of the following values according to which value fits your deployment: <ul style="list-style-type: none"> • 0x1—B node • 0x2—P-node • 0x4—M-node • 0x8—H-node The default setting is 0x0 (unconfigured).
<i>ntp-servers {A.B.C.D} [...]</i>	Specifies the IPv4 address of one or more NTP servers. You can assign a maximum of eight servers. With multiple servers, list each server address on the same command line with a space between each address entry.
<i>tftp boot-file name WORD<0-128></i>	Specifies the file name of the global TFTP boot image file for DHCP Server.

Variable	Definition
<code>tftp server-ip {A.B.C.D}</code>	Specifies the IP address of the global TFTP server for DHCP Server. This parameter uses DHCP Server Option 66. You must assign an IP address to the server before you can assign other TFTP settings such as a server name and boot image file.
<code>tftp server-name WORD<0-64></code>	Specifies the server name of the global TFTP server for DHCP Server.
<code>tftp-server-list {A.B.C.D} [... {A.B.C.D}]</code>	Specifies a list of up to eight TFTP server IP addresses. This parameter uses DHCP Server Option 150, which clients, such as IP Phones, can request from a DHCP Server to obtain configuration files or other information from a TFTP server.

Configure DHCP Custom Options

About This Task

Configure DHCP custom options to deploy non-standard DHCP options:

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Define the custom option by assigning a code, data type, and text label:

```
ip dhcp-server custom-option-def <224-254> type <ipv4-address | uint32 | string> [name WORD<0-128>]
```
3. Assign a data value to the custom option you defined:

```
ip dhcp-server custom-option-data <224-254> value WORD<0-128>
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server** command.

Variable	Definition
<i>custom-option-data</i> <code><224-254> value WORD<0-128></code>	Assigns a data value to a DHCP Custom Option that you defined previously. <ul style="list-style-type: none"> <code><224-254></code>—Specifies the DHCP code that identifies this option. <code>value WORD<0-128></code>—Assigns a data value to the option. The data value must use the data type that is specified in the custom option definition.
<i>custom-option-def</i> <code><224-254> type <ipv4-address uint32 string> [name WORD<0-128>]</code>	Defines the format of a customized option that you want the DHCP Server to provide to clients. <ul style="list-style-type: none"> <code><224-254></code>—Assigns the DHCP code that calls this custom option. Enter an integer in the unassigned range from 224 to 254. <code>type <ipv4-address uint32 string></code>—Assigns the data type for all data values that you assign to this custom option. The available data types are: IPv4 address, integer, or string. <code>name <WORD<0-128></code>—Specifies a text label that identifies this custom option.

Display DHCP Server Custom Options Data

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display DHCP Server custom options data:

```
show ip dhcp-server custom-option-data
```

Example

Display DHCP Server custom options data information:

```
switch:1#show ip dhcp-server custom-option-data
=====
DHCP Server Custom Option Data
=====
SCOPE                CODE    TYPE          VALUE
-----
Global               250    ipv4-address  203.0.113.5
198.51.100.0/24      250    ipv4-address  203.0.113.23
vendor-class-20      246    ipv4-address  192.0.2.24
=====
Total Number of DHCP Server custom-options displayed: 3
=====
```

Configure a DHCP Subnet

Configure parameters for a DHCP subnet. Configure only the parameters that apply to your deployment.

Before You Begin

Before you configure a custom subnet data value, you must define the custom option globally.

About This Task

For subnet options that you leave unconfigured, the global DHCP Server configuration, if one exists, applies to the subnet.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create a DHCP subnet:
`ip dhcp-server subnet {A.B.C.D/X}`
3. Configure the IP address pool that DHCP Server can use to assign IP addresses to clients within this subnet:
`pool {A.B.C.D} {A.B.C.D}`
4. Configure a subnet lease expiration timer:
`lease-time <60-3000000000>`
5. (Optional) Configure the router address to specify the default gateway for subnet clients:
`router {A.B.C.D} [...]`
6. (Optional) Configure subnet DNS servers:
`domain-name-servers {A.B.C.D} [...]`
7. (Optional) Configure a subnet domain name:
`domain-name WORD<1-128>`
8. (Optional) Configure subnet NTP servers:
`ntp-servers {A.B.C.D} [...]`
9. (Optional) Configure subnet TFTP settings:
`tftp server-ip {A.B.C.D}`
`tftp server-name WORD<0-64>`
`tftp boot-file-name WORD<0-128>`
10. Configure a TFTP server list:
`tftp-server-list {A.B.C.D} [... {A.B.C.D}]`
11. (Optional) Configure a subnet NetBIOS name server and node type:
`netbios-name-server {A.B.C.D}`
`netbios-node-type [0x1|0x2|0x4|0x8]`
12. (Optional) Assign a subnet data value to custom options that you defined globally:
`custom-option-data <224-254> value WORD<0-128>`

Variable Definitions

Use the following table to configure DHCP subnets.

Variable	Value
<i>custom-option-data</i> <224-254> <i>value WORD</i> <0-128>	Assigns a data value to an existing DHCP Custom Option. <ul style="list-style-type: none"> <224-254>—Specifies the DHCP code that identifies the option. <i>value WORD</i><0-128>—Assigns a data value to this option. The value must use the data type that was specified in the custom option definition.
<i>domain-name WORD</i> <1-128>	Specifies a domain name for this DHCP subnet.
<i>domain-name-servers</i> {A.B.C.D} [...]	Specifies the IPv4 address of one or more DNS servers. You can assign a maximum of eight server addresses. Enter each server address on the same command line with a space separating each address entry.
<i>lease-time</i> <60-300000000>	Specifies the length of time, in seconds, before a host lease expires. The default value is 0 (unconfigured).
<i>netbios-name-server</i> {A.B.C.D}	Specifies the IPv4 address of the NetBIOS name server for this subnet.
<i>netbios-node-type</i> [0x1 0x2 0x4 0x8]	Specifies the node type of the Net BIOS name server. Enter one of the following options: <ul style="list-style-type: none"> 0x1—B node 0x2—P node 0x4—M node 0x8—H node The default value is 0x0—unconfigured.
<i>ntp-servers</i> {A.B.C.D} [...]	Specifies the IPv4 address of one or more NTP servers. You can assign a maximum of eight server addresses. Enter each server address on the same command line with a space separating each address entry.
<i>pool</i> {A.B.C.D} {A.B.C.D}	Assigns a range of IPv4 addresses that DHCP Server can use to assign addresses to clients. Enter the starting IP address followed by the ending IP address. Note the following: <ul style="list-style-type: none"> The pool is a mandatory setting. The start IP address must be lower than the end IP address. Both IP addresses must be non-zero. Broadcast addresses cannot be leased.
<i>router</i> {A.B.C.D} [...]	Specifies one or more default router IPv4 addresses for this subnet. You can assign a maximum of eight router addresses. Enter each address on the same command line with a space separating each address entry.
<i>subnet</i> {A.B.C.D/X}	Assigns a DHCP subnet that the DHCP Server services. Note that for the subnet address, the host bits must be zero.
<i>tftp boot-file-name</i> <i>WORD</i> <0-128>	Specifies the file name of the TFTP boot image file for the subnet.

Variable	Value
<code>tftp server-ip {A.B.C.D}</code>	Specifies the IPv4 address of the TFTP server for the subnet. This parameter uses DHCP Server Option 66.
<code>tftp server-name</code> <code>WORD<0-64></code>	Specifies the server name of the TFTP server for the subnet.
<code>tftp-server-list {A.B.C.D}</code> <code>[... {A.B.C.D}]</code>	Specifies a list of up to eight TFTP server IP addresses for the subnet. This command uses DHCP Server Option 150, which clients, such as IP Phones, can request from a DHCP Server to obtain configuration files or other information from a TFTP server.

Configure High Availability for the DHCP Server

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the address for the HA peer:

```
ip dhcp-server ha-peer {A.B.C.D}
```
3. (Optional) Configure the role of this switch:

```
ip dhcp-server ha-role {primary | standby}
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server** command.

Variable	Definition
<code>ha-peer {A.B.C.D}</code>	Specifies the IPv4 address of the redundant peer switch in a High Availability (HA) pair.
<code>ha-role {primary standby}</code>	<p>Assigns the High Availability (HA) role of this switch. Within an HA pair, you must have one primary server and one standby server that do the following:</p> <ul style="list-style-type: none"> • <code>primary</code>—This switch provides DHCP services while the switch is active. • <code>standby</code>—This switch provides DHCP services only when the primary switch is inactive. <p>Note: Assigning roles is optional. By default, the switch with the lower IP address is primary.</p>

Configure a Static DHCP Server Lease

Complete this procedure to reserve static DHCP options such as a specific IP address or TFTP boot file for a subnet host. Configure only the options that apply to your deployment.

About This Task

The host configuration overrides conflicting configuration from the subnet or global DHCP Server. For any host configuration that you leave unconfigured, the host inherits the configuration from the subnet.

Procedure

1. Enter DHCP Server Subnet Configuration mode for the applicable subnet:

```
enable
```

```
configure terminal
```

```
ip dhcp-server subnet {A.B.C.D/X}
```

2. Define the static lease with the host MAC address:

```
host <0x00:0x00:0x00:0x00:0x00:0x00>
```

3. Assign an IP address to the host:

```
host <0x00:0x00:0x00:0x00:0x00:0x00> ip-address {A.B.C.D}
```

4. Assign a hostname to the host:

```
host <0x00:0x00:0x00:0x00:0x00:0x00> hostname WORD<0-64>
```

5. Assign a TFTP server and boot file to this host:

```
host <0x00:0x00:0x00:0x00:0x00:0x00> tftp server-ip {A.B.C.D}
```

```
host <0x00:0x00:0x00:0x00:0x00:0x00> tftp server-name WORD<0-64>
```

```
host <0x00:0x00:0x00:0x00:0x00:0x00> tftp boot-file-name WORD<0-128>
```

Variable Definitions

Use the variable definitions in the following table to configure static host reservations within a DHCP subnet.

Variable	Definition
<i>host</i> <i><0x00:0x00:0x00:0x00:0x00:0x00:0x00></i>	Specifies the MAC address for the host.
<i>ip-address {A.B.C.D}</i>	Assigns an IPv4 address to the host. The IP address must be taken from the address pool that is configured for the DHCP subnet where the host resides.
<i>hostname WORD<0-64></i>	Specifies a hostname for the host.
<i>tftp server-ip {A.B.C.D}</i>	Specifies the IP address of the TFTP server that you want to assign to the host. You must assign the TFTP IP address before you can assign a server name or boot file name. Note: For this host, these TFTP settings override conflicting TFTP settings from the subnet or DHCP Server.

Variable	Definition
<code>tftp server-name</code> <code>WORD<0-64></code>	Specifies the server name of the TFTP server that you want the host to use.
<code>tftp boot-file-name</code> <code>WORD<0-128></code>	Specifies the boot image file name that you want the host to use.

Enable DHCP Server

Enable the hosted DHCP Server on the switch.

Before You Begin

- A Management CLIP IP interface must exist.
- A DHCP subnet must have been configured with an IP address pool.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable the DHCP Server:

```
ip dhcp-server enable
```

What to Do Next

If you are deploying High Availability, repeat the full DHCP Server configuration on the peer switch. Otherwise, configure DHCP Relay.

Display DHCP Server Information

Perform this procedure to view existing DHCP Server settings.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. View general DHCP Server information:

```
show ip dhcp-server
```

3. View information on a specific DHCP subnet:

```
show ip dhcp-server subnet {A.B.C.D/X} detail
```

4. View current lease information:

```
show ip dhcp-server leases [ip {a.b.c.d} | subnet {a.b.c.d/x}]
```

5. View information on a specific host:

```
show ip dhcp-server host {mac} detail
```

6. View log files:

```
show ip dhcp-server log
```

Variable Definitions

The following table defines parameters for the **show ip dhcp-server** command.

Variable	Value
<code>subnet {A.B.C.D/X} [detail]</code>	Specifies the subnet network address and mask. Include the optional <i>detail</i> variable to view a more detailed summary.
<code>leases [ip {A.B.C.D} subnet {A.B.C.D/X}]</code>	Displays existing leases for the DHCP Server. To limit the display to specific leases, append one of these optional variables to the command: <ul style="list-style-type: none"> <code>ip {A.B.C.D}</code>—View the lease for a specific IP address. <code>subnet {A.B.C.D/X}</code>—View existing leases within a subnet as specified by the subnet address and mask.
<code>host mac [detail]</code>	Specifies the host for which you want to view information. <ul style="list-style-type: none"> <code>mac</code>—Specifies the MAC address of the host. <code>detail</code>—Include this optional variable to view a more detailed summary.

Clear DHCP Server Leases

Perform this procedure to clear all leases from the hosted DHCP Server.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```
2. Run the following command:

```
clear ip dhcp-server leases
```

Clear DHCP Server Logs

Perform this procedure to clear log files from the hosted DHCP Server.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```
2. Run the following command:

```
clear ip dhcp-server log
```


Define Vendor Options Definitions

About This Task

Perform this procedure to define vendor options definitions.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options. For more information, see [DHCP Server](#) on page 820.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Define the vendor options definitions:

```
ip dhcp-server vendor-option-def <1-255> type {ipv4-address | string |  
uint32} [name WORD<0-128>]
```

Example

Define the vendor options definitions for DHCP server:

```
Switch:1>enable  
Switch:1#config t  
Switch:1(config)#ip dhcp-server vendor-option-def 1 type string  
Switch:1(config)#ip dhcp-server vendor-option-def 2 type uint32 name firmware  
<cr>
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server vendor-option-def** command.

Variable	Definition
<code><1-255></code>	Specifies the vendor option definition.
<code>type ipv4-address uint32 string</code>	Specifies the vendor option definition data type as follows: <ul style="list-style-type: none"> • <code>ipv4address</code>—IPv4 address format • <code>string</code>—text string up to 128 characters • <code>uint32</code>—32-bit integer
<code>name WORD<0-128></code>	Specifies a text label that identifies the vendor option definition.

Display Vendor Options Definitions for DHCP Server

About This Task

Perform this procedure to display vendor options definitions for DHCP Server.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display DHCP vendor options definitions:

```
show ip dhcp-server vendor-option-def
```

Example

Display the DHCP server vendor options definitions information:

```
Switch:1>show ip dhcp-server vendor-option-def
=====
                        DHCP Server Defined Vendor Options
=====
CODE   TYPE           NAME
-----
1      string        -
2      uint32         firmware
-----
```

Configure Vendor Option Data Globally

Before You Begin

Define the vendor option.

About This Task

Perform this procedure to configure DHCP vendor option data global parameters.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options. For more information, see [DHCP Server](#) on page 820.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the vendor option data:

```
ip dhcp-server vendor-option-data <1-255> value WORD<0-128>
```

Example

Configure the vendor option data:

```
Switch:1>enable
Switch:1#config t
Switch:1(config)#ip dhcp-server vendor-option-data 2 value 46
<cr>
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server vendor-option-data** command.

Variable	Definition
<1-255>	Specifies the vendor ID.
value WORD<0-128>	Specifies a value to apply to this vendor option.

Display Vendor Options Data for DHCP Server

About This Task

Perform this procedure to display vendor options data for DHCP Server.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display vendor options data:

```
show ip dhcp-server vendor-option-data
```

Example

Display DHCP Server vendor option data information:

```
Switch:1#show ip dhcp-server vendor-option-data
=====
                        DHCP Server Vendor Option Data
=====
SCOPE                CODE   TYPE      VALUE
-----
Global                1     string   voss
Global                2     uint32    46
vendor-class-20       2     uint32    1337
=====
Total Number of DHCP Server custom-options displayed: 4
=====
```

Configure a Vendor Class

About This Task

Perform this procedure to configure a vendor class. If the string matches the vendor class identifier (VCI) provided by the DHCP client, then the DHCP client is part of that vendor class.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options. For more information, see [DHCP Server](#) on page 820.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create the vendor class that matches the specified VCI:

```
ip dhcp-server vendor-class <1-255> matches WORD<1-128>
```

3. (Optional) Configure the TFTP server parameters that you want the vendor class to use:

```
ip dhcp-server vendor-class <1-255> tftp server-ip {A.B.C.D}
```

```
ip dhcp-server vendor-class <1-255> tftp server-name WORD<0-64>
```

```
ip dhcp-server vendor-class <1-255> tftp boot-file-name WORD<0-128>
```

4. (Optional) Configure a domain name that you want the vendor class to use:

```
ip dhcp-server vendor-class <1-255> domain-name WORD<1-128>
```

Example

Configure a vendor class:

```
Switch:1>enable
Switch:1#config t
Switch:1(config)#ip dhcp-server vendor-class 20 matches printer
<cr>
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server vendor-class** command.

Variable	Definition
<i>WORD<1-255></i>	Specifies the vendor class ID.
<i>domain-name WORD<1-128></i>	Specifies the domain name advertised on the vendor class.
<i>matches WORD<1-128></i>	Specifies the vendor class identifier for the vendor class.
<i>server-name WORD<0-64></i>	Specifies the server name of the TFTP server advertised on the vendor class.
<i>tftp boot-file-name WORD<0-128></i>	Specifies the boot file name.
<i>tftp server-ip {A.B.C.D}</i>	Specifies the IPv4 address of the TFTP server advertised on the vendor class.
<i>vendor-option-data <1-255></i>	Specifies the file name of the TFTP boot image advertised on the vendor class.
<i>value WORD<0-128></i>	Specifies the data value for the vendor class.

*Display DHCP Server Vendor Class Information***Procedure**

1. To enter User EXEC mode, log on to the switch.
2. Display DHCP Server vendor class information:
show ip dhcp-server vendor-class [detail]

Example

Display the DHCP Server vendor class information:

```
Switch:1>show ip dhcp-server vendor-class
=====
DHCP Server Vendor Class Info
=====
VENDOR CLASS      MATCHES
-----
vendor-class-4
telephone
vendor-class-20    printer
-----
Total Number of DHCP Server vendor classes displayed: 2
=====
```

Display DHCP Server detailed vendor class information:

```
Switch:1#show ip dhcp-server vendor-class detail
=====
DHCP Server Vendor Classes Detailed Info
=====

Vendor Class 4
=====
Matches           : telephone
TFTP Server IP     : 192.0.2.10
TFTP Server Name   : dial.extremenetworks.com
TFTP Bootfile-Name : sim.cfg
Domain Name        :

Vendor Class 20
=====
Matches           : printer
TFTP Server IP     : 0.0.0.0
TFTP Server Name   :
TFTP Bootfile-Name :
Domain Name        : printer.extreme.com
Vendor-Option-2    : 1337
-----
Total Number of DHCP Server vendor classes displayed: 2
=====
```

*Configure Custom Option Data for a Vendor Class***Before You Begin**

- Create the vendor class.
- Define the custom option.

About This Task

Perform this procedure to configure custom vendor option data for the vendor class.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options. For more information, see [DHCP Server](#) on page 820.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Configure custom option data for the vendor class:

```
ip dhcp-server vendor-class <1-255> custom-option-data <224-254> value
WORD<0-128>
```

Example

Configure custom option data for the vendor class:

```
Switch:1>enable
Switch:1#config t
Switch:1(config)#ip dhcp-server vendor-class 20 custom-option-data 246 value 192.0.2.24
<cr>
```

Variable Definitions

The following table defines parameters for the **ip dhcp-server vendor-class** command.

Variable	Definition
<i>WORD<1-255></i>	Specifies the vendor class ID.
<i>custom-option-data</i> <i>WORD<224-254></i>	Specifies the DHCP code that identifies the custom option.
<i>value WORD<0-128></i>	Specifies a data value to apply to this vendor class.

Display DHCP Server Custom Options Data

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display DHCP Server custom options data:

```
show ip dhcp-server custom-option-data
```

Example

Display DHCP Server custom options data information:

```
switch:1#show ip dhcp-server custom-option-data
=====
                        DHCP Server Custom Option Data
=====
SCOPE                CODE   TYPE                VALUE
```

```

-----
Global                250    ipv4-address    203.0.113.5
198.51.100.0/24       250    ipv4-address    203.0.113.23
vendor-class-20       246    ipv4-address    192.0.2.24
-----
Total Number of DHCP Server custom-options displayed: 3
-----

```

Configure Vendor Option Data for a Vendor Class

Before You Begin

- Define the vendor option.
- Create the vendor class.

About This Task

Perform this procedure to configure vendor option data for the specified vendor class.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options. For more information, see [DHCP Server](#) on page 820.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure vendor option data for a vendor class:

```
ip dhcp-server vendor-class <1-255> vendor-option-data <1-255> value
WORD<0-128>
```

Example

Configure vendor option data for a vendor class:

```

Switch:1>enable
Switch:1#config t
Switch:1(config)#ip dhcp-server vendor-class 20 vendor-option-data 2 value 1337
<cr>

```

Variable Definitions

The following table defines parameters for the **ip dhcp-server vendor-class** command.

Variable	Definition
<1-255>	Specifies the vendor class ID.
vendor-option-data WORD<224-254>	Specifies the DHCP code that identifies the vendor option.
value WORD<0-128>	Specifies a data value to apply to this vendor class.

DHCP Relay Configuration

Complete the following tasks to configure DHCP Relay on the switch.

Table 94: DHCP Relay Configuration

Steps	Task	Description
1	Configure DHCP Relay Parameters Globally on page 850	Configure DHCP forwarding path entries as part of the DHCP Relay global configuration.
2	Configure DHCP option 82 on page 852	(Optional). Enable the circuit ID to encode an agent-local identifier of the circuit from which a DHCP client-to-server packet is received.
3	Configure DHCP Relay on a Port or VLAN on page 853	Enable DHCP Relay on a port or VLAN interface.

Configure DHCP Relay Parameters Globally

Configure DHCP Relay global parameters that will be used by ports or VLANs that have the DHCP Relay feature enabled.

Before You Begin

Configure an IP address on the interface to be used as the DHCP relay interface.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Create the forwarding path from the client to the server:

```
ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D>
```
- Enable the forwarding path from the client to the server:

```
ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> enable
```



Note

- (Optional) Modify DHCP mode to forward BOOTP messages only, DHCP messages only, or both:

```
ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> mode <bootp|bootp_dhcp|dhcp>
```


5. (Optional) Configure the forwarding path with source port 67 from client to the server.

```
ip dhcp-relay fwd-path {A.B.C.D} {A.B.C.D} src-port-67
```



Note

If the agent IP address (the first `<A.B.C.D>` variable) is a VLAN or port IP address, you must also enable DHCP Relay on that VLAN or port by running **ip dhcp-relay** within the VLAN context. However, if the first `<A.B.C.D>` variable is a VRRP address, you do not need to enable DHCP Relay on the VLAN or port in which the VRRP address resides. See [Configure DHCP Relay on a Port or VLAN](#) on page 853.

Example

Create the forwarding path from the client to the server. Enable the forwarding path from the client the server. Modify DHCP mode to forward both BOOTP and DHCP messages. Configure the forwarding path with source port 67 for BOOTP request.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip dhcp-relay fwd-path 192.0.2.120 192.0.2.50
Switch:1(config)#ip dhcp-relay fwd-path 192.0.2.128 192.0.2.50 enable
Switch:1(config)#ip dhcp-relay fwd-path 192.0.2.128 192.0.2.50 mode bootp_dhcp
Switch:1(config)#ip dhcp-relay fwd-path 192.0.2.128 192.0.2.50 src-port-67
```

Variable Definitions

The following table defines parameters for the **ip dhcp-relay fwd-path** command.

Variable	Value
<code>{A.B.C.D}</code>	Configures the forwarding path from the client to the server. The first <code>{A.B.C.D}</code> variable is the agent IP address configured on an interface (a locally configured IP address). You cannot specify a Secondary IP Interface address as the agent IP address in the forwarding path. To make DHCP aware of Secondary IP Interfaces, you must enable DHCP Smart Relay. The second <code>{A.B.C.D}</code> variable is the IP address of the DHCP server in the network.
<code>disable</code>	Disables DHCP Relay globally.
<code>enable</code>	Enables DHCP Relay globally.
<code>mode {bootp bootp_dhcp dhcp></code>	Modifies DHCP mode to forward BOOTP messages only, DHCP messages only, or both. The default is both.
<code>src-port-67</code>	Configures the UDP source port to 67 for BOOTP request. The default is 68.

Configure DHCP option 82

Configure the DHCP option 82 to enable the circuit ID to encode an agent-local identifier of the circuit from which a DHCP client-to-server packet is received. Configure the DHCP option 82 to enable the remote ID to encode the MAC address of the interface on which the packet is received. By default, the DHCP option 82 is disabled.

Before You Begin

- You must enable ip and dhcp-relay on the VLAN.

About This Task

To configure the DHCP option 82 on a VLAN, you must enter the VLAN Interface Configuration mode.

To configure the DHCP option 82 on a brouter port, you must enter the GigabitEthernet Interface Configuration mode.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable the circuit ID:

```
ip dhcp-relay circuitID
```

3. Enable the remote ID:

```
ip dhcp-relay remoteID
```

4. Configure the circuit as trusted:

```
ip dhcp-relay trusted
```

5. Show statistics for option 82, which is the relay agent information option:

```
show ip dhcp-relay counters option82 [vrf WORD<1-16>] [vrfids WORD <0-512>]
```

Example

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)# interface gigabitethernet 1/10
```

Enable the circuit ID:

```
Switch:1(config-if)# ip dhcp-relay circuitID
```

Enable the remote ID:

```
Switch:1(config-if)# ip dhcp-relay remoteID
```

Configure the circuit as trusted:

```
Switch:1(config-if)# ip dhcp-relay trusted
```

Show statistics for option 82, which is the relay agent information option:

```
Switch:1(config-if)# show ip dhcp-relay counters option82
```

Variable Definitions

Use the data in the following table to configure the DHCP option 82.

Variable	Value
<i>circuitID</i>	Enables the Circuit ID.
<i>remoteID</i>	Enables the Remote ID.
<i>trusted</i>	Sets the circuit as trusted.

Use the data in the following table to use the **show ip dhcp-relay counters option82 [vrf WORD<1-16>] [vrfids WORD <0-512>]** command.

Variable	Value
<i>vrf WORD<1-16></i>	Displays DHCP counters for a particular VRF. WORD<1-16> specifies the VRF name.
<i>vrfids WORD <0-512></i>	Displays a DHCP forward path for a particular VRF. WORD <0-512> specifies the VRF ID.

Configure DHCP Relay on a Port or VLAN

You can view and configure the DHCP parameters on a specific port or VLAN.

Before You Begin

- You must configure an IP interface on the interface.

Procedure

- Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable DHCP parameters on a specified port or VLAN:

```
ip dhcp-relay
```

3. (Optional) Enable DHCP Smart Relay on a VLAN:

```
ip dhcp-relay include-secondary <parallel | sequential>
```

4. (Optional) Configure all other optional parameters as required.

For more information see, [Variable Definitions](#) on page 854.

Variable Definitions

Use the data in the following table to use the **ip dhcp-relay** command.



Note

The **no ip dhcp-relay** command disables DHCP Relay, it does not delete the DHCP entry.

Variable	Value
<i>broadcast</i>	Enables the device to send the server reply as a broadcast to the end station. After you disable this variable, the device sends the server reply as a unicast to the end station.
<i>circuitId</i>	Enables Option 82 circuit ID on the interface.
<i>clear-counters</i>	Clears DHCP Relay counters for the interface.
<i>fwd-path <A.B.C.D> [vrid <1-255>]</i>	Creates a forward path server with a virtual router ID (or VRRP ID), a mode, and a state. <i>A.B.C.D</i> is the IP address. <i>vrid <1-255></i> is the ID of the virtual router and is an integer from 1 to 255.
<i>fwd-path <A.B.C.D> disable [vrid <1-255>]</i>	Disables a forward path server with a specific value and virtual router ID. <i>A.B.C.D</i> is the IP address. <i>vrid <1-255></i> is the ID of the virtual router (or VRRP ID).
<i>fwd-path <A.B.C.D> enable [vrid <1-255>]</i>	Enables a forward path server with a specific value and virtual router ID (or VRRP ID). <i>A.B.C.D</i> is the IP address. <i>vrid <1-255></i> is the ID of the virtual router (or VRRP ID).
<i>fwd-path <A.B.C.D> mode <bootp bootp_dhcp dhcp> [vrid <1-255>]</i>	Configures the forward path mode for a VLAN. This command string is available only in VLAN Interface Configuration mode. <i>A.B.C.D</i> is the IP address. <i>vrid <1-255></i> is the ID of the virtual router (or VRRP ID).
<i>max-hop <1-16></i>	Configures the maximum number of hops before a BootP/DHCP packet is discarded. The default is 4.

Variable	Value
<code>min-sec <0-65535></code>	Configures the minimum seconds count for DHCP. If the secs field in the BootP/DHCP packet header is greater than this value, the device relays or forwards the packet; otherwise, the packet is dropped. The default is 0 seconds.
<code>mode <bootp bootp_dhcp dhcp></code>	Configures DHCP mode to forward BootP messages only, DHCP messages only, or both. The default is both.
<code>remoteId</code>	Enables Option82 remote ID on the interface.
<code>trusted</code>	Configures the DHCP circuit as trusted.
<code>include-secondary <parallel sequential></code>	<p>Enables DHCP Smart Relay in one of the following modes:</p> <ul style="list-style-type: none"> parallel— the relay agent makes multiple copies of DHCP requests and inserts each IP address of the client- facing interface (VLAN or VRRP) as the gateway address in each copy of the request. The switch simultaneously relays multiple packets to the DHCP server, assigning each packet the correct IP address in the correct subnet. sequential—the relay agent forwards received DHCP packets from all IP addresses on the VLAN or VRRP interface one at a time, starting with the primary address and continuing with each secondary address. <p>By default, Smart Relay is disabled.</p> <p>Note: This parameter is optional and applies only to a VLAN interface.</p>

Show DHCP Relay Information

Perform this procedure to show relay information about DHCP routes and counters.

For scaling information on DHCP Relay forwarding (IPv4 or IPv6), see *Fabric Engine Release Notes*.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about DHCP relay forward paths:
show ip dhcp-relay fwd-path [vrf WORD<1-16>] [vrfids WORD<0-512>]
3. Display information about DHCP relay counters:
show ip dhcp-relay counters [vrf WORD<1-16>] [vrfids WORD<0-512>]

4. Display the options for each listed interface:

```
show ip dhcp-relay interface [gigabitethernet {slot/port[/sub-port]}[-
slot/port[/sub-port]][,...][slot/all][all]} [vlan <1-4059>] [vrf WORD
<1-16>] [vrfids WORD <0-512>]
```

Example

Display information for all DHCP relay interfaces:

```
Switch:1>show ip dhcp-relay interface
```

```
=====
                        Port Dhcp
=====
PORT      VRF          MAX MIN      ALWAYS  CIRCUIT  REMOTE  TRUST
NUM      NAME          ENABLE HOP SEC   MODE    BCAST   ID      ID      CIRC
-----
1/5      GlobalRouter    false  4    0    dhcp & bootp false  false  false  false

=====
                        Vlan Dhcp
=====
VLAN VRF          MAX MIN      ALWAYS  CIRCUIT  REMOTE  TRUST  SMART
ID   NAME          ENABLE HOP SEC   MODE    BCAST   ID      ID      CIRC   RELAY
-----
2    GlobalRouter    false  4    0    dhcp & bootp false  false  false  false  disabled
5    GlobalRouter    true   4    0    dhcp & bootp false  false  false  false  disabled
```

```
All 2 out of 2 of Vlan Dhcp Entries displayed
```

Variable definitions

Use the data in the following table to use the **show ip dhcp-relay** command.

Variable	Value
<i>vrf</i> WORD<1-16>	The name of the VRF.
<i>vrfids</i> WORD<0-512>	The ID of the VRF. The value is an integer in the range of 0–512.

Use the data in the following table to use the **show ip dhcp-relay interface** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>[vlan <1-4059>]</code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>[vrf WORD<1-16>]</code>	Specifies the name of the VRF.
<code>[vrfids WORD<0-512>]</code>	Specifies the ID of the VRF. The value is an integer from 0–512.

Display DHCP-Relay Statistics for Specific Ports

Display individual DHCP-relay statistics for specific ports to manage network performance.



Note

Slot and port information can differ depending on hardware platform.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View DHCP-relay statistics for a specific port or VRF.
show interfaces GigabitEthernet statistics dhcp-relay [vrf WORD<1-16>]
[vrfids WORD<0-255>] [{slot/port[-slot/port] [,...] [slot/all] [all]}

Example

View DHCP-relay statistics:

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet statistics dhcp-relay

=====
                        Port Stats Dhcp
=====
PORT_NUM VRF NAME          NUMREQUEST NUMREPLY
-----
```

1/12	GlobalRouter	0	2
1/13	GlobalRouter	3	2
2/3	GlobalRouter	0	2

Variable Definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics dhcp-relay** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF instance by VRF name.
<code>vrfids WORD<0-255></code>	Specifies the ID of the VRF.
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats: a single slot and port (1/1).</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display DHCP Relay Statistics for All Interfaces

About This Task

Display DHCP-relay statistics for all interfaces to manage network performance.



Note
Slot and port information can differ depending on hardware platform.

Procedure

1. Show the number of requests and replies for each interface:
`show ip dhcp-relay counters [vrf WORD<1-16>] [vrfids WORD<0-512>]`
2. Show counters for Option 82:
`show ip dhcp-relay counters option82 [vrf WORD<1-16>] [vrfids WORD<0-512>]`

Example

```
Switch:1>show ip dhcp-relay counters option82
=====
DHCP Counters Option82 - GlobalRouter
=====
```


INTERFACE	IP ADDR	FOUND OP82	DROP PKT	CIRC ID	ADD CIRC	DEL CIRC	REMOTE ID	ADD REMID	DEL REMID	ACCESS
Port 1/12	192.0.1.10	0	0	395	0	0	00:24:7f:9d:0a:00	0	0	GlobalRouter
Vlan40	192.0.1.20	0	0	2088	0	0	00:24:7f:9d:0a:01	0	0	GlobalRouter

Variable Definitions

Use the data in the following table to use the **show ip dhcp-relay counters** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies a VRF instance by the VRF name.
<i>vrfids WORD<0-512></i>	Specifies the ID of the VRF.

UDP Broadcast Forwarding Configuration

About This Task

By default, routers do not forward broadcasts. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts. You must setup UDP broadcast forwarding on the system. Configure UDP broadcast forwarding to forward the UDP broadcasts of network applications to the required server through physical or virtual router interfaces.

Procedure

1. Enter protocols into a table.
2. Create policies (protocol/server pairs).
3. Assemble the policies into lists or profiles.
4. Apply the list to the appropriate interfaces.

Configure UDP protocols

About This Task

Configure UDP protocols to determine which UDP broadcasts are forwarded.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure a UDP protocol:

```
ip forward-protocol udp <1-65535> WORD<1-15>
```

3. Confirm your configuration:

```
show ip forward-protocol udp interface [vrf WORD<1-16>][vrfids
WORD<0-512>] portfwd [vrf WORD<1-16>][vrfids WORD<0-512>]
portfwdlist <1-1000>[vrf WORD<1-16>][vrfids WORD<0-512>] vrf
WORD<1-16> vrfids WORD<0-512>
```

Example

Configure a UDP protocol and confirm your configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip forward-protocol udp 53 DNS
Switch:1(config)#show ip forward-protocol udp

=====
                        Udp Protocol Tbl - GlobalRouter
=====
UDP_PORT  PROTOCOL_NAME
-----
37         Time Service
49         TACACS Service
53         DNS
69         TFTP
137        NetBIOS NameSrv
138        NetBIOS DataSrv
```

Variable Definitions

Use the data in the following table to use the **ip forward-protocol udp** command.

Variable	Value
<1-65535> WORD<1-15>	Creates a new UDP protocol. <1-65535> WORD<1-15> is the UDP protocol name as a string.
portfwd	Displays portfwd information.
portfwdlist	Displays port forward list information.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF.

Configure a UDP Port Forward Entry

Configure a UDP port forward entry to add or remove a port forward entry.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure a UDP port forward entry:

```
ip forward-protocol udp portfwd <1-65535> {A.B.C.D}
```

3. Confirm your configuration:

```
show ip forward-protocol udp portfwd [vrf WORD<1-16>] [vrfids  
WORD<0-512>]
```

Example

Configure a UDP port forward entry:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip forward-protocol udp portfwd 150 192.0.2.10
```

Variable Definitions

Use the data in the following table to use the **ip forward-protocol udp portfwd** command.

Variable	Value
<code><1-65535> {A.B.C.D}</code>	Adds a UDP protocol port to the specified port forwarding list. <i>1-65535</i> is a UDP protocol port in the range of 1-65535. <i>A.B.C.D</i> is an IP address in a.b.c.d format.
<code>vrf WORD<1-16></code>	Specifies the name of the VRF.
<code>vrfids WORD<0-512></code>	Specifies the ID of the VRF.

Configure the UDP Port Forwarding List

Configure the UDP port forwarding list to assign protocols and servers to the port forward list.

About This Task

You can perform this procedure in Global Configuration mode, VLAN Interface Configuration mode, or VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the UDP port forwarding list:

```
ip forward-protocol udp portfwdlist <1-1000>
```



Important

The following two steps are not available in the Global Configuration or VRF Router Configuration mode. The following two commands are available in VLAN Interface Configuration mode only.

3. Enter VLAN Interface Configuration mode:

```
interface vlan <1-4059>
```

4. Configure the broadcast mask:

```
ip forward-protocol udp broadcastmask {A.B.C.D}
```

5. Configure the maximum time to live:

```
ip forward-protocol udp maxttl <1-16>
```

6. Confirm your configuration:

```
show ip forward-protocol udp portfwdlist <1-1000> [vrf WORD<1-16>]  
[vrfids WORD<0-512>]
```

Example

Configure the UDP port forwarding list:

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)#ip forward-protocol udp portfwdlist 1
```

Log on to the VLAN interface configuration mode:

```
Switch:1(config)# interface vlan 3
```

Configure the broadcast mask:

```
Switch:1(config-if)# ip forward-protocol udp broadcastmask 192.0.2.255
```

Configure the maximum time to live:

```
Switch:1(config-if)# ip forward-protocol udp maxttl 10
```

Confirm the configuration:

```
Switch:1(config-if)# show ip forward-protocol udp portfwdlist
```

Variable Definitions

Use the data in the following table to use the **ip forward-protocol udp portfwdlist** command.

Variable	Value
<code><1-1000></code>	Creates a UDP port forwarding list in the range of 1-1000.
<code><1-65535> {A.B.C.D}</code>	Adds a UDP protocol port to the specified port forwarding list. <i>1-65535</i> is a UDP protocol port in the range of 1-65535. <i>A.B.C.D</i> is an IP address in a.b.c.d format.
<code>name WORD<0-15></code>	Changes the name of the port forwarding list.

Use the data in the following table to use the **ip forward-protocol udp** command.

Variable	Value
<code>broadcastmask {A.B.C.D}</code>	Configures the interface broadcast mask (the interface broadcast mask can be different from the interface mask). <i>A.B.C.D</i> is an IP address in a.b.c.d format.
<code>maxttl <1-16></code>	Configures the maximum time-to-live value (TTL) for the UDP broadcast forwarded by the interface. The range is 1-16.
<code>portfwdlist <1-1000></code>	Assigns the list to the VLAN.
<code>vlan <1-4059></code> <code>[portfwdlist <1-1000>]</code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. If you use the portfwdlist variable with the VLAN variable, it assigns the list to the specified VLAN, regardless of which VLAN context you currently configure.

Display UDP Forward Information

Display information about the UDP forwarding characteristics of the device. UDP forwarding only supports 128 entries.

Procedure

1. Enter Privileged EXEC mode:
enable

2. Display information about the UDP interface for all IP addresses or a specified IP address:

```
show ip forward-protocol udp interface [<A.B.C.D>] [vrf WORD<1-16>]
[vrfids WORD<0-512>]
```

3. Display the UDP port forwarding table:

```
show ip forward-protocol udp portfwd [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```

4. Display the UDP port forwarding list table for the specified list or all lists on the device:

```
show ip forward-protocol udp portfwdlist [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```

5. Display the UDP protocol table with the UDP port numbers for each supported or designated protocol:

```
show ip forward-protocol udp [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Display the UDP protocol table with the UDP port numbers for each supported or designated protocol:

```
Switch:1>enable
Switch:1#show ip forward-protocol udp

=====
                        Udp Protocol Tbl - GlobalRouter
=====
UDP_PORT  PROTOCOL_NAME
-----
37        Time Service
49        TACACS Service
53        DNS
69        TFTP
137       NetBIOS NameSrv
138       NetBIOS DataSrv
```

Variable Definitions

Use the data in the following table to use the **show ip forward-protocol udp interface** command.

Variable	Value
<A.B.C.D>	Specifies the IP address for the interface in a.b.c.d format.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0 to 512.

IPv6 DHCP Relay Configuration using CLI

This section provides procedures to configure IPv6 DHCP relay using the command line interface (CLI).

Configure a DHCP Relay Forwarding Path

Configure a forwarding path to specify the relay agent address and the DHCP server address to which to forward packets.

To use DHCP Relay for IPv6, you must configure at least one forwarding path and enable the relay on one interface.

Before You Begin

For a VRF other than GlobalRouter, the interface must be first associated to that VRF.

About This Task

The relay agent can use the IPv6 address of the interface or the VRRP global address linked to that interface. The relay forwards the DHCP messages only if VRRP is in the master state, otherwise the relay discards the messages.

You can configure only one relay agent on an interface. If you need to change the relay agent, you must delete all the forwarding paths with the old relay agent, and then configure the new relay agent.

For scaling information on DHCP Relay forwarding paths, see *Fabric Engine Release Notes*.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure a forwarding path:

```
ipv6 dhcp-relay fwd-path WORD<0-255> WORD<0-255> [enable]
```

If you configure the forwarding path globally, the relay agent address can be any configured IP address of the relay interface or the VRRP global address linked to the relay interface.

3. To configure a forwarding path on an interface, enter Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

OR

```
interface vlan <1-4059>
```

4. Configure a forwarding path:

```
ipv6 dhcp-relay fwd-path WORD<0-255> [enable] [vrid WORD<1-255>]
```

If you configure the forwarding path on an interface, the relay agent address is either the smallest IP configured on the interface or the first VRRP global address configured, if the relay is the VRRP master. You do not specify the relay agent address as part of the command.

**Note**

IPv6 DHCP Relay is established only between agents within the context of each VRF.

Examples

Configure a forwarding path globally:

```
Switch:1(config)#ipv6 dhcp-relay fwd-path 1111::1111 1234::1234 enable
```

Configure a forwarding path on an interface:

```
Switch:1(config)#interface GigabitEthernet 1/1
Switch:1(config-if)#ipv6 dhcp-relay fwd-path 1234::1234 enable
```

Configure the VRRP master as the relay:

```
Switch:1(config-if)#ipv6 dhcp-relay fwd-path 1234::1234 vrid 12 enable
```

Variable Definitions

Use the data in the following table to use the **ipv6 dhcp-relay fwd-path** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>enable</i>	Enables the forwarding path. The default is disabled.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>vrid WORD<1-255></i>	Specifies the VRRP ID to use the VRRP master as the relay agent interface.

Variable	Value
<i>WORD</i> <0-255>	Specifies the IPv6 address of the DHCP server for the interface configuration.
<i>WORD</i> <0-255> <i>WORD</i> <0-255>	Specifies the IPv6 address of the relay agent interface and the IPv6 address of the DHCP server for the global configuration.

Configure DHCP Relay for an Interface

Configure the DHCP relay behavior on the interface.

About This Task

You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable DHCP on the interface:

```
ipv6 dhcp-relay
```

3. Configure the maximum hop count:

```
ipv6 dhcp-relay max-hop <1-32>
```

4. Enable the remote ID:

```
ipv6 dhcp-relay remote-id
```

Example

Configure the maximum hop count and disable the remote ID:

```
Switch:1(config-if)#ipv6 dhcp-relay max-hop 30  
Switch:1(config-if)#no ipv6 dhcp-relay remote-id
```

Variable Definitions

Use the data in the following table to use the **ipv6 dhcp-relay** command.

Variable	Value
<i>max-hop</i> <1-32>	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
<i>remote-id</i>	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled

Use the data in the following table to use the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View DHCP Relay Information

View DHCP Relay information to display the current configuration for the forwarding path and the interface configuration.

About This Task

Not all parameters are available in non-default VRFs.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View DHCP Relay global information:

```
show ipv6 dhcp-relay {counters [vrf WORD<1-16> | vrfids WORD<0-512>] |
fwd-path [vrf WORD<1-16> | vrfids WORD<0-512>]}
```

3. View IPv6 DHCP Relay interface configuration:

```
show ipv6 dhcp-relay interface {gigabitEthernet {slot/port[/sub-port]}
[-slot/port[/sub-port]][,...][slot/all][all]} | vlan <1-4059>}
```

**Note**

The **no ipv6 dhcp-relay** command disables DHCP on the interface but does not delete the entry.

Example

```
Switch:1(config-if)#show ipv6 dhcp-relay fwd-path
```

```
=====
                        DHCPv6 Fwd-path - GlobalRouter
=====
INTERFACE                SERVER                ENABLE
-----
1111:0:0:0:0:0:1111      1234:0:0:0:0:0:1234      enable
```

Variable Definitions

Use the information in the following table to help you use the **show ipv6 dhcp-relay** command.

Variable	Value
<code>{slot/port[/sub-port]} [-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>vlan<1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

View IPv6 DHCP Relay Statistics

Display individual IPv6 DHCP Relay statistics for specific interfaces to manage network performance.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View statistics:

```
show ipv6 dhcp-relay counters
```

**Note**

Use the `sys action reset counters` command to clear DHCP Relay statistics.

Example

```
Switch:1#show ipv6 dhcp-relay counters
```

```
=====
                        DHCPv6 Counters
=====
INTERFACE                REQUESTS    REPLIES
-----
1111:0:0:0:0:0:1111      1         1
```

DHCP and UDP Configuration Using EDM

Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), dynamically provides host configuration information to workstations. To lower administrative overhead, network managers prefer to configure a small number of DHCP servers in a central location. Using few DHCP servers requires the routers connecting to the subnets or bridge (or VLAN) domains to support the BootP/DHCP relay function so that hosts can retrieve the configuration information from servers several router hops away.

User datagram protocol (UDP) is a connectionless protocol that adds reliability and multiplexing to IP. It describes how messages reach application programs within a destination computer. Some network applications, such as the NetBIOS name service, rely on a UDP broadcast to request a service or to locate a service. By default, broadcasts are not forwarded by a router. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts.

**Important**

BootP/DHCP relays are supported only on IP routed port-based VLANs and protocol-based VLANs.

Before You Begin

You must enable DHCP relay on the path for port or VLAN configuration to take effect.

DHCP Server Configuration using EDM

Complete the following tasks to configure a hosted DHCP Server within your network. With this configuration, the switch hosts the DHCP server.

Before You Begin

- You must configure a management CLIP IP interface.
- Plan the DHCP Server configuration. Many of the parameters are optional and you can apply them either globally, within a subnet, or to a specific host only.

Table 95: DHCP Server EDM configuration task flow

Step	Procedure	Description
1	Configure the DHCP Server Globally on page 872	Configure global values for the DHCP Server.
2	Configure a DNS Address for the DHCP Server on page 873	Configure DNS servers for the DHCP Server globally.
3	Configure an NTP Address for the DHCP Server on page 874	Configure NTP servers for the DHCP Server globally.
4	Configure Custom Option Definitions for the DHCP Server on page 875	Optional. Define Custom Options to assign non-standard DHCP parameters.
5	Configure Custom Option Data for DHCP Server Globally on page 876	Optional. Assign global data values for any custom options you configure.
6	Configure Vendor Options or a Vendor Class	Optional. Configure a vendor class and vendor options.
7	Create a DHCP Server Subnet on page 877	Create a DHCP subnet and assign an IP address pool to the subnet.
8	Configure Subnet Routers for DHCP Server on page 878	Optional. Assign IP addresses to subnet routers.
9	Configure DNS for a DHCP Server Subnet on page 878	Optional. Assign DNS servers for subnets.
10	Configure an NTP Server for a DHCP Server Subnet on page 879	Optional. Assign NTP servers for subnets.
11	Configure Custom Option Data for a DHCP Server Subnet on page 880	Optional. Assign subnet-specific data values for any custom options you configure.
12	Configure a Static DHCP Server Lease on page 881	Optional. Configure hosts for whom you want to reserve details, such as a specific IP address.
13	Enable DHCP Server on page 882	Enable the DHCP Server on the switch.
14	For High Availability, repeat the DHCP Server configuration on the peer switch.	

Table 96: Vendor options and vendor class EDM configuration task flow

	Procedure	Description
1	Define Vendor Option Definitions on page 882	Define DHCP Server vendor options definitions.
2	Configure Vendor Option Data Globally on page 883	Configure DHCP vendor option data global parameters.

Table 96: Vendor options and vendor class EDM configuration task flow (continued)

	Procedure	Description
3	Configure a Vendor Class on page 884	Configure a vendor class.
4	Configure Custom Option Data for a Vendor Class on page 885	Configure custom options data for the vendor class.
5	Configure Vendor Option Data for a Vendor Class on page 886	Configure a value for the vendor option for the specified vendor class.

What to do Next

Configure DHCP Relay

Configure the DHCP Server Globally

Configure DHCP global parameters for the hosted DHCP Server. Configure only the parameters that apply to your deployment.

About This Task

It is not mandatory to configure all of the parameters at the global level. You can also apply many of these parameters to specific subnets or hosts.

To configure a TFTP Server List, see [Configure a Global TFTP Server List for the DHCP Server](#) on page 874.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Global** tab.
4. In **HA Role**, select the role to apply to this switch.
5. In **HA Peer Ip**, type the IPv4 address of the peer switch.
6. In **Lease Duration**, type the lease expiration timer in seconds.
7. In **TFTP Server Ip (opt 66)**, type the IPv4 address of the TFTP server to advertise on all subnets.
8. In **TFTP Server Name**, type the server name for the TFTP sever.
9. In **TFTP Bootfile Name**, type the boot image file name for the TFTP server.
10. Select **Authoritative** if you want to assign this server as an authoritative server.
11. In **NetBios Type**, select the option that matches the node type.
12. In **NetBios Server Ip**, type the IPv4 address of the NetBIOS server.
13. In **Domain Name**, type the domain name.
14. Select **Apply**.

Global Field Descriptions

Use the data in the following table to use the **Global** tab.

Name	Description
Enable	Specifies that the DHCP Server is enabled on the switch. By default, DHCP Server is disabled.
HA Role	Specifies the role of this switch in a DHCP Server High Availability (HA) pairing. In an HA pair, one server must be primary and the other must be secondary. Select one of the following options: <ul style="list-style-type: none"> • Default—The HA Role is assigned automatically. By default, the switch with the lower IP address is primary. • Primary—This switch provides a DHCP service while DHCP Server is active. • Secondary—This switch provides a DHCP service only when the primary switch is down.
HA Peer Ip	Specifies the IPv4 address of the peer switch in the HA pair.
Lease Time	Specifies the global lease duration timer in seconds. The default is 86400.
TFTP Server Ip (opt 66)	Specifies the IPv4 address of the TFTP server for the DHCP Server. This parameter uses DHCP Server Option 66. You must assign an IP address to the server before you can assign other TFTP settings such as a server name and boot image file.
TFTP Server Name	Specifies the server name of the TFTP server for the DHCP Server.
TFTP Bootfile Name	Specifies the file name of the TFTP boot image for the DHCP Server.
Authoritative	Assigns the DHCP Server as authoritative. The default setting is non-authoritative.
NetBios Type	Specifies the global NetBIOS node type. Select one of the following: <ul style="list-style-type: none"> • B-Node • P-Node • M-Node • H-Node
NetBios Server Ip	Specifies the IPv4 address of the global NetBIOS server.
Domain Name	Specifies the DHCP server domain name.

Configure a DNS Address for the DHCP Server

Use this procedure to configure a DNS address for the DHCP Server to advertise to clients.

Procedure

1. In the navigation pane, expand **Configuration > IP**.

2. Select **DHCP Server**.
3. Select the **Global DNS** tab.
4. Select **Insert**.
5. In **IP address**, type an IPv4 address for the DNS server.
6. Select **Insert**.

Global DNS Field Descriptions

Use the data in the following table to use the **Global DNS** tab.

Name	Description
IP Address	Specifies the IPv4 address of the DNS server. You can enter a maximum of eight DNS servers.
Origin	Specifies the origin of DHCP Server configuration, either manually configured through CLI or EDM, or dynamically configured through RADIUS.

Configure an NTP Address for the DHCP Server

Use this procedure to configure an NTP address for the DHCP Server to advertise to clients.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Global NTP** tab.
4. Select **Insert**.
5. In **IP Address**, type the IPv4 address of the NTP server.
6. Select **Insert**.

Global NTP Field Descriptions

Use the data in the following table to use the **Global NTP** tab.

Name	Description
IP Address	Specifies the IPv4 address of the NTP server. You can enter a maximum of eight DNS servers.
Origin	Specifies the origin of DHCP Server configuration, either manually configured through CLI or EDM, or dynamically configured through RADIUS.

NEW! *Configure a Global TFTP Server List for the DHCP Server*

About This Task

Configure a list of up to eight TFTP server IP addresses. Clients, such as IP Phones, can request Option 150 from a DHCP Server to obtain configuration files or other information from a TFTP server.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Global TFTP** tab.
4. Select **Insert**.
5. Type the IP address.
6. Select **Insert**.
7. Select **Apply**.

NEW! Global TFTP Field Descriptions

Use the data in the following table to use the **Global TFTP** tab.

Name	Description
Ip (opt 150)	Specifies a list of up to eight TFTP server IP addresses. This parameter uses DHCP Server Option 150, which clients, such as IP Phones, can request from a DHCP Server to obtain configuration files or other information from a TFTP server.

Configure Custom Option Definitions for the DHCP Server

Use this procedure to configure custom options for the DHCP Server to advertise to clients.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Custom Option Definitions** tab.
4. Select **Insert**.
5. In **ID**, type an integer from the unassigned range.
6. In **Type**, select the data type for this custom option.
7. In **Name**, type a text label that identifies this option.
8. Select **Insert**.



Note

You can use the **Custom Option Data** tab to assign a data value to the custom option that you define. You can assign the value to the DHCP Server globally, or within a DHCP subnet. See [Configure Custom Option Data for DHCP Server Globally](#) on page 876 and [Configure Custom Option Data for a DHCP Server Subnet](#) on page 880.

Custom Option Definitions Field Descriptions

Use the data in the following table to use the **Custom Option Definitions** tab.

Name	Description
ID	Specifies the DHCP code of the option to which you want to assign this data. Enter an integer in the unassigned range from 224 to 254.
Type	Assigns a data type to this Custom Option. All data values that use this custom option must follow the data type that you assign. Select one of three options: <ul style="list-style-type: none">• Ipv4 address—An IP address in IPv4 format• String—A text string.• Unit32—A 32-bit integer.
Name	Assigns a readable text label that identifies the custom option.

Configure Custom Option Data for DHCP Server Globally

Use this procedure to configure global custom option data for DHCP Server. Unless you create a subnet-specific configuration, subnets inherit the global configuration.

Before You Begin

You must first configure Custom Option Definitions. For more information, see [Configure Custom Option Definitions for the DHCP Server](#) on page 875.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Custom Option Data** tab.
4. Select **Insert**.
5. In **ID**, type the DHCP code that you defined for the option.
6. In **Value**, type the data value to assign to the option.
7. Select **Insert**.

Custom Option Data Field Descriptions

Use the data in the following table to use the **Custom Option Data** tab.

Name	Description
ID	Specifies the DHCP code that identifies this Custom Option. Must be an integer in the range from 224 to 254.
Value	Assigns a data value to this Custom Option. The value must use the data type format that is specified in the Custom Option definition.

Create a DHCP Server Subnet

About This Task

Use this procedure to create a DHCP Server subnet. The DHCP Server configuration must include at least one subnet.

To configure a TFTP Server List for the subnet, see [Configure a TFTP Server List for the DHCP Server Subnet](#) on page 879.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Subnet** tab.
4. Select **Insert**.
5. In **Subnet IP**, type the subnet network address.
6. In **Subnet Mask**, type the mask.
7. In **Pool Range Start**, type the IPv4 address that marks the start of the address range.
8. In **Pool Range End**, type the IPv4 address that marks the end of the address range.
9. (Optional) Configure other optional parameters for the subnet to override the global DHCP Server configuration.
10. Select **Insert**.

Subnet Field Descriptions

Use the data in the following table to configure the **Subnet** tab for DHCP Server.

Name	Description
Subnet IP	Assigns the network address of the DHCP subnet.
Subnet Mask	Assigns the subnet mask for the DHCP subnet.
Pool Range Start	Assigns the starting address in the IP address range.
Pool Range End	Assigns the ending IP address in the address range. This address must be higher than the starting address.
Lease Duration	Assigns the lease expiration timer in seconds.
TFTP Server IP (opt 66)	Specifies the IPv4 address of the TFTP server for the subnet. This parameter uses DHCP Server Option 66. You must assign an IP address to the server before you can assign other TFTP settings such as a server name and boot image file.
TFTP Server Name	Specifies the server name of the TFTP server.
TFTP Bootfile Name	Specifies the file name of the TFTP boot image.
Domain Name	Specifies the DHCP server domain name.

Name	Description
NetBios Type	Specifies node type of the NetBIOS server. . Select one of the following: <ul style="list-style-type: none"> • B-Node • P-Node • M-Node • H-Node
NetBios Server Ip	Specifies the IPv4 address of the NetBIOS server for the subnet.

Configure Subnet Routers for DHCP Server

Use this procedure to configure subnet routers (default gateways) to advertise to clients within a DHCP Server subnet.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Subnet Routers** tab.
4. Select **Insert**.
5. In **Subnet IP**, type the subnet network address.
6. In **Subnet Mask**, type the subnet mask.
7. In **IP**, type the router IPv4 address.
8. Select **Insert**.

Subnet Routers Field Descriptions

Use the data in the following table to use the **Subnet Routers** tab.

Name	Description
Subnet IP	Specifies the network address of the DHCP subnet.
Subnet Mask	Specifies the subnet mask for the DHCP subnet.
IP	Specifies the IPv4 address for the subnet router (default gateway). You can assign up to eight router addresses.

Configure DNS for a DHCP Server Subnet

Use this procedure to configure a DNS address to advertise to clients within a DHCP Server subnet.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Subnet DNS** tab.
4. Select **Insert**.
5. In **Subnet IP**, type the subnet network address.
6. In **Subnet Mask**, type the subnet mask.

7. In **IP**, type the IPv4 address of the DNS server.
8. Select **Insert**.

Subnet DNS Field Descriptions

Use the data in the following table to use the **Subnet DNS** tab.

Name	Description
Subnet IP	Specifies the network address of the DHCP subnet.
Subnet Mask	Specifies the subnet mask for the DHCP subnet.
IP	Specifies the IPv4 address for the DNS server that you want to assign to this subnet. You can assign up to eight DNS servers to a DHCP subnet.

Configure an NTP Server for a DHCP Server Subnet

Use this procedure to configure an NTP server address to advertise to clients within a DHCP Server subnet.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Subnet NTP** tab.
4. Select **Insert**.
5. In **Subnet IP**, type the subnet network address.
6. In **Subnet Mask**, type the subnet mask.
7. In **IP**, type the IPv4 address of the NTP server.
8. Select **Insert**.

Subnet NTP Field Descriptions

Use the data in the following table to use the **Subnet NTP** tab.

Name	Description
Subnet IP	Specifies the network address of the DHCP subnet.
Subnet Mask	Specifies the subnet mask for the DHCP subnet.
IP	Specifies the IPv4 address for the NTP server that you want to assign to this subnet. You can assign up to eight NTP servers to a subnet.

NEW! *Configure a TFTP Server List for the DHCP Server Subnet*

About This Task

Configure a list of up to eight TFTP server IP addresses. Clients, such as IP Phones, can request Option 150 from a DHCP Server to obtain configuration files or other information from a TFTP server.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Subnet TFTP** tab.
4. Select **Insert**.
5. Type the subnet IP address.
6. Type the subnet mask in bits.
7. Type the TFTP server IP address.
8. Select **Insert**.
9. Select **Apply**.

NEW! Subnet TFTP Field Descriptions

Use the data in the following table to use the **Subnet TFTP** tab.

Name	Description
Subnet Ip	Specifies the subnet IP address.
Subnet Bitmask	Specifies the subnet mask in bits.
Ip (opt 150)	Specifies a list of up to eight TFTP server IP addresses. This parameter uses DHCP Server Option 150, which clients, such as IP Phones, can request from a DHCP Server to obtain configuration files or other information from a TFTP server.

Configure Custom Option Data for a DHCP Server Subnet

Use this procedure to configure custom option data for a DHCP Server subnet. Unless you create a subnet-specific configuration, subnets inherit the global configuration.

Before You Begin

You must first configure Custom Option Definitions. For more information, see [Configure Custom Option Definitions for the DHCP Server](#) on page 875.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Subnet Custom Option Data** tab.
4. Select **Insert**.
5. In **Subnet IP**, type the subnet network address.
6. In **Subnet Mask**, type the subnet mask.
7. In **ID**, type the DHCP code for the custom option.
8. In **Value**, type the data value to assign to the custom option for this subnet.
9. Select **Insert**.

Subnet Custom Option Data Field Descriptions

Use the data in the following table to use the **Subnet Custom Option Data** tab.

Name	Description
Subnet IP	Specifies the network address of the DHCP subnet.
Subnet Mask	Specifies the subnet mask for the DHCP subnet.
ID	Specifies the DHCP code that identifies the Custom Option. The code must be an integer in the range from 224 to 254.
Value	Assigns a data value to this Custom Option for use within this subnet. The value must use the data type that you configured in the Custom Option definition.

Configure a Static DHCP Server Lease

Complete this procedure to reserve static DHCP options such as a specific IP address or TFTP boot file for a subnet host. Configure only the options that apply to your deployment.

About This Task

The host configuration overrides conflicting configuration from the subnet or global DHCP Server. For any host configuration that you leave unconfigured, the host inherits the configuration from the subnet.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Hosts** tab.
4. Select **Insert**.
5. Type the **Subnet IP** address and **Subnet Mask**.
6. Type the host **Mac Address**.
7. Configure the options to reserve for this host:
 - In **IP**, type the IPv4 address to reserve for this host.
 - In **Name**, type the hostname to assign to this host.
 - In **TFTP Server IP**, type the IPv4 address of the TFTP server that you want the host to use.
 - In **TFTP Server Name**, type the server name for the TFTP server that you want the host to use.
 - In **TFTP Bootfile Name**, type the file name for the boot image that you want the host to use.
8. Select **Insert**.

Hosts Field Descriptions

Use the data in the following table to use the **Hosts** tab.

Name	Description
Subnet IP	Specifies the network address of the subnet.
Subnet Mask	Specifies the subnet mask for the DHCP subnet.
Mac Address	Specifies the Mac address of the host.
IP	Assigns an IPv4 address to the host.
Name	Assigns a hostname to the host.
TFTP Server IP	Specifies the IP address of the TFTP server that you want this host to use.
TFTP Server Name	Specifies the server name of the TFTP server that you want this host to use.
TFTP Bootfile Name	Specifies the file name of the TFTP boot image that you want this host to use.

Enable DHCP Server

After you complete the DHCP Server configuration, enable the server.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Global** tab.
4. Select **Enabled**.
5. Select **Apply**.

Define Vendor Option Definitions

About This Task

Perform this procedure to define DHCP Server vendor options definitions.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options.

The switch process the DHCP options in the following order of precedence:

1. host—Host configuration applies only to a specific host.
2. DHCP subnet—DHCP subnet options apply to only those clients within the specified subnet.
3. vendor class—Vendor class options apply to only those clients with a string that matches a specified vendor class identifier (VCI).
4. global—Global options apply to all clients.

DHCP Server builds a cumulative configuration from all scopes but applies precedence rules in instances of conflict.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Global Vendor option Defines** tab.
4. Select **Insert**.
5. In **Id**, type the integer from the unassigned range.
6. In **Type**, select the data type for the vendor option.
7. In **Name**, type a text label that identifies this option.
8. Select **Insert**.

Global Vendor option Defines Field Descriptions

Use the data in the following table to use the **Global Vendor option Defines** tab.

Name	Description
Id	Specifies the DHCP code of the option to which you want to assign this data.
Type	Specifies a data type to this vendor option as follows: <ul style="list-style-type: none">• ipv4address—IPv4 address format• string—text string up to 128 characters• uint32—32-bit integer
Name	Specifies a text label that identifies the vendor option.

Configure Vendor Option Data Globally

Before You Begin

- Define the vendor option.

About This Task

Perform this procedure to configure DHCP vendor option data global parameters.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options.

The switch process the DHCP options in the following order of precedence:

1. host—Host configuration applies only to a specific host.
2. DHCP subnet—DHCP subnet options apply to only those clients within the specified subnet.
3. vendor class—Vendor class options apply to only those clients with a string that matches a specified vendor class identifier (VCI).
4. global—Global options apply to all clients.

DHCP Server builds a cumulative configuration from all scopes but applies precedence rules in instances of conflict.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Global Vendor option Data** tab.
4. Select **Insert**.
5. In **Id**, type the DHCP code that you defined for the option.
6. In **Value**, type the data value to assign to the option.
7. Select **Insert**.

Global Vendor option Data Field Descriptions

Use the data in the following table to use the **Global Vendor option Data** tab.

Name	Description
Id	Specifies the DHCP code that identifies this vendor option.
Value	Specifies a value to apply to this vendor option.

Configure a Vendor Class

About This Task

Perform this procedure to configure a vendor class. If the string you configure matches the vendor class identifier (VCI) provided by the DHCP client, the client is part of that vendor class.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options.

The switch process the DHCP options in the following order of precedence:

1. host—Host configuration applies only to a specific host.
2. DHCP subnet—DHCP subnet options apply to only those clients within the specified subnet.
3. vendor class—Vendor class options apply to only those clients with a string that matches a specified vendor class identifier (VCI).
4. global—Global options apply to all clients.

DHCP Server builds a cumulative configuration from all scopes but applies precedence rules in instances of conflict.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Vendor Class** tab.
4. Select **Insert**.
5. In **ID**, type the vendor class ID.
6. In **Matches**, type the vendor class identifier.

7. In **TFTP Server IP**, type the IPv4 address of the TFTP server that you want the vendor class to use.
8. In **TFTP Server Name**, type the TFTP server name that you want the vendor class to use.
9. In **TFTP Bootfile Name**, type the file name for the boot image that you want the vendor class to use.
10. In **Domain Name**, type the domain name that you want the vendor class to use.
11. Select **Insert**.

Vendor Class Field Descriptions

Use the data in the following table to use the **Vendor Class** tab.

Name	Description
ID	Specifies the vendor class.
Matches	Specifies the vendor class identifier for this vendor class.
TFTP Server IP	Specifies the IPv4 address of the TFTP server advertised on this vendor class.
TFTP Server Name	Specifies the server name of the TFTP server advertised on this vendor class.
TFTP Bootfile Name	Specifies the file name of the TFTP boot image advertised on this vendor class.
Domain Name	Specifies the domain name of the TFTP server advertised on this vendor class.

Configure Custom Option Data for a Vendor Class

Before You Begin

- Create the vendor class.
- Define the custom option.

About This Task

Perform this procedure to configure custom vendor option data for the vendor class.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options.

The switch process the DHCP options in the following order of precedence:

1. **host**—Host configuration applies only to a specific host.
2. **DHCP subnet**—DHCP subnet options apply to only those clients within the specified subnet.
3. **vendor class**—Vendor class options apply to only those clients with a string that matches a specified vendor class identifier (VCI).
4. **global**—Global options apply to all clients.

DHCP Server builds a cumulative configuration from all scopes but applies precedence rules in instances of conflict.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Vendor Class Custom Option** tab.
4. Select **Insert**.
5. In **ID**, type the Vendor Class ID.
6. In **Code**, type the DHCP code that you defined for the custom option.
7. In **Value**, type the value to assign to the custom option.
8. Select **Insert**.

Vendor Class Custom Option Field Descriptions

Use the data in the following table to use the **Vendor Class Custom Option** tab.

Name	Description
ID	Specifies the vendor class ID.
Code	Specifies the DHCP code that identifies the custom option.
Value	Specifies a data value to apply to this vendor class.

Configure Vendor Option Data for a Vendor Class

Before You Begin

- Define the vendor option.
- Create the vendor class.

About This Task

Perform this procedure to configure vendor option data for a vendor class.



Note

The switch follows inheritance and precedence rules for configured DHCP Server options.

The switch process the DHCP options in the following order of precedence:

1. **host**—Host configuration applies only to a specific host.
2. **DHCP subnet**—DHCP subnet options apply to only those clients within the specified subnet.
3. **vendor class**—Vendor class options apply to only those clients with a string that matches a specified vendor class identifier (VCI).
4. **global**—Global options apply to all clients.

DHCP Server builds a cumulative configuration from all scopes but applies precedence rules in instances of conflict.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Server**.
3. Select the **Vendor Class Vendor Option** tab.
4. Select **Insert**.
5. In **ID**, type the DHCP code that you defined for the vendor option.
6. In **Vendor class id**, type the value to apply to the vendor option.
7. In **Value**, type the value to assign to the vendor option.
8. Select **Insert**.

Vendor Class Vendor Option Field Descriptions

Use the data in the following table to use the **Vendor Class Vendor Option** tab.

Name	Description
ID	Specifies the vendor class ID.
Code	Specifies the DHCP code that identifies the vendor option.
Value	Specifies a data value to apply to this vendor class.

DHCP Relay Configuration

Complete the following tasks to configure DHCP Relay on the switch.

This section includes EDM procedures.

*Configure DHCP on a Brouter Port or a VRF Instance***Before You Begin**

- You must first enable BootP/DHCP relay on a port (or VLAN).
- You must enable DHCP and forwarding path.
- You must enable IP Routing on the interface.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Use the DHCP tab to configure the DHCP behavior on a brouter port or a VRF instance. The DHCP tab is available only if the port is routed (that is, assigned an IP address).

Procedure

1. In the **Device Physical View** tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **DHCP Relay** tab.
5. Select **Enable** to select the DHCP option.
6. Configure the other parameters as needed.
7. Select **Apply**.

DHCP Relay Field Descriptions

Use data from the following table in the **DHCP Relay** tab.

Name	Description
Enable	Lets you use BootP/DHCP on the port. The default is disable.
MaxHop	Sets the maximum number of hops before a BootP/DHCP packet is discarded (1 to 16). The default is 4.
MinSec	The secs field in the BootP/DHCP packet header represents the elapsed time since the client first sent the message. If the secs field in the packet header is greater than this value, the system relays or forwards the packet; otherwise, the packet is dropped. The default is 0 seconds.
Mode	Sets the interface to process only BootP, only DHCP, or both types of packets. The default is both.
AlwaysBroadcast	When enabled, the server reply is sent as a broadcast back to the end station. The default is disable.
CircuitId	Indicates whether DHCP Relay inserted the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.

Name	Description
Remoteld	Indicates whether DHCP Relay inserted the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Trusted	Indicates if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.

Configure BootP/DHCP on a VLAN or VRF Instance

Use the DHCP Relay tab to configure the DHCP behavior on a VLAN.

Before You Begin

- You must enable IP Routing on the interface.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **DHCP Relay** tab.
7. Select **Enable**.
8. (Optional) Configure the parameters as required.
9. Select **Apply**.

DHCP Relay Field Descriptions

Use the data in the following table to use the **DHCP Relay** tab.

Variable	Value
Enable	Enables DHCP Relay on the VLAN. The default is disable.
MaxHop	Specifies the maximum number of hops a BootP/DHCP packet can take from the DHCP client to the DHCP server. The default is 4.

Variable	Value
MinSec	Specifies the minimum number of seconds to wait between receiving a DHCP packet and forwarding the DHCP packet to the DHCP server. A value of 0 indicates that forwarding is done immediately. The default value is 0.
Mode	Specifies the type of DHCP packet required. The options are: <ul style="list-style-type: none"> • bootp • dhcp • both The default is both.
AlwaysBroadcast	When enabled, the DHCP Reply packets are sent as a broadcast to the DHCP client. The default is disable.
CircuitId	Specifies whether DHCP Relay inserted the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
RemoteId	Specifies whether DHCP Relay inserted the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Trusted	Specifies if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.
SmartRelayMode	Specifies one of the following DHCP Smart Relay modes: <ul style="list-style-type: none"> • disabled • parallel - the relay agent makes multiple copies of DHCP requests and inserts each IP address of the client-facing interface (VLAN or VRRP) as the gateway address in each copy of the request. The switch simultaneously relays multiple packets to the DHCP server, assigning each packet the correct IP address in the correct subnet. • sequential - the relay agent forwards received DHCP packets from all IP addresses on the VLAN or VRRP interface one at a time, starting with the primary address and continuing with each secondary address. The default is disabled.

Configure DHCP Relay

Before You Begin

If the agent IP address is a VLAN or port IP address, you must enable DHCP Relay on that VLAN or port within the VLAN context. However, if the agent IP address is a VRRP address, you do not need to enable DHCP Relay on the VLAN or port in which the VRRP address resides.

About This Task

After you configure the BOOTP/DHCP relay on an IP interface, you can configure forwarding paths to indicate where packets are forwarded. The forwarding paths are based on the type of packet and where the packet is received.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Relay**.
3. Select the **Globals** tab.
4. Select **Insert**.
5. In the **AgentAddr** box, type the agent address.
6. In the **ServerAddr** list, type the server address.
7. Select **Enable** to enable BOOTP/DHCP relay.
8. In the **Mode** box, select the type of messages to relay.

Both the mode setting for the DHCP interface and the mode setting for the agent interface determine which packets are forwarded.

9. Select **SrcPort67** to configure the source port for the BOOTP/DHCP relay request.
10. Select **Insert**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AgentAddr	Specifies the IP address of the input interface (agent) on which the BOOTP/DHCP request packets are received for forwarding. You cannot specify a Secondary IP Interface address as the agent IP address in the forwarding path. To make DHCP aware of Secondary IP Interfaces, you must enable DHCP Smart Relay.
ServerAddr	<p>This parameter is either the IP address of the BOOTP/DHCP server or the address of another local interface.</p> <ul style="list-style-type: none"> • If it is the address of the BOOTP/DHCP server, the request is unicast to the server address. • If the address is one of the IP addresses of an interface on the system, the BOOTP/DHCP requests are broadcast out of that local interface.
Enable	Enables BOOTP/DHCP relay.
Mode	<p>Specifies the type of messages relayed:</p> <ul style="list-style-type: none"> • Only BOOTP • Only DHCP • Both types of messages <p>The default is to forward both BOOTP and DHCP messages.</p>
SrcPort67	Assigns source port 67 to the UDP forwarding path. The default is source port 68.

Configure DHCP Relay on an IP Interface

Use this procedure to view or modify configuration information about the DHCP relay on IP interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Relay**.
3. Select the **Interfaces** tab.
4. To change the configuration information, select the field in the row for the required interface, and then type or select a new value.
5. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Variable	Value
IfIndex	Display a read-only interface number that represents a physical interface, or the VLAN logical interface.
MaxHop	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 4.
MinSec	Specifies the minimum number of seconds to wait between receiving a DHCP packet and forwarding the DHCP packet to the DHCP server. A value of 0 indicates that forwarding is done immediately. The default value is 0.
Mode	Specifies the type of DHCP packet required. The options are: <ul style="list-style-type: none">• bootp• dhcp• both The default is both.
AlwaysBroadcast	Specifies if DHCP Reply packets can be sent as a broadcast to the DHCP client. The default is false.
CircuitId	Specifies whether DHCP Relay inserts the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is false.
RemoteId	Specifies whether DHCP Relay inserts the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is false.

Variable	Value
Trusted	Specifies if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.
SmartRelayMode	<p>Enables DHCP Smart Relay on a VLAN interface in one of the following modes:</p> <ul style="list-style-type: none"> parallel—the relay agent makes multiple copies of DHCP requests and inserts each IP address of the client-facing interface (VLAN or VRRP) as the gateway address in each copy of the request. The switch simultaneously relays multiple packets to the DHCP server, assigning each packet the correct IP address in the correct subnet. sequential—the relay agent forwards received DHCP packets from all IP addresses on the VLAN or VRRP interface one at a time, starting with the primary address and continuing with each secondary address. <p>By default, Smart Relay is disabled.</p>

View DHCP Statistics for an Interface

View DHCP statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Relay**.
3. Select the **Interfaces Stats** tab.

Interfaces Stats Field Descriptions

Use the data in the following table to use the **Interfaces Stats** tab.

Name	Description
IfIndex	Identifies the physical interface.
AgentAddr	Shows the IP address configured as the relay on this interface. This address is either the IP of the physical interface or the IP of the VRRP address.
NumRequests	Shows the number of DHCP and BootP requests on this interface.
NumReplies	Shows the number of DHCP and BootP replies on this interface.

Graph DHCP Statistics for a Port

View DHCP statistics to manage network performance.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **DHCP** tab.
5. Select one or more values.
6. Select the type of graph to create.

DHCP Field Descriptions

The following table describes parameters on the **DHCP** tab.

Name	Description
NumRequests	The number of DHCP and/or BootP requests on this interface.
NumReplies	The number of DHCP and/or BootP replies on this interface.

View DHCP Statistics for a Port

View DHCP statistics to manage network performance.

Procedure

1. In the Device Physical view, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **DHCP Relay** tab.
5. Select **Graph**.
6. Select one or more values.
7. Select the type of graph.

DHCP Stats Field Descriptions

Use the data in the following table to use the **DHCP Stats** tab.

Name	Description
NumRequests	The number of DHCP and BootP requests on this interface.
NumReplies	The number of DHCP and BootP replies on this interface.

Display DHCP-Relay Statistics for Option 82

Display DHCP-relay statistics for all interfaces to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > IP**.

2. Select **DHCP-Relay**.
3. Select the **Option 82 Stats** tab.

Option 82 Stats Field Descriptions

Use the data in the following table to use the **Option 82 Stats** tab.

Name	Description
IfIndex	Shows the name of the interface on which you enabled option 82. Shows the port number if the interface is a brouter port or the VLAN number if the interface is a VLAN.
AgentAddr	Shows the IP address configured as the relay on this interface. This address is either the IP of the physical interface or the IP of the VRRP address.
FoundOp82	Shows the number of packets that the interface received that already had option82 in them.
Dropped	Shows the number of packets the interface dropped because of option 82-related issues. These reasons could be that the packet was received from an untrusted source or spoofing was detected. To determine the cause of the drop, you must enable trace on level 170.
CircuitId	Shows the value inserted in the packets as the circuit ID. The value is the index of the interface.
AddedCircuitId	Shows how many packets (requests from client to server) the circuit ID was inserted for that interface. If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
RemovedCircuitId	Shows how many packets (replies from server to client) the circuit id was removed for that interface.
RemotId	Shows the value inserted in the packets as the remote ID. The value is the MAC address of the interface.

Name	Description
AddedRemoteld	Shows how many packets (requests from client to server) the remote ID was inserted for that interface. If you expect this value to increase but it does not, and the interface does not drop a packet, it is possible the packet does not have enough space to insert the option. You must enable trace on level 170 to determine the cause.
RemovedRemoteld	Shows how many packets (replies from server to client) the remote ID was removed for that interface.

Graph DHCP Statistics for a VLAN

View DHCP statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. On the **Basic** tab, select a VLAN.
4. Select **IP**.
5. Select the **DHCP Relay** tab.
6. Select **Graph**.
7. Select one or more values.
8. Select the type of graph.

DHCP Stats Field Descriptions

Use the data in the following table to use the **DHCP Stats** tab.

Name	Description
NumRequests	The number of DHCP and BootP requests on this interface.
NumReplies	The number of DHCP and BootP replies on this interface.

UDP Broadcast Forwarding Configuration

By default, routers do not forward broadcasts. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts. You must setup UDP broadcast forwarding on the system. Configure UDP broadcast forwarding to forward the UDP broadcasts of network applications to the required server through physical or virtual router interfaces.

1. Enter protocols into a table.
2. Create policies (protocol/server pairs).
3. Assemble the policies into lists or profiles.
4. Apply the list to the appropriate interfaces.

Manage UDP Forwarding Protocols

About This Task

The switch configures the following protocols, by default:

- Time Service
- Terminal Access Controller Access Control System (TACACS) Service
- Domain Name System (DNS)
- Trivial file transfer protocol (TFTP)
- Network Basic Input/Output System (NetBIOS) NameSrv
- NetBIOS DataSrv

You can use these protocols to create forwarding entries and lists but you cannot delete them; you can add or remove other protocols to the list of protocols.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **UDP Forwarding**.
3. Select **Insert**.
4. In **PortNumber**, type a UDP port number.

This number defines the UDP port used by the server process as its contact port. The range is from 1 to 65535 and cannot be one of the UDP port numbers or a number previously assigned.

5. In **Name**, type a name for the protocol.
6. Select **Insert**.

The protocol is added to the Protocol table. After you create a protocol, you cannot change its name or number.

Protocols Field Descriptions

Use the data in the following table to use the **Protocols** tab.

Name	Description
PortNumber	Defines the UDP port.
Name	Specifies an administratively assigned name for this list.

Manage UDP Forwarding

About This Task

You manage UDP forwarding by defining the destination addresses for the UDP protocol.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **UDP Forwarding**.

3. Select the **Forwardings** tab.
4. Select **Insert**.
5. In the Insert Forwardings dialog box, select a destination UDP port from the defined protocols in the **DestPort** box.
6. Enter a destination IP address in the **DestAddr** box.
The destination address can be any IP server address for the protocol application or the IP address of an interface on the router.
7. Select **Insert**. The information is added to the Forwarding tab.

Forwardings Field Descriptions

Use the data in the following table to use the **Forwardings** tab.

Name	Description
DestPort	Specifies the port number defined for UDP, depending upon the protocol type.
DestAddr	Specifies the destination address can be any IP server address for the protocol application or the IP address of an interface on the router: <ul style="list-style-type: none"> • If the address is that of a server, the packet is sent as a unicast packet to this address. • If the address is that of an interface on the router, the frame is rebroadcast.
Id	Specifies an integer that identifies this entry internally.
NumFwdPackets	Specifies the total number of UDP broadcast packets forwarded using this policy.
NumDropPacketsTtlExpired	Specifies the total number of UDP broadcast packets dropped because the time-to-live value (TTL) expired.
NumDropPacketsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the specified destination address was unreachable.

Create the Forwarding Profile

About This Task

A forwarding profile is a collection of port and destination pairs. When you configure UDP forwarding list entries, be sure to first configure the UDP forwarding list. Then, configure your UDP forwarding list entries and assign them to a UDP forwarding list. If you do not assign a UDP forwarding list entry to at least one UDP forwarding list, the UDP forwarding list is lost after a restart.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **UDP Forwarding**.
3. Select the **Forwarding Lists** tab.

4. Select **Insert**.
5. In **Id**, type the forwarding list ID.
6. (Optional) In **Name**, type the name of the forwarding list, if required.
The system displays the forwarding list in **FwdIdList**.
7. Select **Insert**.

Forwarding Lists Field Descriptions

Use the data in the following table to use the **Forwarding Lists** tab and **Insert Forwarding Lists** dialog box.

Name	Description
Id	Specifies a value that uniquely identifies this list of entries (1 to 1000).
Name	Specifies an administratively assigned name for this list (0 to 15 characters).
FwdIdList	Specifies the zero or more port forwarding entries associated with this list. Each list identifier is stored as 2 bytes in this array, starting from 0 bytes (size=64). Select the ellipsis (...) to display the ID list.

Manage the Broadcast Interface

About This Task

Manage the broadcast interface by specifying and displaying which router interfaces can receive UDP broadcasts to forward.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **UDP Forwarding**.
3. Select the **Broadcast Interfaces** tab.
4. Select **Insert**.
5. In **LocalIfAddr**, select the ellipsis (...) to select a local interface IP address from the list.
6. Select **OK**.
7. In **UdpPortFwdListId**, select the ellipsis (...) to select a forwarding list ID from the list.
8. Select **OK**.
9. In **MaxTtl**, type the maximum number of hops an IP broadcast can take from the source device to the destination device (the default is 4; the range is 1 to 16).
10. In **BroadCastMask**, enter the subnet mask of the local interface that broadcasts the UDP broadcast packets.

When you configure the UDP forwarding broadcast mask, the broadcast mask must be less specific (shorter in length) or equally specific (equal in length) to the subnet mask of the IP interface on which it is configured. If the UDP forwarding broadcast mask is more specific than the subnet mask of the corresponding IP interface, UDP forwarding does not function properly.

11. Select **Insert**.

Broadcast Interfaces Field Descriptions

Use the data in the following table to use the **Broadcast Interfaces** tab.

Name	Description
LocalIfAddr	Specifies the IP address of the local router interface that receives forwarded UDP broadcast packets.
UdpPortFwdListId	Specifies the number of the UDP lists or profiles that this interface is configured to forward (0 to 100). A value of 0 indicates that the interface cannot forward any UDP broadcast packets.
MaxTtl	Specifies the maximum number of hops an IP broadcast packet can take from the source device to the destination device (the default is 4; the range is 1 to 16).
NumRxPkts	Specifies the total number of UDP broadcast packets received by this local interface.
NumFwdPkts	Specifies the total number of UDP broadcast packets forwarded by this local interface.
NumDropPktsMaxTtlExpired	Specifies the total number of UDP broadcast packets dropped because the time-to-live (TTL) value expired.
NumDropPktsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the destination was unreachable.
NumDropPktsUnknownPort	Specifies the total number of UDP broadcast packets dropped because the destination port or protocol specified has no matching forwarding policy.
BroadCastMask	Specifies the subnet mask of the local interface that broadcasts the UDP broadcast packets.

View UDP Endpoint Information

View UDP Endpoints to confirm correct configuration.

About This Task

You can use UDP endpoint information to display local and remote UDP activity.

Since UDP is a protocol used to establish connectionless network sessions, you need to monitor local and remote UDP activity and to know which applications are running over UDP.

You can determine which applications are active by checking the port number.

Processes are further identified with a UDP session to enable multiplexing of a port mapping for UDP.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **TCP/UDP**.
3. Select the **UDP Endpoints** tab.

UDP Endpoints Field Descriptions

Use the data in the following table to use the **UDP Endpoints** tab.

Name	Description
LocalAddressType	Displays the local address type (IPv6 or IPv4).
LocalAddress	Displays the local IPv6 address.
LocalPort	Displays the local port number.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IPv6 address.
RemotePort	Displays the remote port number.
Instance	Distinguishes between multiple processes connected to the UDP endpoint.
Process	Displays the ID for the UDP process.

IPv6 DHCP Relay Configuration using EDM

This section provides procedures to configure IPv6 DHCP relay using Enterprise Device Manager (EDM).

Configure a DHCP Relay Forwarding Path

Configure a forwarding path to specify the relay agent address and the DHCP server address to which to forward packets.

To use DHCP Relay for IPv6, you must configure at least one forwarding path and enable the relay on one interface.

Before You Begin

Change the VRF instance as required to configure a DHCP Relay forwarding path on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

The relay agent can use the IPv6 address of the interface or the VRRP global address linked to that interface. The relay forwards the DHCP messages only if VRRP is in the Master state, otherwise the relay discards the messages.

You can configure only one relay agent on an interface. If you need to change the relay agent, you must delete all the forwarding paths with the old relay agent, and then configure the new relay agent.

For scaling information on DHCP Relay forwarding paths, see *Fabric Engine Release Notes*.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **DHCP Relay**.
3. Select the **Forward Path** tab.
4. Select **Insert**.
5. In **AgentAddr**, type the address of the input interface that forwards the packets.
6. In **ServerAddr**, type the address of the DHCP server.
7. Select **Enabled**.
8. Select **Insert**.

Forward Path Field Descriptions

Use the data in the following table to use the **Forward Path** tab.

Name	Description
AgentAddr	Specifies the IP address of the input interface (relay agent) on which the DHCP request packets are received for forwarding. This address is the IPv6 or VRRP global address of either a brouter port or a VLAN for which forwarding is enabled.
ServerAddr	Specifies the IP address of the DHCP server. The request is unicast to the server address.
Enabled	Enables DHCP Relay for the system. The default is disabled (clear).

Configure DHCP Relay for an Interface

Configure the DHCP relay behavior on the interface.

Before You Begin

Change the VRF instance as required to configure DHCP Relay for an interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

You can modify the DHCP Relay configuration for a brouter port through the **Edit > Port > IPv6** navigation path, and for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can configure both types of interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **DHCP Relay**.
3. Select the **Interface** tab.
4. Select **Insert**.
5. Beside **IfIndex**, select **Port** or **Vlan**.
6. Select a port or VLAN.
7. Select **OK**.
8. Select **Insert**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface. For the brouter port, the value is the ifindex of the port and, in the case of the VLAN, the value is the ifindex of the VLAN.
MaxHop	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
RemoteIdEnabled	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled (clear or false).

Modify DHCP Relay for a VLAN

Modify the existing DHCP relay behavior on the VLAN interface.

About This Task

You can configure only one relay for a VLAN, regardless of how many addresses are configured on that VLAN. The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.

3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IPv6**.
6. Select the **DHCP Relay** tab.
7. Double-click a cell to change the value.
8. Select **Apply**.

DHCP Relay Field Descriptions

Use the data in the following table to use the **DHCP Relay** tab.

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface.
MaxHop	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
RemoteIdEnabled	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled (clear or false).
DhcpEnabled	Enables (true) or disables (false) DHCP Relay for an interface with an existing DHCP Relay configuration.

Modify DHCP Relay for a Port

Modify the existing DHCP relay behavior on the brouter port interface.

About This Task

The default address is the smallest address configured. If the relay is a VRRP address, the default value is the first VRRP address configured.

Procedure

1. In the Device Physical View, select the port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select **DHCP Relay**.
5. Double-click a cell to change the value.
6. Select **Apply**.

DHCP Relay Field Descriptions

Use the data in the following table to use the **DHCP Relay** tab.

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface. For the brouter port, the value is the ifindex of the port and, in the case of the VLAN, the value is the ifindex of the VLAN.
MaxHop	Specifies the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The default is 32.
RemoteIdEnabled	Enables the relay agent to add information about the relay to DHCPv6 messages before relaying the messages to the DHCP server. The default is disabled (clear or false).
DhcpEnabled	Enables (true) or disables (false) DHCP Relay for an interface with an existing DHCP Relay configuration. The system displays this parameter on the DHCP Relay tab for a brouter port only if you modify an existing configuration. The system does not display this parameter if you create a new DHCP Relay port configuration.

View DHCP Statistics for an IPv6 Interface

View IPv6 DHCP statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **DHCP Relay**.
3. Select the **Interfaces Stats** tab.

Interfaces Stats Field Descriptions

Use the data in the following table to use the **Interfaces Stats** tab.

Name	Description
IfIndex	Identifies the physical interface.
NumRequests	Shows the number of DHCP and BootP requests on this interface.
NumReplies	Shows the number of DHCP and BootP replies on this interface.

View IPv6 DHCP Relay Statistics for a Port

Display individual IPv6 DHCP Relay statistics for specific ports to manage network performance. You can also create a graph of selected statistical values.

Procedure

1. On the Device Physical view, select a port.
2. In the navigation pane, expand **Configuration > IPv6**.
3. Select the **DHCP Relay** tab.
4. Select the **Interface** tab.
5. Select the interface on which you want to view the IPv6 DHCP Relay statistics.
6. Select **Statistics**.
7. Select one or more values.
8. Select the type of graph.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.

Name	Description
NumRequests	Shows the number of DHCP and BootP requests on this interface.
NumReplies	Shows the number of DHCP and BootP replies on this interface.



Domain Name Service

[DNS fundamentals](#) on page 909

[DNS configuration using CLI](#) on page 910

[DNS configuration using EDM](#) on page 912

Table 97: Domain Name Service product support

Feature	Product	Release introduced
Domain Name Service (DNS) client (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
DNS client (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 97: Domain Name Service product support (continued)

Feature	Product	Release introduced
IPv6 Router Advertisement options for DNS	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.2
	5420 Series	Fabric Engine 9.2
	5520 Series	Fabric Engine 9.2
	5720 Series	Fabric Engine 9.2
	7520 Series	Fabric Engine 9.2
	7720 Series	Fabric Engine 9.2
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 9.2
	VSP 7400 Series	VOSS 9.2

The following sections provide information on the Domain Name Service (DNS) implementation for the switch.

DNS fundamentals

This section provides conceptual material on the Domain Name Service (DNS) implementation for the switch. Review this content before you make changes to the configurable DNS options.

DNS client

Every equipment interface connected to a Transmission Control Protocol over IP (TCP/IP) network is identified with a unique IPv4 or IPv6 address. You can assign a name to every machine that uses an IPv4 or IPv6 address. The TCP/IP does not require the usage of names, but these names make the task easier for network managers in the following ways:

- An IP client can contact a machine with its name, which is converted to an IP address, based on a mapping table. All applications that use this specific machine do not depend on the addressing scheme.
- It is easier to remember a name than a full IP address.

To establish the mapping between an IP name and an IPv4 or an IPv6 address you use the Domain Name Service (DNS). DNS is a hierarchical database that you can distribute on several servers for backup and load sharing. After you add a new hostname, update this database. The information is sent to all the different hosts. An IP client that resolves the mapping between the hostname and the IP address sends a request to one of the database servers to resolve the name.

After you establish the mapping of IP name and IP address, the application is modified to use a hostname instead of an IP address. The switch converts the hostname to an IP address.

If the entry to translate the hostname to IP address is not in the host file, the switch queries the configured DNS server for the mapping from hostname to IP address. You can configure connections for up to three different DNS servers—primary, secondary and tertiary. First the primary server is queried, and then the secondary, and finally the tertiary.

The DNS client tracks any server addresses or domain names provided from a DHCP server. If a DHCP server provides info to the DNS client, the DNS configuration is classified as dynamic. You can manually delete dynamic DNS entries, but cannot manually add dynamic DNS entries. You can view the Dynamic DNS entries with **show ip dns** or **show sys dns**. Dynamic DNS entries are not saved in the configuration file. The status monitoring of DNS occurs every 60 seconds.

DNS modifies Ping, Telnet, and copy applications. You can enter a hostname or an IP address to invoke Ping, Telnet, and copy applications.

A log/debug report is generated for all the DNS requests sent to DNS servers and all successful DNS responses received from the DNS servers.

IPv6 Support

The Domain Name Service (DNS) used by the switch supports both IPv4 and IPv6 addresses with no difference in functionality or configuration.

DNS configuration using CLI

This section describes how to configure the Domain Name Service (DNS) client using Command Line Interface (CLI).

DNS supports IPv4 and IPv6 addresses.

Configure the DNS Client

About This Task

Configure the Domain Name Service to establish the mapping between an IP name and an IPv4 or IPv6 address. DNS supports IPv4 and IPv6 addresses with no difference in

functionality or configuration using CLI.

You can configure connection for up to three different DNS servers—primary, secondary and tertiary. First the primary server is queried, and then the secondary, and finally the tertiary.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```

2. Configure the DNS client:
`ip domain-name WORD<0-255>`
3. (Optional) Add addresses for primary, secondary, or tertiary DNS servers:
`ip name-server <primary|secondary|tertiary> WORD<0-46>`
4. View the DNS client system status:
`show ip dns`

Example

Add addresses for a tertiary DNS server:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip name-server tertiary 254.104.201.141
```

Variable Definitions

The following table defines parameters for the **ip domain-name** command.

Variable	Value
<i>WORD<0-255></i>	Configures the default domain name. <i>WORD<0-255></i> is a string 0–255 characters.

The following table defines parameters for the **ip name-server** command.

Variable	Value
<i>primary secondary tertiary WORD<0-46></i>	Configures the primary, secondary, or tertiary DNS server address. Enter the IP address in a.b.c.d format for IPv4 or hexadecimal format (string length 0–46) for IPv6. You can specify the IP address for only one server at a time; you cannot specify all three servers in one command.

Query the DNS Host

About This Task

Query the DNS host for information about host addresses.

You can enter either a hostname, an IPv4 or IPv6 address. If you enter the hostname, this command shows the IP address that corresponds to the hostname and if you enter an IP address, this command shows the hostname for the IP address. DNS supports IPv4 and IPv6 addresses with no difference in functionality or configuration using CLI.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View the host information:
`show hosts WORD<0-256>`

Variable Definitions

The following table defines parameters for the **show hosts** command.

Variable	Value
<i>WORD</i> <0–256>	Specifies one of the following: <ul style="list-style-type: none">the name of the host DNS server as a string of 0–256 characters.the IP address of the host DNS server in a.b.c.d format.The IPv6 address of the host DNS server in hexadecimal format (string length 0–46).

DNS configuration using EDM

This section describes how to configure the Domain Name Service (DNS) using Enterprise Device Manager (EDM).

DNS supports IPv4 and IPv6 addresses with no difference in functionality or configuration except for the following. Under the **DNS Servers** tab, in the **DnsServerListAddressType** box, you must select **ipv4** or **ipv6**.

Configure the DNS Client

About This Task

You can configure connections for up to three different DNS servers—primary, secondary and tertiary. First the primary server is queried, and then the secondary, and finally the tertiary.

DNS supports IPv4 and IPv6 addresses. Under the **DNS Servers** tab, in the **DnsServerListAddressType** box, you must select **ipv4** or **ipv6**.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **DNS**.
3. Select the **DNS Servers** tab.
4. Select **Insert**.
5. In the **DnsServerListType** box, select the DNS server type.
6. In the **DnsServerListAddressType** box, select the IP version.
7. In the **DnsServerListAddress** box, enter the DNS server IP address.
8. Select **Insert**.

DNS Servers *Field Descriptions*

Use the data in the following table to use the **DNS Servers** tab.

Name	Description
DnsServerListType	Configures the DNS server as primary, secondary, or tertiary. OR Removes a DNS server as primary, primaryDynamic, secondary, secondaryDynamic, tertiary, or tertiaryDynamic.
DnsServerListAddressType	Configures the DNS server address type as IPv4 or IPv6.
DnsServerListAddress	Specifies the DNS server address.
DnsServerListStatus	Specifies the status of the DNS server.
DnsServerListRequestCount	Specifies the number of requests sent to the DNS server.
DnsServerListSuccessCount	Specifies the number of successful requests sent to the DNS server.

Query the DNS Host

About This Task

Query the DNS host for information about host addresses.

You can enter either a hostname or an IPv4 or IPv6 address. If you enter the hostname, this command shows the IP address that corresponds to the hostname and if you enter an IP address, this command shows the hostname for the IP address. DNS supports IPv4 addresses with no difference in functionality or configuration in this procedure.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **DNS**.
3. Select the **DNS Host** tab.
4. In the **HostData** text box, enter the DNS host name, IPv4 or the IPv6 address.
5. Select **Query**.

DNS Host *Field Descriptions*

Use the data in the following table to use the **DNS Host** tab.

Name	Description
HostData	Enter hostname or host IPv4 or IPv6 address to be identified.
HostName	Identifies the host name. This parameter is read-only.
HostAddressType	Identifies the address type of the host.

Name	Description
HostAddress	Identifies the host IP address. This parameter is read-only.
HostSource	Identifies the DNS server IP or host file. This parameter is read-only.

Configure the Default Domain Name

About This Task

Configure the default domain for querying the DNS server.

Procedure

1. In the navigation pane, expand **Configuration > IP > DNS**.
2. Select the **Globals** tab.
3. In **DNSDomainName**, type the domain name.
4. Select **Apply**.

Globals Field Description

Use the data in the following table to use the **Globals** tab.

Name	Description
DnsDomainName	Configures the default domain for querying the DNS server.
DnsDomainNameOrigin	Specifies the origin of the DNS domain name, either manually configured through CLI or dynamic through DHCP.
DNSAdvertisedHostName	Specifies the hostname the DHCP client advertises to the DHCP server.



Distributed Virtual Routing

[Distributed Virtual Routing Fundamentals](#) on page 919

[DvR configuration using the CLI](#) on page 940

[DvR Configuration Using EDM](#) on page 986



Important

The following tables separate Distributed Virtual Routing support based on the role of the switch, as a Controller node or as a Leaf node. You can configure some functionality on both node types while other functionality can exist only on one node type. For example, if you can configure functionality on both node types but the switch can only operate as a Leaf node, the Controller table indicates non support of the functionality while the Leaf table indicates support of the functionality.

Table 98: Distributed Virtual Routing product support on a Controller node

Feature	Product	Release introduced
DvR Controller role	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1.5 Not supported on VSP4900-48P
	VSP 7400 Series	VOSS 8.0

Table 98: Distributed Virtual Routing product support on a Controller node (continued)

Feature	Product	Release introduced
DvR Isolated Domain	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8 Not supported on VSP4900-48P
	VSP 7400 Series	VOSS 8.8
dvr-one-ip	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3 Not supported on VSP4900-48P
	VSP 7400 Series	VOSS 8.3

Table 98: Distributed Virtual Routing product support on a Controller node (continued)

Feature	Product	Release introduced
Distributed Virtual Routing (DvR) In-band Management	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1 Not supported on VSP4900-48P
	VSP 7400 Series	VOSS 8.0
DvR-VRRP Coexistence Important: To use this functionality, you must upgrade the software on all DvR BEBs, including Leaf nodes, to at least the minimum software version to support this functionality.	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8 Not supported on VSP4900-48P
	VSP 7400 Series	VOSS 8.8

Table 99: Distributed Virtual Routing product support on a Leaf node

Feature	Product	Release introduced
DvR Leaf role	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Supported on 5320-48P-8XE,

Table 99: Distributed Virtual Routing product support on a Leaf node (continued)

Feature	Product	Release introduced
		5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
DvR In-band Management	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 99: Distributed Virtual Routing product support on a Leaf node (continued)

Feature	Product	Release introduced
Management I-SID Assignment to DvR Leaf	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

Distributed Virtual Routing (DvR) is a technology for router redundancy in a Fabric deployment where IP subnets are stretched across multiple switches. DvR provides Default Gateway Redundancy and optimizes traffic flows to avoid traffic tromboning due to inefficient routing, thereby increasing the total routing throughput.

The topics in this section provide DvR concepts and configuration procedures.

Distributed Virtual Routing Fundamentals

You can deploy Distributed Virtual Routing (DvR) in Campus environments for stretching IP subnets between multiple aggregation layer switches and also simplifies data center deployments by introducing a Controller-Leaf architecture. In this architecture, Layer 3 configuration is required only on the Controller nodes, whereas the Leaf nodes require only Layer 2 configuration. All Layer 3 configuration is automatically distributed to the Leaf nodes by the Controller nodes.

For typical Campus DvR deployments, configure aggregation layer switches as DvR Controllers. Wiring closet access switches are then typically dual-homed to a pair of DvR Controllers.

IP subnets, which stretch between aggregation layer switches and multiple wiring closets, enable seamless IP roaming for wireless users while at the same time ensure optimal traffic forwarding. To optimize automation, Fabric Attach is typically deployed between wiring closet and aggregation switches. In this construct, there would likely be no DvR Leaf configured.

In Fabric deployments, DvR replaces VRRP (with VRRP-BackupMaster or RSMLT). The operator can choose for each I-SID/IP subnet what router redundancy method to use.

To migrate to a DvR-enabled I-SID/IP subnet, all member Fabric switches of this I-SID must be either DvR Controllers or DvR Leafs. You can connect non Fabric switches to DvR Leafs and DvR Controllers with manual configuration or Fabric Attach configuration. Until all Fabric switches that are members of the I-SID/IP subnet are DvR-enabled, use VRRP or RSMLT as the router redundancy protocol.

DvR Domain

DvR domains enable multi-site DvR deployments. Within a DvR domain, a set of up to eight DvR Controllers control the DvR domain Leaf switches. A domain can also include just DvR Controllers without DvR Leafs. Typically, a DvR domain is restricted to one physical location. Traffic leaving this physical location always passes through DvR Controllers.

A DvR domain is a logical group of switches or nodes that are DvR enabled. These nodes are not physically connected but are connected over the SPB Fabric such that each node is aware of the BMAC addresses of all other nodes within the domain. A DvR domain does not contain nodes that are not DvR enabled. However, those nodes can coexist with other DvR-enabled nodes within the same SPB Fabric network.

You configure a common DvR domain ID for all nodes belonging to a DvR domain. This domain ID translates internally to a Domain Data Distribution (DDD) I-SID. All switch nodes that share the same DvR domain ID or DDD I-SID receive the Layer 3 information that is distributed from all other nodes belonging to that DvR domain.

A DvR domain can contain multiple Layer 3 VSNs and Layer 2 VSNs. Layer 2 and Layer 3 VSNs can span multiple DvR domains.

A DvR domain typically has the following members:

1. DvR Controller(s)
2. DvR Leaf nodes

For more information about DvR Isolated Domains and Restrictions on a DvR Isolated Controller, see [DvR Isolated Domains](#) on page 920 and [Restrictions on a DvR Isolated Controller](#) on page 936.

For scaling information on the number of Controllers and Leaf nodes to configure in a DvR domain, see *Fabric Engine Release Notes*.

DvR Isolated Domains

A DvR isolated domain separates the DvR Controller from the DvR backbone in a Fabric Network. Configure an isolated Controller for the following benefits:

- to create DvR-enabled networks without backbone size limitations
- to prevent flooding from a backbone with many entries
- to improve the performance of less powerful devices as DvR Controllers

For more information about Configure an Isolated Controller on a DvR Isolated Domain, see [Configure an Isolated Controller on a DvR Isolated Domain](#) on page 941.

DvR Controller

In a DvR domain, the Controller nodes are the central nodes on which Layer 3 is configured. They own all the Layer 3 configuration and push the configuration information to the Leaf nodes within the SPB network.

A DvR domain can have one or more controllers for redundancy and you must configure every Layer 2 VSN (VLAN) and Layer 3 VSN within the domain, on the Controller(s). A node that you configure as a DvR Controller is considered the controller for all the Layer 2 and Layer 3 VSNs configured on that node. A Controller is configured with its own subnet IP address for every DvR enabled Layer 2 VSN within the domain.

All Layer 2 VSNs on a DvR Controller need not be DvR enabled. A controller can be configured with individual Layer 2 VSNs that are DvR disabled.

The Layer 3 configuration data that is pushed to the Leaf nodes include the Layer 3 IP subnet information for all Layer 2 VSNs within the DvR domain. It also includes the IP routes learned or redistributed by the Controllers from networks outside the SPB network, into the DvR Domain. Controllers also send information on whether Multicast is enabled on a specific DvR enabled Layer 2 VSN, and the version of IGMP. DvR Controllers inject a default route into the DvR domain for external route reachability. Use route policies to inject specific routes into a DvR domain or inject host routes into OSPF or BGP.



Note

When sFlow operates in a DvR domain and DvR Leaf nodes use the management CLIP address as the sFlow agent IP, DvR Leaf nodes always report the sFlow collector as reachable because DvR Controllers inject a default route into the DvR domain. You can use the **dvr controller inject-default-route-disable** command to withdraw the route and force DvR Leaf nodes to use either a DvR host route, a direct, or a static route that the DvR Controller can redistribute. The best practice is to perform appropriate analysis before you use this setting.

A Controller can only belong to one DvR domain, based on the domain ID that you configure on the node.

DvR Controllers include all DvR Leaf functions, thus a Leaf node free deployment is a valid network deployment. Especially if you use DvR in Campus deployments to replace VRRP or RSMLT, a Controller-only deployment, as Fabric Attach server nodes, is a valid deployment option. For more information about DvR Isolated Domains and Restrictions on a DvR Isolated Controller, see [DvR Isolated Domains](#) on page 920 and [Restrictions on a DvR Isolated Controller](#) on page 936.

DvR Leaf Node

DvR Leaf nodes are typically data center top of the rack (TOR) Fabric switches that aggregate physical and virtual servers or storage devices. DvR Leaf nodes operate in a reduced configuration mode, where Layer 3 is not configured locally, but pushed to

them from the DvR Controller(s) within the domain. You need to configure only the IS-IS infrastructure and the Layer 2 VSNs on the Leaf nodes.

A DvR Leaf node also monitors local host attachments and communicates updates about the current state of those host attachments to the DvR domain. All DvR nodes exchange host attachment information using the DvR host distribution protocol, which leverages a DvR domain I-SID.

DvR leaf nodes are managed in-band through a local loopback address, which is exchanged using the IP Shortcut protocol.

Eligibility Criteria for a Leaf Node

A Leaf node must support the following criteria:

- configuration of basic parameters of IP Multicast over Fabric Connect, such as the system ID, nickname, B-VLANs, SPBM instance, area, peer system ID and virtual BMAC
- configuration of a physical port as either an SPB network-to-network interface (NNI), a FLEX-UNI interface or an FA interface
- configuration of an MLT as either an SPB NNI, a FLEX-UNI interface or an FA interface
- configuration of an SMLT as a Flex-UNI Interface or an FA Interface
- configuration of Layer 2 VSN I-SID instances of type ELAN
- configuration of FLEX-UNI end-points as part of a Layer 2 VSN
- FA Server functionality on FA enabled interfaces
- SMLT and vIST
- configuration of a in-band management interface for in-band management of the node

Summary of Controller and Leaf Node Functions

A DvR Controller performs the following functions:

- pushes Layer 3 configuration data (IPv4 Unicast and Multicast) to the Leaf nodes for all the Layer 2 VSNs or subnets within the DvR domain.
- pushes the Layer 3 learned host routes (host routes learned on its own UNI ports) and route data learned through route redistribution or route policies, to the Leaf nodes.
- configures learned remote host routes from other Controllers and Leaf nodes, on its own device.

A DvR enabled Leaf node performs the following functions:

- configures the gateway MAC when the gateway IPv4 address is learned.
- pushes the Layer 3 learned remote host routes to other Controllers and Leaf nodes in the domain.
- configures learned remote host routes from other Controllers and Leaf nodes on its own device.

- configures ECMP routes (in the datapath only) for the Layer 2 VSN subnets, with each next hop as the Controller in the DvR domain.
- configures learned routes from the Controllers that are redistributed using DvR.
- handles host route response packet interception based on the Controller VLAN MAC or the gateway MAC.

DvR backbone

The DvR backbone is automatically established among the DvR Controllers from all DvR domains. Every Controller node has an edge gateway to its DvR domain, to the DvR backbone and all other non-DvR domains within the network.

Controllers exchange host route information such that any host can be reached in a shortcut switched manner, irrespective of its location. For these host route information exchanges, controllers use an automatically assigned backbone I-SID. Local subnets to the Controllers are automatically injected into the DvR host route exchanges.

To redistribute DvR host routes into OSPF or BGP, you can configure route policies. These host routes are not injected into IS-IS.

DvR Backbone Members

You can configure a non-DvR backbone edge bridge (BEB) to join the DvR backbone. This enables the node to receive redistributed DvR host routes from all DvR Controllers in the SPB network, just like a DvR Controller. However, unlike the Controller, you can neither configure a DvR interface on this node nor can the node inject its host routes into the DvR domain.

DvR operation

In a DvR domain, DvR enabled Controller(s) handle the learning and distribution of Layer 3 configuration and route data to the DvR enabled Leaf nodes. The Leaf nodes in turn, use this data to automatically create distributed Layer 3 datapaths on themselves. In this way, Layer 3 configuration and learning remains only with the Controller(s) and there are distributed Layer 3 datapaths at the edges of the fabric network. This allows for destination lookups at the edge to happen quickly, and traffic is sent directly to their destinations without multiple lookups.

An important benefit of DvR is that only minimal configuration is required on the Leaf node. Based on the Layer 2 VSN that the Leaf node is a part of, all Layer 3 configuration information (IPv4 Unicast and Multicast configuration) is pushed from the Controllers in the domain. Thus the leaf nodes, although basically Layer 2 configured switches, become fully layer 3 capable devices.




Figure 49: SPB Fabric network with central Layer 3 Controller and distributed Layer 3 datapath at the edges

ARP Learning

When you enable Distributed Virtual Routing (DvR) on a Controller, it initiates Address Resolution Protocol (ARP) requests for traffic to be routed to unknown destination hosts.

A DvR-enabled Controller learns ARP requests from:

- DvR-enabled Leaf nodes (here the Leaf node owns the ARP requests)
- its own local UNI ports. Here, the controller owns the ARP requests
- other DvR-enabled Controllers

A DvR-enabled Leaf node learns ARP requests from:

- its own local UNI ports (here the Leaf node owns the ARP requests)
- other DvR-enabled Leaf nodes (that own the ARP requests) and respond to ARP requests on their UNI ports
- DvR-enabled Controllers (that own the local UNI ARP requests)

A Controller only distributes ARP entries that are locally learned on its own UNI ports to other DvR-enabled nodes in the domain.

dvr-leaf-mode boot flag

To configure a node to operate as a DvR Leaf node, you must first enable the `dvr-leaf-mode boot` flag.

- The `dvr-leaf-mode` boot flag is disabled by default. You must explicitly enable this flag before you configure a switch node to operate as a Leaf node.

When you enable the `dvr-leaf-mode` boot flag, you can configure the node as a DvR leaf node without rebooting, as long as there is no unsupported configuration discovered on the switch.

- After you enable or disable the boot flag, you must save the configuration.



Important

A node on which the `dvr-leaf-mode` boot flag is enabled cannot be configured as a DvR Controller.

In-band management

Use in-band management to manage a DvR enabled Leaf node that does not have an out-of-band management port or a console port.

For in-band management of the node within the management subnet (for example, from a Controller node), you must configure a unique IPv4 address to be used as the in-band management IP address, on that node. This IPv4 address functions like a CLIP address.

DvR deployment scenarios

The following sections describe typical deployments of the DvR infrastructure.


DvR Deployment in a Single Data Center

The following topology shows DvR deployment in a single data center. This deployment consists of a single DvR domain comprising a Controller layer and a Leaf node layer. The Controller layer has two controllers (for redundancy), which are deployed closer to the boundary of the DvR domain and the rest of the SPB Fabric network. The DvR Leaf nodes or Top of Rack (TOR) switches are typically access or edge switches.

All switches that belong to the DvR domain are configured with the same DvR domain ID and communicate with each other over a predefined I-SID.


The Controller nodes control the Leaf nodes and also build the gateway between the DvR domain and the rest of the Fabric infrastructure. So traffic is either routed between the Leaf nodes, or through the Controllers, to the rest of the fabric infrastructure.

Two IP subnets (Layer 2 VSNs), yellow and green, span the Leaf nodes. Each subnet is configured with a virtual IP address that is shared among all Controller and Leaf nodes that belong to the subnet. The Controller and Leaf nodes are configured with routing interfaces to the subnets, as shown in the figure.



DvR works by enabling each Leaf node or Top of Rack (TOR) switch to bi-directionally route traffic for each IP subnet of which it is a member. This is done by distributing the Layer 3 configuration information (IP Unicast, IP Multicast and virtual IP configuration) needed to handle Layer 3 routing, from the Controllers to the Leaf nodes. Configuration information is pushed over the DvR Domain I-SID, as indicated by the blue arrows in the previous figure.

Routing between the two IP subnets is achieved directly at the Leaf nodes when the Layer 3 distributed datapath is programmed at the Leaf Nodes, based on the Layer 3 configuration data that is pushed. Thus traffic within and between IP subnets is shortcut switched without having to traverse the central routing nodes, as shown in the following figure, if there are direct physical connections between them.



Thus, in a DvR deployment, all virtual IP and Layer 3 configuration is performed on the Controller nodes and pushed to the Leaf nodes, so that the Leaf nodes though basically Layer 2 configured switches, become fully layer 3 capable devices.

DvR deployment in a dual data center

The following example deployment shows two data centers each having its own DvR domain, connected through a backbone.

All nodes in data center Campus 1 belong to DvR domain shown in green, and the nodes in the data center Campus 3 belong to the DvR domain shown in orange. The two DvR domains are individually managed, so in this scenario, the controllers colored orange manage the orange Leaf nodes and the controllers colored green manage the green Leaf nodes. However, subnets can still be stretched across the DvR domains (and possibly between buildings), as shown in the figure.

Each DvR domain learns its own Layer 3 data and distributes this information to its own Leaf nodes. Layer 3 host information that is redistributed from other DvR Domains is learned by the Controllers only (through inter-DvR domain redistribution) and is programmed on the Leaf nodes in the same domain, but not in the other Domain. For example, Layer 3 information redistributed from domain 2 is learned by all controllers including the domain 1 controllers, but this information is not distributed to the Leaf nodes in domain 1. Hosts in one DvR domain can reach the hosts in the other DvR domain only through the Controllers.




Figure 50: Shortest path routing between servers in different data centers

All controllers in all domains are always part of the DvR backbone by default, as they are connected by the SPB Fabric. The DvR backbone connects many DvR domains.

Thus DvR can scale to multiple campuses, allowing a simplified way to deploy a large scale fully-routed infrastructure.

DvR Route Redistribution

The following sections describe redistribution of IPv4 local and static routes from DvR Controllers into the DvR domain. It also describes redistribution of host routes that are learned on DvR enabled VLANs, to BGP and OSPF. You can configure route policies to control the selection of routes to be distributed. You can also configure IS-IS accept policies on DvR Controllers and non-DvR BEBs, to determine which DvR host routes to accept into the routing-table from the DvR backbone.

Redistribution of IPv4 Local and Static Routes

The DvR feature supports redistribution of IPv4 local and static routes into the DvR domain.

**Note**

For every VRF instance and the Global Router, the Controller automatically injects a default route to the Leaf node, with a next hop as the advertising Controller. However, if you require only local or static routes to be advertised to the Leaf nodes, you can manually disable the injection of default routes on the Controller.

On a DvR Controller, you can configure (enable or disable) the redistribution of direct or static routes. Direct routes are redistributed with the route type as internal. Static routes are redistributed with the route type as external. You can apply route policies on the Controller to selectively permit the redistribution of these routes and also configure a metric value for the route that is redistributed. The default metric for imported local routes is 1. For static routes, the configured route metric or cost is honored.

You can configure redistribution of static and direct routes from the Global Router, or within a VRF instance. For redistributed routes, the Controller configures the Layer 3 VSN as that of the VRF redistributing the route, and the next hop BEB as the system ID of the Controller injecting the route into the DvR domain.

The following example demonstrates how a DvR Leaf node benefits from the redistribution of local and static routes.

By default, if the injection of default routes is enabled on a DvR Controller, the DvR Leaf node can only route traffic to other nodes within the DvR enabled subnet. For the Leaf node to reach networks outside of the DvR enabled subnet, the Controllers must redistribute local and static routes from non-DvR subnets into the DvR domain. In the following figure, the DvR Leaf L1 can route traffic only to nodes in the DvR enabled subnet 10.10.10.0/24. To be able to reach hosts in VLAN 20 (20.20.20.0/24) or VLAN 30 (30.30.30.0/24), redistribution of local routes into DvR is required at each of the Controllers C1 and C2. For the Leaf node to reach hosts in remote networks 40.40.40.0/24 or 50.50.50.0/24, redistribution of static routes to the DvR domain is required.

You can apply route policies to control which local or static routes are to be redistributed into the DvR domain.




Figure 51: Redistribution of IPv4 local and static routes

Redistribution of Routes to OSPF or BGP

For non-SPB routers to benefit from the host accessibility information learned within a DvR domain, DvR supports the redistribution of host routes into OSPF or BGP. Redistribution of these host routes is only by the DvR Controllers and only for the intra-domain host routes within the DvR enabled subnets.

A DvR Controller can redistribute host routes for all hosts from a DvR domain into OSPF or BGP. You can also apply route policies on the Controller to select the routes to be redistributed. The Controller supports redistribution of routes from the Global Router or within a VRF instance. You can also configure the metric of the route before redistribution.

The following example demonstrates the benefit of redistribution of routes to BGP.

Consider a 10.1.0.0/16 network with a stretched Layer 2 VSN spanning two data centers. On the campus side of the network, BGP peering is configured between a non-Extreme Networks router and one or more routers in the data center. BGP advertises the network route 10.1.0.0/16 to the campus BGP routers. Depending on which edge router the traffic is delivered to, it is possible that traffic from a host on the campus traverses the WAN a second time to reach the server that is physically connected to one segment of the data center, as shown in the following figure.




Figure 52: Inefficient traffic flow

Redistribution of the host routes from the DvR Controller to BGP solves this problem.

The following figure shows two DvR domains (shown in green and orange) configured at each data center. Each campus edge router establishes a BGP peering session with one or more Controllers in each data center (DvR domain). This enables BGP to advertise more specific routes to the campus BGP router so that the optimal routing path is always taken. So, there is no need for traffic to traverse the WAN multiple times. Also, in the case of server movement within or between data centers, the updated DvR host routes are propagated to BGP, thus ensuring that traffic flowing into the data center continues along the most optimal path.

For example, in the following figure, only the Controller attached to the Leaf node where the 10.1.0.111 server exists, advertises its accessibility over the 10.1.0.111/32 route. Similarly, the DvR Controller associated with the Leaf node connected to the 10.1.0.222 server advertises the 10.1.0.222/32 host route.




Figure 53: Traffic flow optimized with route redistribution

Controllers in each data center learn all host routes through the DvR backbone, but since those routes belong to different DvR domains, they are not all eligible for redistribution to OSPF or BGP.

Route Redistribution and IS-IS Accept Policies

DvR route redistribution leverages IS-IS accept policies to control (accept or reject) DvR routes learned from the DvR backbone. You can configure accept policies on both Controllers and non-DvR BEBs in the SPB network.

For more information about accept policies, see [IS-IS Accept Policies](#) on page 1597.

DvR Deployment for Wireless Roaming in Campus Deployments

In fabric deployments where IP subnets/I-SIDs stretch between multiple buildings, you can use DvR instead of VRRP or RSMLT to avoid traffic tromboning issues. This deployment is supported with non-fabric switches that have Fabric Attach enabled, or switches that do not support Fabric Attach. IP subnets can stretch across one or multiple DvR domains as shown in the following figure.

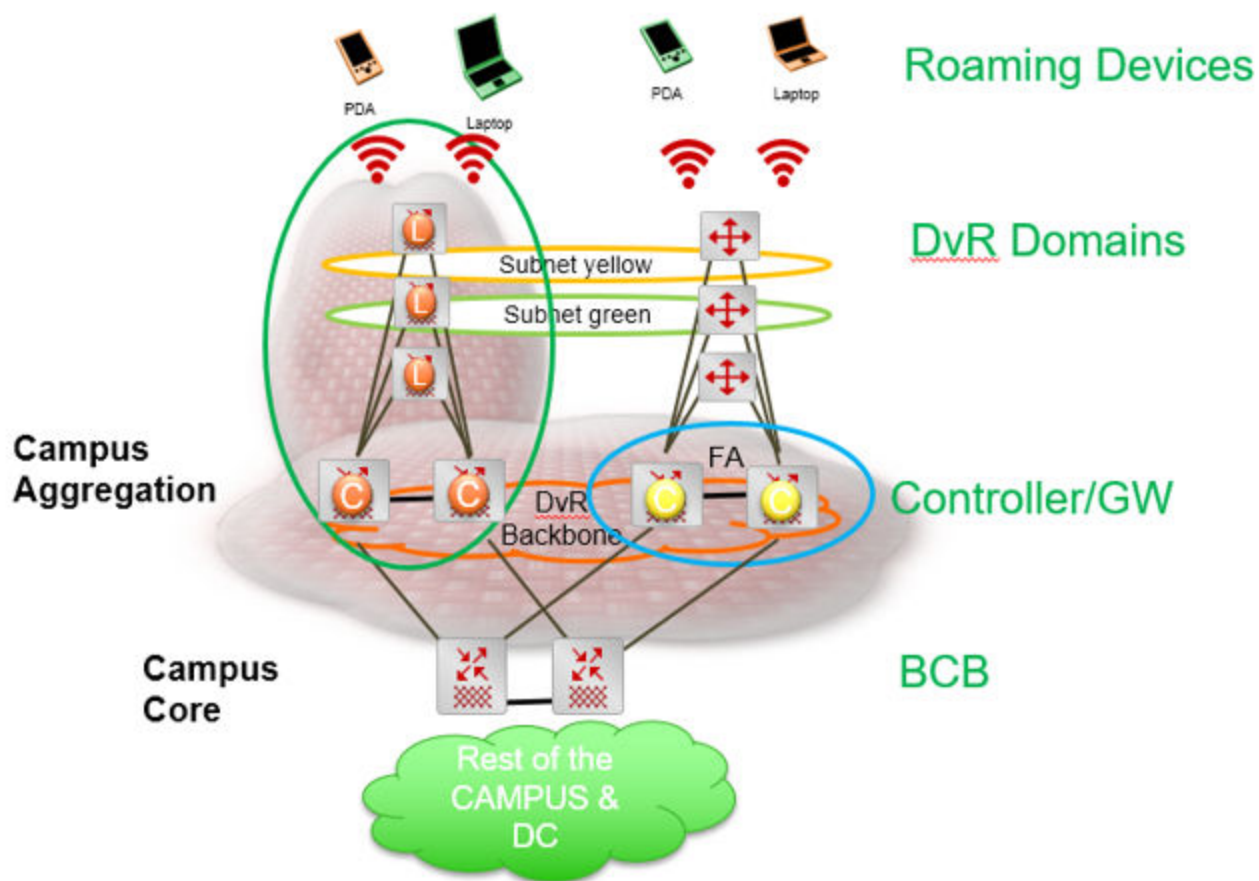


Figure 54: Wireless roaming in Campus

Management I-SID Assignment to DvR Leaf

The Management I-SID Assignment to DvR Leaf feature simplifies the process of creating a Management Instance VLAN interface on a DvR Leaf.

You can configure a Management Instance VLAN on a DvR Leaf node by specifying the I-SID. When you specify the I-SID, an internal VLAN is created and associated with the I-SID.

Operational Considerations

The following section describes operational considerations for assigning a Management I-SID to DvR Leaf node.

- If the specified I-SID is not associated with a VLAN, the Management Instance VLAN interface is created with the specified I-SID and an internal VLAN.



Note

The internal VLAN is not configurable.

- You cannot create a Management Instance VLAN interface if the I-SID is already associated with a VLAN. However, you can configure a Management Instance VLAN interface if the I-SID is associated with the onboarding VLAN.
- The I-SID cannot be learned dynamically. You cannot create a Management Instance VLAN interface if the I-SID sent from the DvR Controller is the same as the Management Instance VLAN I-SID.
- When you delete the Management Instance VLAN interface, the internal VLAN is deleted. You cannot delete the onboarding VLAN.
- You can migrate from a previous configuration or to a DvR Leaf only if an I-SID is associated with the Management Instance VLAN. If you disable DvR on a DvR Leaf by disabling the **dvr-leaf-mode** boot flag, the Management Instance VLAN is deleted. You can either restart the onboarding process or configure a Management Instance CLIP or Management Instance OOB interface.
- In DvR Leaf, you can create a Management Instance VLAN interface using the quick-config-mgmt utility script with the Management I-SID but not with port number or VLAN ID. This I-SID cannot be used if the Management Instance interface is already created. You can issue the **convert** command to use this I-SID.
- Migration from DvR leaf to non-DvR leaf deletes the Management Instance VLAN configuration.

Using Ping or IP Traceroute for Hosts in the DvR-One-IP Subnet

To use DvR-One-IP, a circuitless IP (CLIP) must exist in the VRF to which the DvR-One-IP interface belongs. If the DvR-One-IP interface is part of the global router (GRT), a CLIP must exist in the GRT and it must be configured as the IS-IS **ip-source-address**. If these CLIPs exist, pinging hosts in the DvR-One-IP subnet from the DvR Controller works as expected.

If the CLIPs do not exist, pinging hosts in the DvR-One-IP subnet is not possible from DvR Controllers. The ping attempt times out and the switch displays the following warning message: `Warning: For DVR one IP a loopback IP must be configured on the VRF.` If you provide a source IP address with the **ping** command, the switch does not display the warning message but the ping attempt fails.

This same restriction also applies to IP traceroute.

DvR Restrictions

Review the following restrictions and behavioral characteristics associated with DvR.

- The DvR feature does not affect out-of-band management on a switch chassis, if the chassis supports it.
- The DvR feature does not support a non-DvR BEB in a DvR-enabled Layer 2 VSN.



Note

If the DvR Controller supports it, you can configure Virtual Router Redundancy Protocol (VRRP) on DvR I-SIDs to permit traffic routing from non-DvR Backbone Edge Bridges (BEBs) outside of the domain. For more information, see [DvR-VRRP Coexistence](#) on page 938. For information on DvR Controller support, see *Fabric Engine and VOSS Feature Matrix*.

- The number of host route records that can be stored in the data path of a Leaf node is limited to the scaling capacity of the switch node. Different switch platforms have different scaling capacities.

For information on the scaling capacities of different platforms, see *Fabric Engine Release Notes*.

- You must first disable DvR on a Controller or Leaf node, before you attempt to change the domain ID of the node.
- You cannot configure IGMP snooping on DvR-enabled nodes.
- In any DvR ISID the only supported VRRP routers are DvR Controllers with DvR VRRP enabled. Interoperation with non-DvR VRRP routers is not supported.

Configuration Restrictions on a DvR Controller

- If you use two different IP addresses for the DvR VLAN and the DvR GW IP, you must first configure a gateway IPv4 address, and then configure an IP interface for the VLAN before you enable DvR on a Layer 2 VSN (VLAN). Both the VLAN IP address and the gateway IPv4 address must be in the same subnet.

If you use the same IP address for VLAN interface and DvR GW IP, you can use the command **ip address {A.B.C.D/X} dvr-one-ip**.

For more information, see [Enable DvR on a Layer 2 VSN \(VLAN\)](#) on page 4040 and [Configure a Single IP Address for All DvR Controllers on a VLAN Subnet](#) on page 4042.

- You cannot configure IPv4 VRRP on a DvR-enabled VLAN.



Note

If the DvR Controller supports it, you can configure Virtual Router Redundancy Protocol (VRRP) on DvR I-SIDs to permit traffic routing from non-DvR Backbone Edge Bridges (BEBs) outside of the domain. For more information, see [DvR-VRRP Coexistence](#) on page 938. For information on DvR Controller support, see *Fabric Engine and VOSS Feature Matrix*.

- You cannot configure RSMLT on a DvR-enabled VLAN.
- You cannot configure SPB-PIM Gateway (SPB-PIM GW) on a DvR VLAN.

- You cannot configure dynamic routing protocols, such as OSPF, RIP, BGP, IPv6 OSPFv3, IPv6 RIPng, IPv6 MLD, and IPv6 PIM-GW on a DvR-enabled VLAN.
- DvR-enabled VLAN/I-SIDs are for host connectivity only; you cannot connect a router to a DvR-enabled VLAN/I-SID and use dynamic or static routing. Use a non-DvR VLAN/I-SID instead to connect an external router.

The following additional restrictions apply to platforms that support IPv6 VRRP coexistence with DvR on the same I-SID:

- You can configure DvR on a VLAN that has configured IPv6 interface. You must first delete the IPv6 interface, configure DvR, and then reconfigure IPv6 interface.

A DvR VLAN is a VLAN configured on a DvR Controller with a VLAN IP address, a VLAN/I-SID, the DvR gateway IP address, and DvR enabled. This Layer 3 configuration for the DvR VLAN (the DvR gateway IP address and this DvR subnet) is pushed to the DvR Leaf nodes. The DvR gateway IP address must be the same address across all DvR Controllers for that DvR VLAN.

- You cannot configure an IPv6 interface on a DvR-enabled VLAN from a subnet that is used as next-hop in an IPv6 static route.
- You cannot configure an IPv6 address on a DvR-enabled VLAN from a subnet used as an IPv6 BGP Peer.

Configuration Limitations on a DvR Leaf

- Enabling the **DvR-leaf-mode** boot flag before you configure a node as a DvR Leaf, automatically removes all existing non-DvR configuration on the node such as platform VLANs and their IP address configuration, CLIP configuration, routing protocol configuration and VRF configuration. The gateway IPv4 address, if configured, is also removed.
- You cannot configure SPB-PIM GW on a Leaf node. The configuration is supported only on a DvR Controller.
- You cannot perform Layer 3 configuration (for example, IP interfaces, IP routing, and VRFs). You can only perform Layer 2 configuration.

You cannot configure Microsoft NLB on a Leaf node.

- You cannot configure Fabric Extend on a Leaf node.
- You cannot configure the VXLAN Gateway on a Leaf node.
- You cannot configure a T-UNI on a Leaf node.
- You cannot configure IPv4 multicast on a Leaf node. The configuration is supported only on a DvR Controller.
- You can configure only one instance of vIST on a Leaf node pair. Also, you cannot configure vIST on Leaf nodes from different domains.
- Platform VLANs are not supported. You cannot configure a platform VLAN directly on a DvR leaf node. However, you can configure a VLAN Management Instance (non-routing mode) on a DvR leaf node. After you configure the management VLAN, you can configure a platform VLAN.
- You cannot configure IP Shortcuts on Leaf nodes. This configuration is pushed from the DvR Controllers in the domain.

- You cannot configure IP Multicast over Fabric Connect on DvR Leaf nodes. You cannot deploy IP Multicast over Fabric Connect, statically or dynamically, using EAP on a DvR Leaf. This configuration is pushed from the DvR Controllers in the domain.
- You must manually configure an I-SID on a Layer 2 VSN Leaf node. This configuration is not pushed from a DvR Controller.
- DvR-enabled VLAN/I-SIDs are for host connectivity only; you cannot connect a router to a DvR-enabled VLAN/I-SID and use dynamic or static routing. Use a non-DvR VLAN/I-SID instead to connect an external router.

Restrictions on a DvR Isolated Controller

The following list identifies restrictions of a DvR isolated Controller:

- A DvR isolated Controller does not connect to the DvR backbone and therefore does not share routes with other DvR domains.
- You cannot configure a shared subnet between two DvR domains that do not have a backbone connection.
- All DvR Controllers within the same domain must have the same role, either all conventional DvR Controllers or all isolated DvR Controllers.
- You cannot enable Multi-area SPB DvR backbone redistribution on a boundary node that operates as a DvR isolated Controller.
- Controllers in the same domain must be of the same type (isolated or regular).



Note

If you migrate from a conventional DvR deployment to an isolated DvR deployment, or vice versa, there can be a combination of controller types in the domain until the migration is complete.

- You can enable Multi-area SPB DvR backbone redistribution on a boundary node if the home area or remote area includes a DvR isolated Controller.

For more information about DvR Isolated Domains and Configure an Isolated Controller on a DvR Isolated Domain, see [DvR Isolated Domains](#) on page 920 and [Configure an Isolated Controller on a DvR Isolated Domain](#) on page 941.

Migrate from VRRP to DvR

About This Task

If you have a VRRP network with a mix of existing routers that do not support DvR and devices that do support DvR, you can migrate your VRRP network to DvR using this high-level process. This migration process assumes the following design:

- Existing routers are the VRRP masters.
- Existing routers are the default gateways for all subnets.

- Fabric Connect network with DvR-capable nodes where DvR is configured globally, but not on I-SIDs, on the VOSS devices; and VOSS devices operate in Layer 2 mode for the VRRP VLANs that need to be migrated.

**Important**

When you configure DvR on Controllers with existing VRRP VLANs, ensure there is no VRRP VLAN with VRID 37 or VRID 38. VRID 37 conflicts with the DvR gateway MAC used by all DvR nodes. The DvR gateway MAC is a constant value 00:00:5e:00:01:25; VRRP VRID 37 translates to the same MAC. Similarly, VRRP VRID 38 translates to 00:00:5e:00:01:26, and is used within DvR. If you have a VRRP VLAN with either of these VRIDs, change the VRID to a different value.

Procedure

1. Enable VRRP interfaces on the DvR Controllers but keep VRRP mastership on the existing routers.
2. Change VRRP mastership on the VLAN or IP Subnet in question on the DvR Controller by applying a higher priority than the current master.

**Note**

You can easily fall back to the original VRRP master to change VRRP priorities back.

3. Disable VRRP on the existing routers.
4. Subnet by subnet (VLAN/I-SID), delete VRRP interfaces on all DvR Controllers first (this includes removing VRRP and removing the VLAN IP address), and then configure DvR interfaces (this includes adding the DVR-GW-IP, enabling DvR, and then adding the VLAN IP address) on the VLAN/I-SID instead. This might lead to a short traffic interruption.

**Note**

For each VLAN/I-SID, ensure that VRRP is disabled on all nodes before you configure DvR interfaces on the Controllers for the VLAN/I-SID.

Keep in mind that you can only enable DvR on VLAN or I-SIDs where all participating BEBs are DvR-capable.

Anytime when falling back, you can delete the DvR interface on the I-SID (this includes disabling DvR, removing the DVR-GW-IP and removing the VLAN IP address) and configure the VRRP interface again (this includes adding the VLAN IP address and adding VRRP again), however, ensure you delete the DvR interfaces on all Controllers first before you enable VRRP again.

Example

Start with VRRP VLAN:

```
vlan create 250 name vlan_test250 type port-mstprstp 0
vlan i-sid 250 111250
interface vlan 250
ip address 192.0.2.3 255.255.255.0
interface vlan 250
```



```
ip vrrp version 2
ip vrrp address 10 192.0.2.1
ip vrrp 10 priority 180
ip vrrp 10 backup-master enable
ip vrrp 10 enable
exit
```

Change to DvR VLAN:

```
interface vlan 250
no ip vrrp address 10 192.0.2.1
no ip address 192.0.2.3
exit
interface vlan 250
dvr gw-ipv4 192.0.2.1
dvr enable
ip address 192.0.2.3 255.255.255.0
exit
```

Change back to VRRP VLAN:

```
interface vlan 250
no dvr enable
no dvr gw-ipv4
no ip address 192.0.2.3
exit
interface vlan 250
ip address 192.0.2.3 255.255.255.0
ip vrrp version 2
ip vrrp address 10 192.0.2.1
ip vrrp 10 priority 180
ip vrrp 10 backup-master enable
ip vrrp 10 enable
exit
```

DvR-VRRP Coexistence

When a Layer 2 Virtual Service Network (VSN) extends outside of the Distributed Virtual Routing (DvR) domain, DvR Controllers can route traffic that enters on the network-to-network interface (NNI) by using Virtual Router Redundancy Protocol (VRRP) to advertise the DvR gateway MAC in Layer 2 VSN.

The primary Virtual Router Redundancy Protocol (VRRP) router, in this case, a DvR Controller, sends VRRP advertisements on the DvR-election VLAN using the VRRP gateway MAC as the source and a unique nickname-derived IP address as the source IP. The primary VRRP router sends VRRP advertisements on all other DvR VLANs using the VRRP gateway MAC as the source and the configured interface IP address as the source IP. The primary or backup VRRP routers process VRRP advertisements only on the selected VRRP election VLAN/I-SID. (The primary VRRP router is sometimes referred to as the *master*.)



Note

The VRRP gateway MAC is the same as the DvR gateway MAC.

Non-DvR Backbone Edge Bridges (BEBs) learn the gateway MAC from the advertisements and encapsulate packets destined to the DvR gateway MAC with the

primary VRRP router backbone MAC (BMAC) destination address. The primary VRRP router routes the packets to the destination.

The destination IP address can be any of the following:

- The DvR host route on any DvR Controller or Leaf node in any domain.
- An ARP entry in another DvR VLAN and I-SID for a host that attaches to a non-DvR BEB.
- A subnet route from IS-IS like IP Shortcuts or Layer 3 VSN.
- An ARP entry or subnet route from any protocol in a non-DvR routable VLAN or I-SID.

Restrictions

DvR-VRRP Coexistence has the following restrictions:

- You can only configure this functionality on a DvR Controller.
- Controllers that participate in DvR VRRP must meet the following conditions:
 - Controllers must use the same DvR I-SIDs and the same election I-SID. For example, if Controller 1 and Controller 2 both have DvR VLAN 200 with I-SID 20200 and DvR VLAN 202 with I-SID 20202, they must both use the same I-SID to determine their VRRP role (Primary / Backup).
 - Controllers must be in the same IS-IS area.
- The next-hop IP address for hosts that attach to non-DvR BEBs on a DvR I-SID must be the same as the DvR gateway IP address for that I-SID.
- VRRP Primary and Backup routers must have all interconnected I-SIDs (I-SIDs extended to non-DvR BEBs) configured as DvR interfaces in their domain.
- In any DvR ISID the only supported VRRP routers are DvR Controllers with DvR VRRP enabled. Interoperation with non-DvR VRRP routers is not supported.

DvR and IPv6 VRRP Coexistence on the Same I-SID

You can configure IPv6 VRRP and IPv6 DHCP relay on a DvR-enabled VLAN or I-SID.

You must perform the following steps in order:

1. Enable DvR on a VLAN interface.
2. Configure an IPv6 interface.
3. Configure IPv6 VRRP.

Configuration Examples

The following example shows how you can configure an IPv6 interface and IPv6 VRRP on a DvR-enabled VLAN:

```
vlan create 10 name vlan_test10 type port-mstprstp 0
vlan i-sid 10 111010
interface vlan 10
dvr gw-ipv4 192.0.2.1
dvr enable
ip address 192.0.2.2 255.255.255.0
ipv6 interface enable
ipv6 interface address 2001:DB8:0::1/64
ipv6 vrrp address 1 link-local fe80::1234
```

```
ipv6 vrrp address 1 global 2001:DB8:0::1234/64
ipv6 vrrp 1 enable
exit
```

You cannot enable DvR on a VLAN interface that has IPv6 VRRP configuration. You must first delete IPv6 interface from the VLAN, configure DvR and then reconfigure IPv6 interface.

```
interface vlan 10
no ipv6 vrrp 1
no ipv6 interface
dvr gw-ipv4 192.0.2.1
dvr enable
ip address 192.0.2.2 255.255.255.0
ipv6 interface enable
ipv6 interface address 2001:DB8:0::1/64
ipv6 vrrp address 12 link-local fe80::1234
ipv6 vrrp address 1 global 2001:DB8:0::1234/64
ipv6 vrrp 1 enable
exit
```

DvR configuration using the CLI

The following sections describe configuration of Distributed Virtual Routing (DvR) using the Command Line Interface (CLI).

Configure a DvR Controller

Before You Begin

- Ensure that you configure IP Shortcuts on the node. This is necessary for proper functioning of the node as a DvR Controller.
- Ensure that the `dvr-leaf-mode` boot flag is disabled on the node.

To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

About This Task

Configure a node as a DvR Controller to enable DvR globally on that node.

Perform this procedure to create a DvR domain with the domain ID that you specify, and configure the role of the node as the Controller of that domain. A Controller can belong to only one DvR domain.



Note

For a node to perform the role of both a Controller and a Leaf within a DvR domain, you must configure it as a Controller.

Disabling DvR on a DvR Controller destroys the domain ID and all dynamic content learned within the DvR domain.

However, the switch retains the VLAN specific configuration and you can view the information using the command **show running-config**.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Configure a DVR domain ID:
dvr controller <1-255>
3. View a summary of the Controller configuration:
show dvr

Examples

Configure a node as a DVR Controller:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#dvr controller 5
Switch:1>show dvr
=====
DVR Summary Info
=====
Domain ID           : 5
Domain ISID         : 16678219
Backbone ISID       : 16678216
Role                : Controller
My SYS ID           : 00:bb:00:00:81:21
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs  : Enabled
VRRP Priority        : 100
VRRP Election       : VLAN 10 / I-SID 10
```

Variable definitions

Use the data in the following table to use the **dvr controller** command.

Variable	Value
<1-255>	Specifies the domain ID of the DVR domain that the controller belongs.

Configure an Isolated Controller on a DVR Isolated Domain

Before You Begin

- Ensure that you configure IP Shortcuts on the node.
- Ensure that the **dvr-leaf-mode** boot flag is disabled on the node.

To verify the configuration, enter **show boot config flags** in Privileged EXEC mode.

About This Task

Configure a node as a DVR isolated Controller to enable DVR globally on that node.

Perform this procedure to create a DvR isolated domain with the domain ID that you specify, and configure the role of the node as an isolated Controller of that domain. For more information about the DvR isolated domain and Restrictions on a DvR Isolated Controller, see [DvR Isolated Domains](#) on page 920 and [Restrictions on a DvR Isolated Controller](#) on page 936.

Disabling DvR on a DvR Controller destroys the domain ID and all dynamic content learned within the DvR domain. However, the switch retains the VLAN specific configuration and you can view the information using the command **show running-config**.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Configure an isolated DvR domain ID:
dvr isolated controller <1-255>
3. (Optional) Disable default route injection for the GRT on an isolated Controller:
dvr isolated controller <1-255> inject-default-route-disable
4. View a summary of the Controller configuration:
show dvr

Examples

Configure a node as a DvR Controller for an isolated domain:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#dvr isolated controller 34
Switch:1>show dvr
=====
DVR Summary Info
=====
Domain ID           : 34
Domain ISID         : 16678250
Backbone            : isolated
Role                : Controller
My SYS ID           : e4:5d:52:3c:64:84
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs  : Disabled
VRRP Priority        : 100
VRRP Election       :
```

Disable injection of default routes for the GRT on an isolated Controller.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#dvr isolated controller 34 inject-default-route-disable
Switch:1(config)#show dvr
=====
DVR Summary Info
=====
Domain ID           : 34
```

```

Domain ISID           : 16678250
Backbone              : isolated
Role                  : Controller
My SYS ID             : e4:5d:52:3c:64:84
Operational State     : Up
GW MAC                : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Enabled
VRRP on DVR I-SIDs    : Disabled
VRRP Priority          : 100
VRRP Election         :

```

Variable Definitions

The following table defines parameters for the **dvr isolated controller** command.

Variable	Value
<1-255>	Specifies the domain ID of the DvR isolated domain that the controller belongs.

Disable Injection of Default Routes on a Controller

About This Task

By default, a DvR Controller injects default routes into the DvR domain and all the Leaf nodes in that domain learn these routes with the next hop as the Controller that advertised them.

You can however disable default route injection for the GRT or a specific VRF on a Controller, to override this behavior.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Disable default route injection for the GRT, or a specific VRF, on the Controller:

Command	Description
dvr controller inject-default-route-disable	Use for the GRT.
dvr inject-default-route-disable	Use for a VRF context.

3. Verify the configuration:

Command	Description
show dvr	Use for the GRT.
show dvr l3vsn	Use for a VRF context.

Example

Disable injection of default routes for the GRT on a Controller.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#dvr controller inject-default-route-disable
Switch:1(config)#show dvr
```

```
=====
                        DVR Summary Info
=====
```

Domain ID	: 5
Domain ISID	: 16678219
Backbone ISID	: 16678216
Role	: Controller
My SYS ID	: 00:bb:00:00:81:21
Operational State	: Up
GW MAC	: 00:00:5e:00:01:25
InjectDefaultRouteDisable (GRT)	: Enabled
VRRP on DVR I-SIDs	: Disabled
VRRP Priority	: 100
VRRP Election	:

Disable injection of default routes for a specific VRF on a Controller.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router vrf vrf3
Switch:1(router-vrf)#dvr inject-default-route-disable
Switch:1(router-vrf)#show dvr l3vsn
```

```
=====
                        DVR L3VSN
=====
```

VRF ID	L3VSN ISID	VRF NAME	INJECT-DEFAULT-ROUTE-DISABLE
1	50	green	Disabled
7	1000003	vrf3	Enabled

2 out of 2 Total Num of DVR L3VSN displayed

Configure DvR Route Redistribution

About This Task

Configure redistribution of direct or static routes into the DvR domain, on the Global Router or for a specific VRF instance.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure route redistribution of direct routes:
 - a. Configure route redistribution of direct routes on a VRF. The route type is internal .
`dvr redistribute direct [metric <0-65535>] |[route-map WORD<1-64>]`
 - b. Enable route redistribution.
`dvr redistribute direct enable`
 - c. Apply the configuration:
`dvr apply redistribute direct`
 - d. (Optional) Disable route redistribution of direct routes.
`no dvr redistribute direct`
3. Configure route redistribution of static routes:
 - a. Configure route redistribution of static routes on a VRF. The route type is external.
`dvr redistribute static [metric <0-65535>] |[route-map WORD<1-64>]`
 - b. Enable route redistribution.
`dvr redistribute static enable`
 - c. Apply the configuration.
`dvr apply redistribute static`
 - d. (Optional) Disable route redistribution of static routes.
`no dvr redistribute static`
4. Verify the route redistribution configuration. You can also verify it on a specific VRF instance.
`show dvr redistribute [vrf WORD<1-16>]`

Example

Configure route redistribution of direct and static routes on the Global Router. Ensure that you apply the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#dvr redistribute static
Switch:1(config)#dvr redistribute static metric 200
```

```
Switch:1(config)#dvr redistribute static enable
Switch:1(config)#dvr apply redistribute static

Switch:1(config)#dvr redistribute direct
Switch:1(config)#dvr redistribute direct metric 100
Switch:1(config)#dvr redistribute direct enable
Switch:1(config)#dvr apply redistribute direct
```

Verify configuration on the Global Router:

```
Switch:1(config)#show dvr redistribute
=====
                        DVR Redistribute List - GlobalRouter
=====

SOURCE MET MTYPE      ENABLE RPOLICY
-----
STAT   200 External   TRUE  -
LOC    100 Internal   TRUE  -
```

Configure redistribution of direct and static routes on the specific VRF instance `vrf1`. Ensure that you apply the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#router vrf vrf1
Switch:1(router-vrf)#dvr redistribute static
Switch:1(router-vrf)#dvr redistribute static metric 20000
Switch:1(router-vrf)#dvr redistribute static enable
Switch:1(router-vrf)#exit
Switch:1(config)#dvr apply redistribute static

Switch:1(router-vrf)#dvr redistribute direct
Switch:1(router-vrf)#dvr redistribute direct metric 10000
Switch:1(router-vrf)#dvr redistribute direct enable
Switch:1(router-vrf)#exit
Switch:1(config)#dvr apply redistribute static
```

Verify configuration on `vrf1`:

```
Switch:1(router-vrf)#show dvr redistribute vrf vrf1
=====
                        DVR Redistribute List - VRF vrf1
=====

SOURCE MET MTYPE      ENABLE RPOLICY
-----
STAT   20000 External   TRUE  -
LOC    10000 Internal   TRUE  -
```


Variable definitions

Use the data in the following table to use the **dvr redistribute direct** or the **dvr redistribute static** commands.

Variable	Value
<i>enable</i>	Enables DvR route redistribution on the VRF instance. Route redistribution is enabled by default.
<i>metric <0-65535></i>	Specifies the DvR route redistribution metric.
<i>route-map WORD<1-64></i>	Specifies the route policy for DvR route redistribution.

Use the data in the following table to use the **show dvr redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF name.

Clear DvR Host Entries

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

Clear DvR host entries (IPv4 remote host routes) on a Controller. The host entries are learned on the switch, either locally on its UNI port or dynamically from other nodes in the DvR domain.



Note

You can clear DvR host entries only on a DvR Controller.

An error message displays if you attempt clearing of host entries on a DvR Leaf node.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Clear the DvR host entries.

```
clear dvr host-entries [ipv4 {A.B.C.D}] | [l2isid <0-16777215>] |  
[l3isid <0-16777215>]
```

Example

In this example, you clear host entries for IP address 50.0.1.0 to clear host entries for IP addresses 50.0.1.2 and 50.0.1.3.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clear dvr host-entries 50.0.1.0
```

Variable definitions

Use the data in the following table to use the **clear dvr host-entries** command.

Variable	Value
<i>ipv4</i>	Specifies the IP address (IPv4) of the DvR host entries to clear.
<i>l2isid</i>	Specifies the Layer 2 VSN I-SID of the DvR host entries to clear The range is 1 to 16777215.
<i>l3isid</i>	Specifies the Layer 3 VSN I-SID of the DvR host entries to clear. The range is 0 to 16777215.

Configure a DvR Leaf

Before You Begin



Caution

Ensure that you save the current configuration on the switch, before you enable the flag.

Enabling the flag removes all existing non-DvR configuration on the switch, such as platform VLANs and their IP address configuration, circuitless IP (CLIP) configuration, routing protocol configuration and VRF configuration. The gateway IPv4 address, if configured, is also removed.

About This Task

Perform this procedure to create a DvR domain with the domain ID that you specify, and configure the role of the node as a Leaf node. Configuring a node as a DvR Leaf automatically enables DvR globally on the node.

A Leaf node can belong to only one DvR domain.



Note

For a node to perform the role of both a Controller and a Leaf within the domain, you must configure it as a Controller.

**Note**

You must enable the VRF-scaling boot configuration flag on a DvR Leaf node, if more than 24 VRFs are required in the DvR domain.

For additional scaling information, see *Fabric Engine Release Notes*.

**Caution**

Disabling DvR on a Leaf node removes its membership with the DvR domain and all the dynamic content learned from the Controllers of that domain.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable the DvR boot configuration flag:

```
boot config flags dvr-leaf-mode
```

**Note**

When you enable the `dvr-leaf-mode` boot flag, you can configure the node as a DvR leaf node without rebooting, as long as there is no unsupported configuration discovered on the switch.

3. Configure a node as a DvR Leaf.

```
dvr leaf <1-255>
```
4. View a summary of the Leaf configuration.

```
show dvr
```
5. Restart the switch for your change to take effect.

Example

Configure a node as a DvR Leaf:

```
Switch2:1>enable
Switch2:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config)#boot config flags dvr-leaf-mode
Switch2:1(config)#save config
Switch2:1(config)#dvr Leaf 5
Switch2:1(config)#show dvr
```

```
=====
                        DVR Summary Info
=====
```

Domain ID	: 5
Domain ISID	: 16678219
Role	: Leaf
My SYS ID	: 00:00:72:54:44:00
Operational State	: Up
GW MAC	: 00:00:5e:00:01:25
Inband Mgmt Clip IP	:
Virtual Ist local address	:
Virtual Ist local subnet mask	:

```
Virtual Ist peer address      :  
Virtual Ist cluster-id      :  
Virtual Ist ISID            :
```

Variable Definitions

Use the data in the following table to use the **dvr leaf** command.

Variable	Value
<1-255>	Specifies the domain ID of the DvR domain to which the Leaf node belongs.

Configure vIST on a DvR Leaf Node Pair

Before You Begin

Ensure that the nodes are configured as DvR Leaf nodes, before you configure vIST.

About This Task

When you configure vIST on a DvR Leaf node pair, the switch generates an I-SID from the configured cluster ID. This I-SID is unique across the SPB network as long as the cluster ID is unique across the SPB network, for the vIST pair. You can configure only one instance of vIST on the Leaf node pair.

To configure vIST, both nodes must be Leaf nodes. You cannot configure vIST, for example, on a Controller-Leaf node pair.

Also both the nodes must belong to the same DvR domain. vIST configuration over Leaf nodes in different domains is not supported.

Procedure

1. Enter Global Configuration mode:

```
enable  
  
configure terminal
```
2. Configure vIST on the Leaf nodes:

```
dvr leaf virtual-ist {<A.B.C.D/X|<A.B.C.D> <A.B.C.D>} peer-ip  
{A.B.C.D} cluster-id <1-1000>
```
3. (Optional) Disable vIST on the DvR Leaf node pair.

```
no dvr leaf virtual-ist
```



Caution

Disabling DvR on a Leaf node in a vIST pair removes all vIST configuration on that node, but not on the pair. The node on which DvR is disabled also loses its membership with the DvR domain and all the dynamic content learned from the Controllers in that domain.

If DvR is re-enabled on the node, you must manually configure vIST on that node again.

4. View a summary of vIST configuration on the Leaf nodes.

```
show dvr
```

Example

Configure vIST on DvR Leaf nodes, with IP addresses 51.51.51.1 and 51.51.51.2 respectively:

```
Switch2:1>enable
Switch2:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config)#dvr leaf virtual-ist 51.51.51.1 peer-ip 51.51.51.2 cluster-id 255
Switch2:1#show dvr

=====
                        DVR Summary Info
=====
Domain ID                : 5
Domain ISID              : 16678219
Role                     : Leaf
My SYS ID                : 00:bb:00:00:71:23
Operational State        : Up
GW MAC                   : 00:00:5e:00:01:25
Inband Mgmt Clip IP      :
Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address  : 51.51.51.2
Virtual Ist cluster-id    : 255
Virtual Ist ISID         : 16677226
```

Variable definitions

Use the data in the following table to use the **dvr leaf virtual-ist** command.

Variable	Value
{<A.B.C.D/X <A.B.C.D> <A.B.C.D>}	Specifies the local IP (IPv4) address and subnet mask of the node.
{<A.B.C.D>}	Specifies the IP address (IPv4) of the vIST peer.
<1-1000>	Specifies the cluster ID of vIST. It is set to 0 if vIST is not configured.

Configure a Management VLAN on a DvR Leaf Node

About This Task

On a DvR leaf node, you can configure a Management Instance VLAN (non-routing mode) by specifying the I-SID. Use the following procedure to configure a Management Instance VLAN on a DvR leaf node.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a Management Instance VLAN I-SID:

```
mgmt vlan i-sid <1-16777215>
```

Variable Definitions

The following table defines parameters for the **mgmt vlan i-sid** command.

Variable	Value
1-16777215	Specifies the VLAN I-SID to associate with the management VLAN.

Move a vIST Leaf Node Pair from one Domain to Another

About This Task

Use this procedure to move a vIST Leaf node pair from one DvR domain to another.

For vIST to work properly, both Leaf nodes must be in the same domain.

Procedure

1. Disable IS-IS on each vIST peer Leaf node, to remove the node from the SPB network.

```
no router isis enable
```

2. Disable DvR on each Leaf node.

```
no dvr leaf
```



Caution

Disabling DvR on a Leaf node in a vIST pair automatically removes all vIST configuration on that node, but not on the pair. The node on which DvR is disabled also loses its membership with the DvR domain and all the dynamic content learned from the Controllers in that domain.

When you re-enable DvR on the node, you must manually configure vIST on that node again.

3. Configure each node as a DvR Leaf node, with the new domain ID.

Ensure that you configure both nodes as Leaf nodes and with the same domain ID.

```
dvr leaf <1-255>
```

4. Configure vIST on the DvR Leaf nodes.

```
dvr leaf virtual-ist {<A.B.C.D/X|<A.B.C.D> <A.B.C.D>} peer-ip  
{A.B.C.D} cluster-id <1-1000>
```

5. Enable IS-IS on each vIST peer Leaf node, to add back the node to the SPB network.

```
router isis enable
```

Example

Consider two vIST peer Leaf nodes Switch1 (IP address 51.51.51.1) and Switch2 (51.51.51.2) that belong to a DVR domain (with domain ID 4), that you need to move to another domain (with domain ID 5).

View a summary of existing Leaf configuration on each node.

```
Switch1:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 4
Domain ISID         : 16678220
Role                : Leaf
My SYS ID           : 00:00:72:54:44:00
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
Inband Mgmt Clip IP :
Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID     : 16677226

Switch2:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 4
Domain ISID         : 16678220
Role                : Leaf
My SYS ID           : 00:00:72:55:45:00
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
Inband Mgmt Clip IP :
Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID     : 16677226
```

Disable IS-IS globally on each Leaf node.

```
Switch1:1>en
Switch1:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config)#router isis
Switch1:1(config-isis)#no router isis enable
Switch1:1(config-isis)#exit
Switch1:1(config)#

Switch2:1>en
Switch2:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config)#router isis
Switch2:1(config-isis)#no router isis enable
Switch1:1(config-isis)#exit
Switch1:1(config)#
```

Disable DvR on each node. This automatically removes all vIST configuration on the node, but not on the vIST pair. The node also loses its membership with the DvR domain and all the dynamic content learned from the Controllers in that domain.

```
Switch1:1(config)#no dvr leaf
Switch2:1(config)#no dvr leaf
```

Configure each node as a DvR Leaf, with domain ID 5.

```
Switch1:1(config)#dvr leaf 5
Switch2:1(config)#dvr leaf 5
```

Configure vIST on each of the DvR Leaf nodes.

```
Switch1:1(config)#dvr leaf virtual-ist 51.51.51.1 peer-ip 51.51.51.2 cluster-id 255
Switch2:1(config)#dvr leaf virtual-ist 51.51.51.2 peer-ip 51.51.51.1 cluster-id 255
```

Enable IS-IS globally on each Leaf node.

```
Switch1:1(config)#router isis
Switch1:1(config-isis)#router isis enable
Switch1:1(config-isis)#exit
Switch1:1(config)#

Switch2:1(config)#router isis
Switch2:1(config-isis)#router isis enable
Switch2:1(config-isis)#exit
Switch2:1(config)#
```

View a summary of Leaf configuration on each node.

```
Switch1:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 5
Domain ISID         : 16678221
Role                : Leaf
My SYS ID           : 00:00:72:54:44:00
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
Inband Mgmt Clip IP :
Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID     : 16677226

Switch2:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 5
Domain ISID         : 16678221
Role                : Leaf
My SYS ID           : 00:00:72:55:45:00
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
Inband Mgmt Clip IP :
Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
```



```
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id   : 255
Virtual Ist ISID         : 16677226
```

Move a vIST Controller Pair from One Domain to Another

Use this procedure to move a vIST Controller node pair from one DvR domain to another.

About This Task

For vIST to work properly, both Controller nodes must be in the same domain.

Procedure

1. Disable IS-IS on each vIST peer Controller node, to remove the node from the SPB network.

```
no router isis enable
```

2. Disable DvR on each Controller node:

```
no dvr controller
```



Caution

Disabling DvR on a DvR Controller destroys the domain ID and all dynamic content learned within the DvR domain. However, the switch retains the VLAN specific configuration which you can view using the command **show running-config**.

3. Configure each node as a DvR Controller node, with the new domain ID. Ensure that you configure both nodes as Controller nodes and with the same domain ID.

```
dvr controller <1-255>
```

4. Enable IS-IS on each vIST peer Controller node, to add back the node to the SPB network.

```
router isis enable
```

Example

Consider two vIST peer Controller nodes Switch1 (IP address 51.51.51.3) and Switch2 (51.51.51.4) that belong to a DvR domain (with domain ID 4), that you need to move to another domain (with domain ID 5).

View a summary of Controller configuration on each node:

```
Switch1:1(config)#show dvr
```

```
=====
                        DVR Summary Info
=====
Domain ID                : 4
Domain ISID              : 16678220
Backbone ISID            : 16678216
Role                     : Controller
My SYS ID                : 00:bb:00:00:81:21
Operational State        : Up
```

```

GW MAC                : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs    : Disabled
VRRP Priority          : 100
VRRP Election         :

```

```
Switch2:1(config)#show dvr
```

```

=====
                        DVR Summary Info
=====
Domain ID              : 4
Domain ISID            : 16678220
Backbone ISID          : 16678216
Role                   : Controller
My SYS ID              : 00:bb:00:00:82:22
Operational State      : Up
GW MAC                 : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs    : Disabled
VRRP Priority          : 100
VRRP Election         :

```

Disable IS-IS globally on each Controller node:

```

Switch1:1>en
Switch1:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config)#router isis
Switch1:1(config-isis)#no router isis enable
Switch1:1(config-isis)#exit
Switch1:1(config)#

Switch2:1>en
Switch2:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config)#router isis
Switch2:1(config-isis)#no router isis enable
Switch1:1(config-isis)#exit
Switch1:1(config)#

```

Disable DvR on each node:

```

Switch1:1(config)#no dvr controller
Switch2:1(config)#no dvr Controller

```

Configure each node as a DvR Controller, with domain ID 5.

```

Switch1:1(config)#dvr controller 5
Switch2:1(config)#dvr controller 5

```

Enable IS-IS globally on each Controller node.

```

Switch1:1(config)#router isis
Switch1:1(config-isis)#router isis enable
Switch1:1(config-isis)#exit
Switch1:1(config)#

Switch2:1(config)#router isis
Switch2:1(config-isis)#router isis enable
Switch2:1(config-isis)#exit
Switch2:1(config)#

```

View a summary of Controller configuration on each node.

```
Switch1:1(config)#show dvr

=====
                        DVR Summary Info
=====
Domain ID                : 5
Domain ISID              : 16678221
Backbone ISID            : 16678216
Role                     : Controller
My SYS ID                : 00:bb:00:00:81:21
Operational State        : Up
GW MAC                   : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs       : Disabled
VRRP Priority             : 100
VRRP Election            :

Switch2:1(config)#show dvr

=====
                        DVR Summary Info
=====
Domain ID                : 5
Domain ISID              : 16678221
Backbone ISID            : 16678216
Role                     : Controller
My SYS ID                : 00:bb:00:00:82:22
Operational State        : Up
GW MAC                   : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs       : Disabled
VRRP Priority             : 100
VRRP Election            :
```

View the vIST configuration on each of the Controller nodes.

```
Switch1:1>show virtual-ist

=====
                        IST Info
=====
PEER-IP      VLAN      ENABLE   IST
ADDRESS      ID        IST      STATUS
-----
51.51.51.2   4002    true    up

NEGOTIATED                                MASTER/
DIALECT      IST STATE                                SLAVE
-----
NONE         up                                     Master

Switch2:1>show virtual-ist

=====
                        IST Info
=====
PEER-IP      VLAN      ENABLE   IST
ADDRESS      ID        IST      STATUS
-----
51.51.51.1   4002    true    up
```

NEGOTIATED DIALECT	IST STATE	MASTER/ SLAVE
NONE	up	Slave

Configure a non-DvR BEB to Join the DvR Backbone

About This Task

Configure a non-DvR backbone edge bridge (BEB) to join the DvR backbone so that it can receive redistributed DvR host routes from all DvR Controllers in the SPB network.



Note

On a non-DvR BEB, the redistributed host routes from the DvR backbone are not automatically installed in the IP routing table. To utilize the backbone host routes to optimize traffic forwarding (forwarding in the data plane), you must explicitly configure an IS-IS accept policy with a backbone route policy using the command **accept backbone-route-map <route-map-name>**, and specifying a suitable route-map to select the list or range of DvR backbone host routes to be installed in the routing table.

For more information on configuring an IS-IS accept policy with a backbone route policy, see [Configure IS-IS Accept Policies](#) on page 1616.

Procedure

1. Enter IS-IS Router Configuration mode:


```
enable

configure terminal

router isis
```
2. Configure a non-DvR BEB to join the DvR backbone.


```
backbone enable
```
3. Verify the configuration using the following commands.
 - `show dvr backbone-members`
 - `show dvr backbone-members non-dvr-beb`
 - `show dvr backbone-entries`
 - `show isis`

Examples

```
Switch3:1>enable
Switch3:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch3:1(config)#router isis
Switch3:1(config-isis)#show dvr
```

```
=====
                                NON DVR BEB Summary Info
=====
Domain ID                        : 0
```

```

Domain ISID           : 0
Backbone ISID         : 16678216
Role                  : NON DVR BEB
My SYS ID             : 00:00:82:84:40:00
Operational State     : Up

```

Configure the non-DvR BEB to join the DvR backbone.

```
Switch3:1(config-isis)#backbone enable
```

Verify the configuration. View the DvR backbone members.

```
Switch3:1(config-isis)#show dvr backbone-members
```

```

=====
DVR BB Members
=====
System Name      Nick-Name      Nodal MAC      Role      Domain Id      Area      Area-Name
-----
DVR-8284-D2-C1-40  0.82.40      00:00:82:84:40:00  NON-DVR-BEB  9999      HOME      area-0.00.20
DVR-8284-D2-C2-41  0.82.41      00:00:82:84:41:00  Controller  9999      HOME      area-0.00.20

Home: 2 out of 2 Total Num of DVR Backbone Members displayed
=====

Switch3:1(config-isis)#show dvr backbone-members non-dvr-beb
=====
DVR BB Members
=====
System Name      Nick-Name      Nodal MAC      Role      Domain Id      Area      Area-Name
-----
DVR-8284-D2-C1-40  0.82.40      00:00:82:84:40:00  NON-DVR-BEB  9999      HOME      area-0.00.20

Home: 1 out of 2 Total Num of DVR Backbone Members displayed
=====

```

View the backbone DvR host routes that the non-DvR BEB receives from other Controllers in the SPB network.

```
Switch:1(config-isis)#show dvr backbone-entries
```

```

=====
DVR Backbone-Entries
=====
IP-ADDRESS      HOST      L3VSN      L2VSN      DOMAIN      ADV-CONTROLLER      NEXT HOP      AREA      AREA-NAME
-----
39.1.1.4        10:cd:ae:70:5d:01  401      10390      200      DVR-8284-D2-C2-41  DVR-8284-D2-C2-41  REMOTE      area-0.00.30
39.2.1.4        10:cd:ae:70:5d:01  401      10391      200      DVR-8284-D2-C2-41  DVR-8284-D2-C2-41  REMOTE      area-0.00.30
39.3.1.4        10:cd:ae:70:5d:01  401      10392      200      DVR-8284-D2-C2-41  DVR-8284-D2-C2-41  REMOTE      area-0.00.30
39.4.1.4        10:cd:ae:70:5d:01  401      10393      200      DVR-8284-D2-C2-41  DVR-8284-D2-C2-41  REMOTE      area-0.00.30
39.5.1.4        10:cd:ae:70:5d:01  401      10394      200      DVR-8284-D2-C2-41  DVR-8284-D2-C2-41  REMOTE      area-0.00.30
39.6.1.4        10:cd:ae:70:5d:01  401      10395      200      DVR-8284-D2-C2-41  DVR-8284-D2-C2-41  REMOTE      area-0.00.30

Remote: 6 out of 427 Total Num of DVR Backbone Routes displayed
=====

```

View the IS-IS related information.

```
Switch3:1(config-isis)#show isis

=====
                        ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID : 00bb.0000.8121
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name : router_r1
ip source-address :
ipv6 source-address :
ip tunnel source-address :
Tunnel vrf :
ip tunnel mtu :
Num of Interfaces : 3
Num of Area Addresses : 1
Inband Mgmt Clip IP :72.54.44.1
backbone :enabled
Dynamically Learned Area : 00.0000.0000
FAN Member : No
Hello Padding : enabled
Multi-Area OperState : disabled
Multi-Area Flags :
```

Configure DvR-VRRP Coexistence



Note

This procedure does not apply to 5320 Series or 5420 Series.

When a Layer 2 Virtual Service Network (VSN) extends outside of the DvR domain, DvR Controllers can route traffic that enters on the network-to-network interface (NNI) by using Virtual Router Redundancy Protocol (VRRP).

Before You Begin

The VLAN must be a DvR VLAN. A DvR VLAN is a VLAN configured on a DvR Controller with a VLAN IP address, a VLAN I-SID, the DvR gateway IP address, and DvR enabled.

About This Task

If you enable this functionality, the switch automatically creates a VRRPv2 interface with the virtual router ID (VRID) 37 and a *holddown-timer* configuration of 60 on the DvR VLAN that you identify for VRRP election. You can configure one VLAN for election.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```
2. Enable VRRP election on the DvR VLAN:

```
dvr vrrp-election
```
3. Exit to Global Configuration mode:

```
exit
```
4. (Optional) Configure the DvR VRRP priority for the elected VLAN:

```
dvr controller vrrp-on-dvr-isids priority <1-254>
```
5. Enable VRRP advertisements on all DvR VLAN Service Instance Identifiers (I-SIDs):

```
dvr controller vrrp-on-dvr-isids enable
```

Variable Definitions

The following table defines parameters for the **dvr controller vrrp-on-dvr-isids priority** command.

Variable	Value
<1-254>	Configures the Virtual Router Redundancy Protocol (VRRP) priority for the elected VLAN and associated Service Instance Identifier (I-SID). Switches advertise their priority to use in the primary router election process. When multiple Controllers elect the same DvR VLAN and I-SID, the Controller with the higher priority becomes the primary router. The elected Controller sends VRRP advertisements on all DvR VLANs. The default value is 100.

DvR show commands

The following section explains the show commands for DvR.

View DvR Summary

Use this procedure to view a summary of the DvR configuration on a DvR Controller or a DvR Leaf.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. View a summary of DvR configuration:

```
show dvr
```

Example

View the information on a DvR Controller:

```
Switch:1>show dvr
=====
DVR Summary Info
=====
Domain ID           : 5
Domain ISID         : 16678219
Backbone ISID       : 16678216
Role                : Controller
My SYS ID           : 00:bb:00:00:81:21
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs  : Enabled
VRRP Priority        : 100
VRRP Election       : VLAN 10 / I-SID 10
```

View the information on a DvR Leaf:

```
Switch2:1#show dvr
=====
DVR Summary Info
=====
Domain ID           : 5
Domain ISID         : 16678219
Role                : Leaf
My SYS ID           : 00:bb:00:00:71:23
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
Inband Mgmt Clip IP : 72.54.44.1
Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID     : 16677226
```

View Members of a DvR Domain

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the members of all DvR domains, namely the Controllers and Leaf nodes.

You can view this information on either a Controller or a Leaf node. Both the Controller and the Leaf node displays those members of the DvR domain to which it belongs.

Procedure

1. Enter Privileged EXEC mode:
enable
2. show dvr members [controller|leaf]

Example

View all members of a DvR domain:

```
Switch:1#show dvr members
```

```
=====
                        DVR Members (Domain ID: 255)
=====
System Name                Nick-Name          Nodal MAC          Role
-----
Leaf-4:110                  0.41.10         00:bb:00:00:41:10   Leaf
Leaf-1:Q:123                0.71.23         00:bb:00:00:71:23   Leaf
Leaf-2:K:124                0.71.24         00:bb:00:00:71:24   Leaf
Leaf-3:K:125                0.71.25         00:bb:00:00:71:25   Leaf
Ctrl-1:Q:121                0.81.21         00:bb:00:00:81:21   Controller
Ctrl-2:Q:122                0.81.22         00:bb:00:00:81:22   Controller

6 out of 6 Total Num of DVR Members displayed
=====
```

View member DvR Controllers:

```
Switch:1#show dvr members controller
```

```
=====
                        DVR Members (Domain ID: 255)
=====
System Name                Nick-Name          Nodal MAC          Role
-----
Ctrl-1:Q:121                0.81.21         00:bb:00:00:81:21   Controller
Ctrl-2:Q:122                0.81.22         00:bb:00:00:81:22   Controller

2 out of 6 Total Num of DVR Members displayed
=====
```

View member DvR Leaf nodes:

```
Switch:1#show dvr members leaf
```

```
=====
                        DVR Members (Domain ID: 255)
=====
System Name                Nick-Name          Nodal MAC          Role
-----
Leaf-4:110                  0.41.10         00:bb:00:00:41:10   Leaf
Leaf-1:Q:123                0.71.23         00:bb:00:00:71:23   Leaf
Leaf-2:K:124                0.71.24         00:bb:00:00:71:24   Leaf
Leaf-3:K:125                0.71.25         00:bb:00:00:71:25   Leaf

4 out of 6 Total Num of DVR Members displayed
=====
```

View DvR Interfaces

View the DvR interfaces on either a Controller or a Leaf node.

On Controllers, DvR interfaces are created when you configure IP on a DvR enabled Layer 2 VSN (VLAN, I-SID). Only Controllers display the administrative state of the interfaces because this is where you enable or disable the interfaces. The Leaf nodes display DvR interface information that is pushed from the Controllers, for example, subnet routes or gateway IP addresses for the Layer 2 VSNs.

Before You Begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the DvR interface information.

On a Controller:

```
show dvr interfaces [l3isid <0-16777215>] [vrf WORD<1-16>] [vrfrids  
WORD<0-512>]
```

On a Leaf node:

```
show dvr interfaces [l3isid <0-16777215>]
```

Viewing the DvR interface information for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View DvR interfaces on a Controller node:

You can view DvR interface information on all interfaces or for a specific Layer 3 I-SID, VRF, or VRF ID.

```
Switch:1#show dvr interfaces

=====
DVR Interfaces
=====
Interface      Mask      L3ISID  VRFID    L2ISID    VLAN    GW IPv4    Admin  SPBMC  IGMP
State          State          Version
-----
50.0.1.2       255.255.0.0  55500   1         50500     500     50.0.1.1  enable  disable 2
1 out of 1 Total Num of DVR Interfaces displayed
=====
```

View DvR interfaces on a Leaf node:

You can view DvR interface information on all interfaces or for a specific Layer 3 I-SID. Viewing the interface information for a specific VRF or VRF ID is not supported on a DvR Leaf node.

```
Switch:1#show dvr interfaces l3isid 401

=====
DVR Interfaces
=====
Interface      Mask      L3ISID  VRFID    L2ISID    VLAN    GW IPv4
-----
40.1.0.0       255.255.0.0  401     2         10401     77     40.1.1.11
40.2.0.0       255.255.0.0  401     2         10402     78     40.2.1.11
40.3.0.0       255.255.0.0  401     2         10403     79     40.3.1.11
40.4.0.0       255.255.0.0  401     2         10404     80     40.4.1.11
```

```
4 out of 4 Total Num of DVR Interfaces displayed
-----
```

Variable definitions

Use the data in the following table to use the **show dvr interfaces** command.

Variable	Value
<i>l3isid</i>	Specifies the Layer 3 I-SID of the DvR interface. The range is 0 to 16777215.
<i>vrf</i>	Specifies the VRF name.
<i>vrfids</i>	Specifies the VRF ID. The range is 0 to 512.

View DvR Host Entries

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View DvR host entries (IPv4 remote host routes) on either a Controller or a Leaf node. The node displays the host entries learned either locally on its Switched UNI port or dynamically from other nodes within the DvR domain.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the DvR host entries.

On a Controller:

```
show dvr host-entries [domain-id <1-255>] [[ipv4 {A.B.C.D}]] [[l2isid <1-16777215>]] [[l3isid <0-16777215>]] [[nh-as-mac]] [[type <1-2>]] [[vrf WORD<1-16>] [vrfids WORD<0-512>]]
```

On a Leaf node:

```
show dvr host-entries [domain-id <1-255>] [[ipv4 {A.B.C.D}]] [[l2isid <1-16777215>]] [[l3isid <0-16777215>]] [[nh-as-mac]] [[type <1-2>]]
```

Viewing the DvR host entries for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View DvR host entries on either a Controller or a Leaf node.

Viewing the DvR host entries for a specific VRF or VRF ID is not supported on a DvR Leaf node.

```
Switch:1#show dvr host-entries domain-id 255 l3isid 55500
```

```
=====
                        DVR Host-Entries
=====
IP-ADDRESS      HOST          L3VSN    L2VSN      DOMAIN
MAC-ADDRESS     ISID       ISID      VRFID     PORT      ID      TYPE      NEXT HOP
-----
50.0.1.2        b0:ad:aa:42:ed:04 55500    50500     0         1/23    255      DYNAMIC   Cont-1:121
50.0.1.3        b0:ad:aa:4c:3d:01 55500    50500     0         cpp      255      LOCAL     Cont-2:122

2 out of 2 Total Num of DVR Host Entries displayed
=====
```

View DvR host entries for a specific IP address.

In this example, you enter IP address 50.0.1.0 to display host entries for IP addresses 50.0.1.2 and 50.0.1.3.

```
Switch:1#show dvr host-entries ipv4 50.0.1.0
```

```
=====
                        DVR Host-Entries
=====
IP-ADDRESS      HOST          L3VSN    L2VSN      DOMAIN
MAC-ADDRESS     ISID       ISID      VRFID     PORT      ID      TYPE      NEXT HOP
-----
50.0.1.2        b0:ad:aa:42:ed:04 55500    50500     0         1/23     2      DYNAMIC   Cont-1:121
50.0.1.3        b0:ad:aa:4c:3d:01 55500    50005     0         cpp       2      LOCAL     Cont-2:122

2 out of 2 Total Num of DVR Host Entries displayed
=====
```

View DvR host entries where the next hop displays the MAC address instead of the system name.

```
Switch:1#show dvr host-entries nh-as-mac
```

```
=====
                        DVR Host-Entries
=====
IP-ADDRESS      HOST          L3VSN    L2VSN      DOMAIN
MAC-ADDRESS     ISID       ISID      VRFID     PORT      ID      TYPE      NEXT HOP
-----
50.0.1.2        b0:ad:aa:42:ed:04 55500    50500     0         1/23     2      DYNAMIC   00:bb:00:00:01:01
50.0.1.3        b0:ad:aa:4c:3d:01 55500    50500     0         cpp       2      LOCAL     00:bb:00:00:01:02

2 out of 2 Total Num of DVR Host Entries displayed
=====
```

View DvR host entries based on the host type. Type 1 indicates local hosts and type 2 dynamic hosts.

```
Switch:1#show dvr host-entries type 2
```

```
=====
                        DVR Host-Entries
=====
HOST          L3VSN    L2VSN      DOMAIN
```

IP-ADDRESS	MAC-ADDRESS	ISID	ISID	VRFID	PORT	ID	TYPE	NEXT HOP
50.0.1.2	b0:ad:aa:42:ed:04	55500	50500	0	1/23	2	DYNAMIC	00:bb:00:00:01:01
1 out of 2 Total Num of DVR Host Entries displayed								

Variable definitions

Use the data in the following table to use the **show dvr host-entries** command.

Variable	Value
<i>domain-id</i>	Specifies the domain ID of the DvR host entry. The range is 1 to 255.
<i>ipv4</i>	Specifies the IP address (IPv4) of the DvR host entry.
<i>l2isid</i>	Specifies the Layer 2 VSN I-SID of the DvR host entry. The range is 1 to 16777215.
<i>l3isid</i>	Specifies the Layer 3 VSN I-SID of the DvR host entry. The range is 0 to 16777215.
<i>nh-as-mac</i>	Specifies the MAC address of the next hop node instead of the system name.
<i>type</i>	Specifies the host type of the DvR host entry. A value of 1 indicates local hosts and a value of 2 indicates dynamic hosts.
<i>vrf</i>	Specifies the VRF name of the DvR host entry.
<i>vrfids</i>	Specifies the VRF ID of the DvR host entry. The range is 0 to 512.

View DvR Routes

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the DvR routes (IPv4 network routes) on a DvR Controller or a Leaf node.

Controllers display all the IP subnet routes configured for that DvR domain. The Leaf nodes display the IP subnet routes that are learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. Leaf nodes also display routes that are redistributed by Controllers (direct routes, static routes and the default route), into the DvR domain.

Procedure

1. Enter Privileged EXEC mode:
enable

2. View the DvR routes.

On a Controller:

```
show dvr routes [ipv4 {A.B.C.D}]|[l3isid <0-16777215>]|[nh-as-mac] |
[vrf WORD<1-16>]|[vrfids WORD<0-512>]
```

On a Leaf node:

```
show dvr routes [ipv4 {A.B.C.D}]|[l3isid <0-16777215>]|[nh-as-mac]
```

Viewing the DvR routes for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View DvR routes on either a Controller or a Leaf node.

Viewing the DvR routes for a specific VRF or VRF ID is not supported on a DvR Leaf node.

```
Switch:1#show dvr routes
```

DVR Routes							
DEST	MASK	NEXT HOP	VRFDID	L3VSN ISID	L2VSN ISID	TYPE	COST
50.0.0.0	255.255.0.0	Ctrl-1:8400:121	0	55500	50500	-	1
1 out of 1 Total Num of DVR Routes displayed							
TYPE Legend: E=Ecmp Route							

View DvR routes where the next hop MAC address is displayed instead of the system name:

```
Switch:1#show dvr routes nh-as-mac
```

DVR Routes							
DEST	MASK	NEXT HOP	VRFDID	L3VSN ISID	L2VSN ISID	TYPE	COST
50.0.0.0	255.255.0.0	00:bb:00:00:01:02	0	55500	50500	-	1
1 out of 1 Total Num of DVR Routes displayed							
TYPE Legend: E=Ecmp Route							

Variable definitions

Use the data in the following table to use the **show dvr routes** command.

Variable	Value
<i>ipv4 {A.B.C.D}</i>	Specifies the IP address (IPv4) of the DvR route.
<i>l3isid <0-16777215></i>	Specifies the Layer 3 I-SID of the DvR route. The range is 0 to 16777215.
<i>nh-as-mac</i>	Specifies the MAC address of the next hop node instead of the system name.
<i>vrf</i>	Specifies the VRF name of the DvR route.
<i>vrfids</i>	Specifies the VRF ID of the DvR route. The range is 0 to 512.

View DvR Database Information

Before You Begin

Ensure that DvR is enabled globally on the node.

About This Task

View all DvR routes on a Controller or a Leaf node.

The Controller node displays all the IP subnet routes configured for that DvR domain. A Leaf node displays all IP subnet routes learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. It also displays the Host Routes (ARPs) learned from other DvR enabled nodes.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the DvR database.

On a Controller:

```
show dvr database [home] [[ipv4 {A.B.C.D}]] [[l3isid<0-16777215>]] [[nh-as-mac]] [[remote]] [[vrf WORD<1-16>]] [[vrfids WORD<0-512>]]
```

On a Leaf node:

```
show dvr database [home] [[ipv4 {A.B.C.D}]] [[l3isid<0-16777215>]] [[nh-as-mac]] [[remote]]
```



Note

Viewing the DvR database for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View the DvR database on either a Controller or a Leaf node.

```
Switch:1#show dvr database
```

DVR DATABASE												
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB	PREFIX COST	COST	AGE	AREA	AREA-NAME
40.0.0.0	255.255.0.0	Ctrl-1:K:121	0	0	40400	cpp	10	1		0 day(s), 05:44:55	HOME	area-49.0016
40.0.1.2	255.255.255.255	Ctrl-1:K:121	0	0	40400	cpp	10	1		0 day(s), 05:44:55	HOME	area-49.0016
40.0.1.3	255.255.255.255	Ctrl-2:K:122	101	0	40400	MLT-1	10	1		0 day(s), 05:44:30	HOME	area-49.0016

3 out of 3 Total Num of DVR Database entries displayed

View the DvR database for a specific IPv4 address:

```
Switch:1#show dvr database ipv4 40.3.1.2
```

DVR DATABASE												
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB	PREFIX COST	COST	AGE	AREA	AREA-NAME
40.3.1.2	255.255.255.255	Ctrl-1:K:121	0	0	40403	cpp	10	1		0 day(s), 05:50:03	HOME	area-49.0016

1 out of 1225 Total Num of DVR Database entries displayed

View DvR database entries for a specific L3 I-SID.

```
Switch:1#show dvr database l3isid 0
```

DVR DATABASE												
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB	PREFIX COST	COST	AGE	AREA	AREA-NAME
40.0.0.0	255.255.0.0	Ctrl-1:K:121	0	0	40400	cpp	10	1		0 day(s), 05:44:55	HOME	area-49.0016
40.0.1.2	255.255.255.255	Ctrl-1:K:121	0	0	40400	cpp	10	1		0 day(s), 05:44:55	HOME	area-49.0016
40.0.1.3	255.255.255.255	Ctrl-2:K:122	0	0	40400	MLT-1	10	1		0 day(s), 05:44:30	HOME	area-49.0016

3 out of 3 Total Num of DVR Database entries displayed

View DvR database entries with next hop MAC address displayed instead of the system name:

```
Switch:1#show dvr database nh-as-mac
```

DVR DATABASE												
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB	PREFIX COST	COST	AGE	AREA	AREA-NAME


```

40.0.0.0      255.255.0.0      00:bb:00:00:81:21      0      0      40400      cpp      10      1      0 day(s), 05:44:55      HOME      area-49.0016
40.0.1.2      255.255.255.255      00:bb:00:00:81:21      0      0      40400      cpp      10      1      0 day(s), 05:44:55      HOME      area-49.0016
40.0.1.3      255.255.255.255      00:bb:00:00:81:22      0      0      40400      MLT-1     10      1      0 day(s), 05:44:30      HOME      area-49.0016

3 out of 3 Total Num of DvR Database entries displayed
-----

```

Variable definitions

Use the data in the following table to use the **show dvr database** command.

Variable	Value
<i>home</i>	Specifies the DvR database information for home instance.
<i>ipv4 {A.B.C.D}</i>	Specifies the IP address (IPv4) of the DvR database entry.
<i>l3isid <0-16777215></i>	Specifies the Layer 3 I-SID of the DvR database entry. The range is 0 to 16777215.
<i>nh-as-mac</i>	Specifies the MAC address of the next hop node instead of the system name.
<i>remote</i>	Specifies the DvR database information for remote instance.
<i>vrf</i>	Specifies the VRF name of the DvR database entry.
<i>vrfids</i>	Specifies the VRF ID of the DvR database entry. The range is 0 to 512.

View DvR Backbone Entries

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the DvR backbone entries (redistributed host routes) learned from all Controllers in all DvR domains.



Note

DvR backbone entries can be viewed only on a Controller. Viewing backbone entries is not applicable on a Leaf node.

Procedure

1. Enter Privileged EXEC mode:
enable

2. View DvR backbone entries:

```
show dvr backbone-entries [adv-controller WORD<1-255>] [[domain-
id <1-255>] | [home] | [host-mac-address 0x00:0x00:0x00:0x00:0x00:0x00] |
[ipv4 {A.B.C.D}] | [l2isid <1-16777215>] | [l3isid <0-16777215>] | [next-hop
WORD<1-255>] | [nh-as-mac] | [remote]
```

Example

View all DvR backbone entries:

```
Switch:1#show dvr backbone-entries
=====
DVR Backbone-Entries
=====
IP-ADDRESS      HOST      L3VSN      L2VSN      DOMAIN
MAC-ADDRESS     ISID      ISID      ID
ADV-CONTROLLER  NEXT HOP   AREA      AREA-NAME
-----
40.0.1.2        b0:ad:aa:4c:55:00  0          40400      255      Ctrl-2:8200:122  Ctrl-1:8400:121  HOME      area-0.00.20
40.0.1.2        b0:ad:aa:4c:55:00  0          40400      255      Ctrl-1:8400:121  Ctrl-1:8400:121  HOME      area-0.00.20
40.0.1.3        b0:ad:aa:43:31:00  0          40400      255      Ctrl-1:8400:121  Ctrl-2:8200:122  HOME      area-0.00.20
40.0.1.3        b0:ad:aa:43:31:00  0          40400      255      Ctrl-2:8200:122  Ctrl-2:8200:122  HOME      area-0.00.20

Home: 4 out of 4 Total Num of DVR Backbone Routes displayed
=====
```

View DvR backbone entries on a specific DvR Controller:

```
Switch:1#show dvr backbone-entries adv-controller Ctrl-2:8200:122
=====
DVR Backbone-Entries
=====
IP-ADDRESS      HOST      L3VSN      L2VSN      DOMAIN
MAC-ADDRESS     ISID      ISID      ID
ADV-CONTROLLER  NEXT HOP   AREA      AREA-NAME
-----
40.0.1.2        b0:ad:aa:4c:55:00  0          40400      255      Ctrl-2:8200:122  Ctrl-1:8400:121  HOME      area-0.00.20
40.1.1.3        b0:ad:aa:43:31:00  0          40401      255      Ctrl-2:8200:122  Ctrl-2:8200:122  HOME      area-0.00.20

Home: 2 out of 2 Total Num of DVR Backbone Routes displayed
=====
```

View DvR backbone entries for a specific host MAC address:

```
Switch:1#show dvr backbone-entries host-mac-address b0:ad:aa:4c:55:00
=====
DVR Backbone-Entries
=====
IP-ADDRESS      HOST      L3VSN      L2VSN      DOMAIN
MAC-ADDRESS     ISID      ISID      ID
ADV-CONTROLLER  NEXT HOP   AREA      AREA-NAME
-----
40.0.1.2        b0:ad:aa:4c:55:00  0          40400      255      Ctrl-2:8200:122  Ctrl-1:8400:121  HOME      area-0.00.20
40.0.1.2        b0:ad:aa:4c:55:00  0          40400      255      Ctrl-1:8400:121  Ctrl-1:8400:121  HOME      area-0.00.20

Home: 2 out of 2 Total Num of DVR Backbone Routes displayed
=====
```

View DvR backbone entries for a specific IP address:

In this example, you enter IP address 40.0.1.0 to display backbone entries for IP addresses 40.0.1.2 and 40.0.1.3.

```
Switch:1#show dvr backbone-entries ipv4 40.0.1.0
=====
DVR Backbone-Entries
=====
IP-ADDRESS      HOST      L3VSN      L2VSN      DOMAIN
MAC-ADDRESS     ISID      ISID      ID
ADV-CONTROLLER  NEXT HOP   AREA      AREA-NAME
-----
```

```

40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-2:8200:122    Ctrl-1:8400:121  HOME    area-0.00.20
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-1:8400:121    Ctrl-1:8400:121  HOME    area-0.00.20
40.1.1.3      b0:ad:aa:43:31:00  0      40401    255      Ctrl-2:8200:122    Ctrl-2:8200:122  HOME    area-0.00.20
40.1.1.3      b0:ad:aa:43:31:00  0      40401    255      Ctrl-2:8200:121    Ctrl-2:8200:122  HOME    area-0.00.20
Home: 4 out of 4 Total Num of DVR Backbone Routes displayed
-----

```

View DvR backbone entries for a specific L3 VSN I-SID:

```
Switch:1#show dvr backbone-entries l3isid 0
```

```

=====
DVR Backbone-Entries
=====
IP-ADDRESS  HOST      L3VSN  L2VSN  DOMAIN  ADV-CONTROLLER  NEXT HOP      AREA  AREA-NAME
MAC-ADDRESS ISID     ISID    ID      ID
-----
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-2:8200:122    Ctrl-1:8400:121  HOME    area-0.00.20
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-1:8400:121    Ctrl-1:8400:121  HOME    area-0.00.20
40.0.1.3      b0:ad:aa:43:31:00  0      40400    255      Ctrl-1:8400:121    Ctrl-2:8200:122  HOME    area-0.00.20
40.0.1.3      b0:ad:aa:43:31:00  0      40400    255      Ctrl-2:8200:122    Ctrl-2:8200:122  HOME    area-0.00.20
Home: 4 out of 4 Total Num of DVR Backbone Routes displayed
-----

```

View DvR backbone entries for a specific next hop node:

```
Switch:1#show dvr backbone-entries next-hop Ctrl-1:8400:121
```

```

=====
DVR Backbone-Entries
=====
IP-ADDRESS  HOST      L3VSN  L2VSN  DOMAIN  ADV-CONTROLLER  NEXT HOP      AREA  AREA-NAME
MAC-ADDRESS ISID     ISID    ID      ID
-----
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-2:8200:122    Ctrl-1:8400:121  HOME    area-0.00.20
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-1:8400:121    Ctrl-1:8400:121  HOME    area-0.00.20
Home: 2 out of 2 Total Num of DVR Backbone Routes displayed
-----

```

View DvR backbone entries where the next hop nodes are displayed as MAC addresses:

```
Switch:1#show dvr backbone-entries nh-as-mac
```

```

=====
DVR Backbone-Entries
=====
IP-ADDRESS  HOST      L3VSN  L2VSN  DOMAIN  ADV-CONTROLLER  NEXT HOP      AREA  AREA-NAME
MAC-ADDRESS ISID     ISID    ID      ID
-----
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-2:8200:122    00:bb:00:00:81:21  HOME    area-0.00.20
40.0.1.2      b0:ad:aa:4c:55:00  0      40400    255      Ctrl-1:8400:121    00:bb:00:00:81:21  HOME    area-0.00.20
Home: 2 out of 2 Total Num of DVR Backbone Routes displayed
-----

```

Variable definitions

Use the data in the following table to use the **show dvr backbone entries** command.

Variable	Value
<i>adv-controller</i> WORD<1-255>	Specifies the system name of the advertising Controller.
<i>domain-id</i> <1-255>	Specifies the domain ID of the DvR backbone entry. The range is 1 to 255.
<i>home</i>	Display the DvR backbone entries for the home instance.

Variable	Value
<i>host-mac-address</i> <i>0x00:0x00:0x00:0x00:0x00:0x00</i>	Specifies the host MAC address of the DvR backbone entry.
<i>ipv4 {A.B.C.D}</i>	Specifies the IP address (IPv4) of the DvR backbone entry.
<i>l2isid <1-16777215></i>	Specifies the Layer 2 I-SID of the DvR backbone entry. The range is 1 to 16777215.
<i>l3isid <0-16777215></i>	Specifies the Layer 3 I-SID of the DvR backbone entry. The range is 0 to 16777215.
<i>next-hop WORD<1-255></i>	Specifies the system name of the next hop node.
<i>nh-as-mac</i>	Specifies the MAC address of the next hop node instead of the system name.
<i>remote</i>	Display the DvR backbone entries for the remote instance.

View DvR Backbone Members

Before You Begin

Ensure that DvR is enabled globally on the node.

About This Task

DvR backbone members are either DvR Controllers or non-DvR BEBs that receive redistributed host routes from all other DvR Controllers in the SPB network.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View DvR backbone member information:
`show dvr backbone-members [controller | home | non-dvr-beb | remote]`

Example

View all DvR backbone members:

```
Switch:1#show dvr backbone-members
```

```
=====
                        DVR BB Members
=====
System Name      Nick-Name      Nodal MAC      Role      Domain Id  Area      Area-Name
-----
DVR-D2-C1-40     0.82.40       00:00:82:84:40:00  NON-DVR-BEB  2          HOME      area-0.00.20
Ctrl-2:8200:122  0.81.22       00:bb:00:00:81:22  Controller   2          HOME      area-0.00.20

Home: 2 out of 2 Total Num of DVR Backbone Members displayed
=====
```

View backbone members that are DvR controllers:

```
Switch:1#show dvr backbone-members controller

=====
DVR BB Members (Domain ID: 255)
=====
System Name      Nick-Name      Nodal MAC      Role      Domain Id      Area      Area-Name
-----
Ctrl-2:8200:122  0.81.22       00:bb:00:00:81:22  Controller  2             HOME      area-0.00.20

Home: 1 out of 2 Total Num of DVR Backbone Members displayed
=====
```

View backbone members that are non-DvR BEBs:

```
Switch:1#show dvr backbone-members non-dvr-beb

=====
DVR BB Members
=====
System Name      Nick-Name      Nodal MAC      Role      Domain Id      Area      Area-Name
-----
DVR-D2-C1-40     0.82.40       00:00:82:84:40:00  NON-DVR-BEB  2             HOME      area-0.00.20

Home: 1 out of 2 Total Num of DVR Backbone Members displayed
=====
```

Variable definitions

Use the data in the following table to use the **show dvr backbone-members** command.

Variable	Value
<i>controller</i>	Specifies backbone members that are DvR Controllers.
<i>home</i>	Specifies DvR backbone members information for the home instance.
<i>non-dvr-beb</i>	Specifies backbone members that are non-DvR BEBs.
<i>remote</i>	Specifies DvR backbone members information for the remote instance.

View Layer 3 VSN Information

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View VRFs corresponding to Layer 3 (routed) VSN I-SIDs on either a Controller or a Leaf node.

Procedure

1. Enter Privileged EXEC mode:
enable

2. View the Layer 3 VSN information:

```
show dvr l3vsns [l3isid <0-16777215>] | [vrf WORD<1-16>] | [vrfids  
WORD<0-512>]
```

Example

View Layer 3 VSN information on a DvR Controller:

```
Switch:1#show dvr l3vsns

=====
DVR L3VSN
=====
VRF ID          L3VSN ISID      VRF NAME      INJECT-DEFAULT-ROUTE-DISABLE
-----
1                55500           vrf600        Disabled
2                55501           vrf601        Disabled
3                55502           vrf602        Disabled
4                55503           vrf603        Disabled

4 out of 4 Total Num of DVR L3VSN displayed
=====
```

View Layer 3 VSN information on a DvR Leaf node:

```
Switch2:1#show dvr l3vsns

=====
DVR L3VSN
=====
VRF ID          L3VSN ISID      VRF NAME
-----
1                55500           vrf600
2                55501           vrf601
3                55502           vrf602

3 out of 3 Total Num of DVR L3VSN displayed
=====
```

Variable definitions

Use the data in the following table to use the **show dvr l3vsns** command.

Variable	Value
<i>l3isid</i> <0-16777215>	Specifies the Layer 3 VSN I-SID. The range is 0 to 16777215.
<i>vrf</i> WORD<1-16>	Specifies the VRF name of the VRF corresponding to the Layer 3 VSN I-SID.
<i>vrfids</i> WORD<0-512>	Specifies the VRF ID of the VRF.

View DvR Domain Redistribution Information

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View DvR domain redistribution information on a Controller or a Leaf node.



Note

You can view DvR domain redistribution information only on a DvR Controller. An error message displays if you attempt to view this information on a DvR Leaf node.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View DvR domain redistribution information:
show dvr redistribute [vrf *WORD<1-16>*] | [vrfids *WORD<0-512>*]

Example

View DvR domain redistribution information on a Controller:

```
Switch:1#show dvr redistribute
=====
DVR Redistribute List - GlobalRouter
=====
SOURCE MET MTYPE      ENABLE RPOLICY
-----
STAT   1   External   TRUE   -
```

View DvR domain redistribution information for a particular VRF.

```
Switch:1#show dvr redistribute vrf vrf1
=====
DVR Redistribute List - VRF vrf1
=====
SOURCE MET MTYPE      ENABLE RPOLICY
-----
STAT   20000 External   TRUE   -
LOC    10000 Internal   TRUE   -
```


Variable definitions

Use the data in the following table to use the **show dvr redistribute** command.

Variable	Definitions
<i>vrf WORD<1-16></i>	Specifies the VRF name.
<i>vrfids WORD<0-512></i>	Specifies the VRF ID of the VRF.

Configure a DvR Solution

The following section describes a simple configuration example to configure Distributed Virtual Routing (DvR) over a Fabric Connect (SPB) network.



Before You Begin

On the switches to be configured as DvR Controllers:

- Ensure that you configure Fabric Connect.
- Ensure that you configure IP Shortcuts on the node. This is necessary for proper functioning of the node as a DvR Controller.
- Verify that the `dvr-leaf-mode` boot flag is disabled on the node. To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

On the switches to be configured as DvR Leaf nodes:

- Ensure that you configure Fabric Connect.

About This Task

In this example, you configure two DvR Controllers (with IP addresses 10.133.226.101 and 10.133.226.102) and two DvR Leaf nodes (with IP addresses 10.133.226.103 and 10.133.226.104), in a single DvR domain with domain ID 9. Hosts connect to the DvR nodes as shown in the figure.

Procedure

DvR Controller configuration — Controller 1 and Controller 2:

1. Verify configuration of Fabric Connect on each of the switches to be configured as the DvR Controllers.

The following examples show verification on one of the switches. Perform this verification on both switches.

- a. Verify the SPB configuration:

```
Switch1:1>en
Switch1:1#show spbm
```



```

        spbm : enable
        ethertype : 0x8100
        nick-name server : enable
        nick-name allocation : static
        nick-name server range : B.00.00-B.FF.FF

```

```
Switch1:1#show isis spbm
```

```

=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY      NICK      LSDB      IP      IPV6      MULTICAST  SPB-PIM-GW  STP-MULTI  ORIGIN
INSTANCE  VLAN            NAME        TRAP
-----
1          4051-4052  4051        0.10.01  disable  enable  disable  enable     disable    disable    dynamic
=====

                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
INSTANCE
-----
1          primary                00:00:00:00:00:00                N/A                    N/A
=====

Total Num of SPBM instances: 1
=====

```

b. Verify the global IS-IS configuration:

```
Switch1:1#show isis
```

```

=====
                        ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID : 00bb.0000.0101
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name : Cont-1
ip source-address : 10.0.0.101
ipv6 source-address :
ip tunnel source-address :
Tunnel vrf :
ip tunnel mtu :
Num of Interfaces : 4
Num of Area Addresses : 1
Inband Mgmt Clip IP :
backbone : disabled
Dynamically Learned Area : 00.0000.0000
FAN Member : No
Multi-Area OperState : disabled
Hello Padding : enabled
Multi-Area OperState : disabled
Multi-Area Flags :

```

2. Configure the DvR Controllers.

a. Configure Controller 1 (IP address 10.133.226.101) with DvR domain ID 9.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#dvr controller 9
Switch:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 9
Domain ISID         : 16678219
Backbone ISID       :
Role                : Controller
My SYS ID           : 00:bb:00:00:81:21
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs  : Disabled
VRRP Priority        : 100
VRRP Election       :
```

b. Configure Controller 2 (IP address 10.133.226.102), also with DvR domain ID 9.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#dvr controller 9
Switch:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 9
Domain ISID         : 16678219
Backbone ISID       : 16678216
Role                : Controller
My SYS ID           : 00:bb:00:00:81:21
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
VRRP on DVR I-SIDs  : Disabled
VRRP Priority        : 100
VRRP Election       :
```

c. Verify the configuration. View the members of the DvR domain.

```
Switch1:1#show dvr members

=====
DVR Members (Domain ID: 2)
=====
System Name          Nick-Name          Nodal MAC          Role
-----
Cont-1                0.10.01            00:bb:00:00:01:01   Controller
Cont-2                0.10.02            00:bb:00:00:01:02   Controller

2 out of 2 Total Num of DVR Members displayed
=====
```

Layer 2 VSN (VLAN) configuration on the DvR Controllers:

3. Configure Layer 2 VSN on the DVR Controllers, Controller 1 and Controller 2.

a. Configure platform VLANs on Controller 1 (VLAN ID=200 and VLAN ID=202).

Associate the VLANs with the I-SIDs 20200 and 20202 respectively. Configure gateway IPv4 addresses 20.0.1.1 and 20.2.1.1 respectively, and enable DVR on those interfaces.

```
Switch1:1(config)#vlan create 200 type port-mstprstp 0
Switch1:1(config)#vlan i-sid 200 20200
Switch1:1(config)#interface vlan 200
Switch1:1(config)#dvr gw-ipv4 20.0.1.1
Switch1:1(config)#dvr enable
Switch1:1(config)#ip address 20.0.1.2 255.255.0.0

Switch1:1(config)#vlan create 202 type port-mstprstp 0
Switch1:1(config)#vlan i-sid 202 20202
Switch1:1(config)#interface vlan 202
Switch1:1(config)#dvr gw-ipv4 20.2.1.1
Switch1:1(config)#dvr enable
Switch1:1(config)#ip address 20.2.1.2 255.255.0.0
Switch1:1(config)#exit
Switch1:1#
```

b. Configure the platform VLANs on Controller 2. Ensure that you configure the same gateway IPv4 addresses on the corresponding VLANs, as on Controller 1.

```
Switch2:1(config)#vlan create 200 type port-mstprstp 0
Switch2:1(config)#vlan i-sid 200 20200
Switch2:1(config)#interface vlan 200
Switch2:1(config)#dvr gw-ipv4 20.0.1.1
Switch2:1(config)#dvr enable
Switch2:1(config)#ip address 20.0.1.3 255.255.0.0

Switch2:1(config)#vlan create 202 type port-mstprstp 0
Switch2:1(config)#vlan i-sid 202 20202
Switch2:1(config)#interface vlan 202
Switch2:1(config)#dvr gw-ipv4 20.2.1.1
Switch2:1(config)#dvr enable
Switch2:1(config)#ip address 20.2.1.3 255.255.0.0
Switch2:1(config)#exit
Switch2:1#
```

c. Verify Layer 2 VSN (VLAN) configuration on the Controllers. The following example shows the verification on Controller 1. Perform this verification on both Controllers. View the DVR interfaces.

On Controllers, DVR interfaces are created when you configure IP on a DVR enabled Layer 2 VSN (VLAN, I-SID). You can also view the administrative state of these interfaces on the Controller.

```
Switch1:1#show dvr interfaces

=====
DVR Interfaces
=====
```

Interface	Mask	L3ISID	VRPID	L2ISID	VLAN	GW IPv4	Admin State	SPBMC State	IGMP Version
20.0.1.2	255.255.0.0	0	0	20200	200	20.0.1.1	enable	disable	2
20.2.1.2	255.255.0.0	0	0	20202	202	20.2.1.1	enable	disable	2

```
2 out of 2 Total Num of DVR Interfaces displayed
=====
```

View the DvR host entries learned locally on the S-UNI port.

```
Switch1:1#show dvr host-entries
```

DVR Host-Entries							
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	PORT	DOMAIN ID	TYPE	NEXT HOP
20.0.1.2	b0:ad:aa:42:ed:04	0	20200	cpp	9	LOCAL	Cont-1
20.2.1.2	b0:ad:aa:42:ed:04	0	20202	cpp	9	LOCAL	Cont-1
2 out of 2 Total Num of DVR Host Entries displayed							

View the DvR database. All IP subnet routes configured on the Controller, for the DvR domain, are displayed.

```
Switch1:1#show dvr database
```

DVR DATABASE								
DEST	MASK	NEXT HOP	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB COST	PREFIX COST	AGE
20.0.1.2	255.255.255.255	Cont-1	0	20200	cpp	10	1	1 day(s), 06:41:40
20.2.1.2	255.255.255.255	Cont-1	0	20202	cpp	10	1	1 day(s), 06:41:40
2 out of 2 Total Num of DVR Database entries displayed								

View the DvR routes for the subnets 20.0.0.0 and 20.2.0.0.

```
Switch1:1#show dvr routes
```

DVR Routes						
DEST	MASK	NEXT HOP	L3VSN ISID	L2VSN ISID	TYPE	COST
20.0.0.0	255.255.0.0	Cont-1	0	20200	-	1
20.2.0.0	255.255.0.0	Cont-1	0	20202	-	1
2 out of 2 Total Num of DVR Routes displayed						

TYPE Legend: E=Ecmp Route

Layer 3 configuration on the DvR Controllers

4. Configure Layer 3 (VRF) on the DvR Controllers, Controller 1 and Controller 2.
 - a. Configure Layer 3 on Controller 1. As part of this configuration, you configure a VRF `vrf501` and associate it with a DvR VLAN.

```
Switch1:1(config)#ip vrf vrf501 vrfid 501
Switch1:1(config)#vlan create 501 type port-mstprstp 0
Switch1:1(config)#vlan i-sid 501 50501
Switch1:1(config)#interface Vlan 501
Switch1:1(config)#vrf vrf501
Switch1:1(config)#dvr gw-ipv4 50.1.1.1
```

```
Switch1:1(config)#dvr enable
Switch1:1(config)#ip address 50.1.1.2 255.255.0.0

Switch1:1(config)#router vrf vrf501
Switch1:1(router-vrf)#i-sid 55501
Switch1:1(router-vrf)#ipvpn enable
Switch1:1(router-vrf)#exit
Switch1:1(config)#
```

b. Configure Layer 3 on Controller 2.

```
Switch2:1(config)#ip vrf vrf501 vrfid 501
Switch2:1(config)#vlan create 501 type port-mstprstp 0
Switch2:1(config)#vlan i-sid 501 50501
Switch2:1(config)#interface Vlan 501
Switch2:1(config)#vrf vrf501
Switch2:1(config)#dvr gw-ipv4 50.1.1.1
Switch2:1(config)#dvr enable
Switch2:1(config)#ip address 50.1.1.3 255.255.0.0

Switch2:1(config)#router vrf vrf501
Switch2:1(router-vrf)#i-sid 55501
Switch2:1(router-vrf)#ipvpn enable
Switch2:1(router-vrf)#exit
Switch2:1(config)#
```

c. Verify Layer 3 configuration. The following example shows verification on Controller 1. Perform this verification on both Controllers.

View the DVR host entries.

```
Switch2:1(config)#show dvr host-entries l3isid 55501
```

```
=====
                                DVR Host-Entries
=====
```

IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	PORT	DOMAIN ID	TYPE	NEXT HOP
50.1.1.2	b0:ad:aa:42:ed:08	55501	50501	cpp	9	LOCAL	Cont-1
50.1.1.3	b0:ad:aa:4c:3d:02	55501	50501	1/23	9	DYNAMIC	Cont-2

```

2 out of 3267 Total Num of DVR Host Entries displayed
=====
```

View the DVR interfaces.

```
Switch2:1(config)#show dvr interfaces l3isid 55501
```

```
=====
                                DVR Interfaces
=====
```

Interface	Mask	L3ISID	VRFID	L2ISID	VLAN	GW IPv4	Admin State	SPBMC State	IGMP Version
50.1.1.2	255.255.0.0	55501	501	50501	501	50.1.1.1	enable	disable	2

```

1 out of 291 Total Num of DVR Interfaces displayed

```

```
Switch2:1(config)#show dvr database l3isid 55501
```

```
=====
                                DVR DATABASE
=====
```

DEST	MASK	NEXT HOP	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB COST	PREFIX COST	AGE
50.1.0.0	255.255.0.0	Cont-1	55501	50501	cpp	10	1	0 day(s), 01:26:49

```

50.1.1.2      255.255.255.255 Cont-1      55501      50501      cpp          10   1      0 day(s), 01:26:49
50.1.1.3      255.255.255.255 Cont-2      55501      50501      1/23         10   1      0 day(s), 01:24:53

3 out of 3558 Total Num of DVR Database entries displayed
-----

```

DvR Leaf configuration — Leaf 1 and Leaf 2

- Configure the boot flag `dvr-leaf-mode` on the switches to be configured as DvR Leaf nodes.



Caution

Ensure that you save the current configuration on the switch, before you enable the flag.

Enabling the flag removes all existing non-DvR configuration on the switch, such as platform VLANs and their IP address configuration, CLIP configuration, routing protocol configuration and VRF configuration. The gateway IPv4 address, if configured, is also removed.

On switch with IP address 10.133.226.104, configure the boot flag and reboot the switch.

```

Switch3:1>en
Switch3:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch3:1(config)#boot config flags dvr-leaf-mode
Switch3:1(config)#save config
Switch3:1(config)#reset

```

On switch with IP address 10.133.226.105, configure the boot flag and reboot the switch.

```

Switch4:1>en
Switch4:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch4:1(config)#boot config flags dvr-leaf-mode
Switch4:1(config)#save config
Switch4:1(config)#reset

```

- After the switches come back up, configure the nodes as DvR Leaf nodes. Configure switch with IP address 10.133.226.104 as DvR Leaf 1; verify the configuration.

```

Switch3:1(config)#dvr Leaf 9
Switch3:1(config)#show dvr

=====
DVR Summary Info
=====
Domain ID           : 9
Domain ISID         : 16678219
Role                : Leaf
My SYS ID           : 00:bb:00:00:80:05
Operational State   : Up
GW MAC              : 00:00:5e:00:01:25
Inband Mgmt Clip IP :
Virtual Ist local address :
Virtual Ist local subnet mask :
Virtual Ist peer address :
Virtual Ist cluster-id :
Virtual Ist ISID     :

```

Configure switch with IP address 10.133.226.105 as DvR Leaf 2; verify the configuration.

```
Switch4:1(config)#dvr Leaf 9
Switch4:1(config)#show dvr
```

```
=====
                        DVR Summary Info
=====
```

Domain ID	: 9
Domain ISID	: 16678219
Role	: Leaf
My SYS ID	: 00:bb:00:00:80:05
Operational State	: Up
GW MAC	: 00:00:5e:00:01:25
Inband Mgmt Clip IP	:
Virtual Ist local address	:
Virtual Ist local subnet mask	:
Virtual Ist peer address	:
Virtual Ist cluster-id	:
Virtual Ist ISID	:

7. Associate the I-SIDs on the DvR Leaf nodes to the DvR VLANs configured on the Controller.

On Leaf node 1 (IP address 10.133.226.105):

```
Switch3:1(config)#i-sid 20200 elan
Switch3:1(elan:20200)#c-vid 200 port 1/2
Switch3:1(config)#exit
```

View the host connections.

```
Switch3:1#show dvr host-entries nh-as-mac
```

```
=====
                        DVR Host-Entries
=====
```

IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	PORT	DOMAIN ID	TYPE	NEXT HOP
20.0.1.67	00:00:00:00:00:67	0	20200	1/4	9	DYNAMIC	00:bb:00:00:81:21
20.0.1.68	00:00:00:00:00:68	0	20200	1/2	9	DYNAMIC	00:bb:00:00:81:21

2 out of 2 Total Num of DVR Host Entries displayed

On Leaf node 2 (IP address 10.133.226.105):

```
Switch4:1(config)#i-sid 20200 elan
Switch4:1(elan:20200)#c-vid 200 port 1/2
Switch4:1(config)#exit
```

View the host connections.

```
Switch4:1#show dvr host-entries nh-as-mac
```

```
=====
                        DVR Host-Entries
=====
```

IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	PORT	DOMAIN ID	TYPE	NEXT HOP
20.0.1.67	00:00:00:00:00:67	0	20200	1/4	9	DYNAMIC	00:bb:00:00:81:21
20.0.1.68	00:00:00:00:00:68	0	20200	1/2	9	DYNAMIC	00:bb:00:00:81:21

```
2 out of 2 Total Num of DVR Host Entries displayed
```

8. View all members of the DvR domain. You can view this information on either a Leaf node or a Controller node.

```
Switch1:1#show dvr members
```

```
=====
DVR Members (Domain ID: 2)
=====
```

System Name	Nick-Name	Nodal MAC	Role
Cont-1	0.10.01	00:bb:00:00:01:01	Controller
Cont-2	0.10.02	00:bb:00:00:01:02	Controller
Leaf1	0.10.04	00:bb:00:00:80:04	Leaf
Leaf2	0.10.05	00:bb:00:00:80:05	Leaf

```
4 out of 4 Total Num of DVR Members displayed
=====
```

DvR Configuration Using EDM

The following sections describe configuration of Distributed Virtual Routing (DvR) using EDM.

Configure a DvR Controller or a DvR Leaf Globally

Before You Begin



Important

For DvR Leaf Configuration only:

You must enable the `dvr-leaf-mode` boot flag before you configure a node as a DvR Leaf node. Navigate to **Configuration > Edit > Chassis**. On the **Boot Config** tab, select **EnableDvrLeafMode**.

Ensure that you save the current configuration on the switch, before you enable the flag. Enabling the flag removes all non-DvR configuration on the switch.

About This Task

Configure a node to perform the role of either a Controller or a Leaf, within the DvR domain.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Globals** tab.
4. In **DomainId**, type the domain ID.



Note

A Controller or a Leaf node can belong to only one DvR domain.

5. In **Role**, select the role of the node.
6. (Optional) On a Controller node, disable injection of default routes into the DvR domain. Select **InjectDefaultRouteDisable**.

**Note**

This parameter applies only to Controllers. Attempting to select this field on a Leaf node displays an error message.

7. Update the parameters as necessary, and then select **Apply**.

Globals *Field Descriptions*

Use the data in the following table to use the **Globals** tab.

Field	Descriptions
DomainId	Identifies the domain that the node belongs to. The range for a Controller or a Leaf is 1 to 255. Set to 0 if is not configured.
Role	Specifies the role of the node in the domain as one of the following: <ul style="list-style-type: none"> • controller • leaf • isolatedController
Enable	Specifies whether Distributed Virtual Routing (DvR) is enabled on the node. Configuring a Controller or Leaf sets this parameter to true.
DomainIsid	Uniquely identifies the domain I-SID that the node belongs to. 0 indicates that is not configured.
BackboneIsid	Uniquely identifies the backbone I-SID that the node belongs to. The valid backbone I-SID is 16678216; it is configured to 0 if is not configured.
GatewayMac	Specifies the Gateway MAC address used by all domains.
InbandMgmtIp	Specifies the In-band Management IP address configured under IS-IS. You can use this IP address to manage the node, irrespective of whether DvR is enabled on it.

Field	Descriptions
InjectDefaultRouteDisable	Specifies whether injection of default routes is disabled on the Controller in the domain. By default, Controllers inject default routes into the domain so that all Leaf nodes in the domain learn these routes with the next hop as the Controller that advertised it. Select this parameter to disable this behavior.
VirtualIstLocalAddr	Specifies the local IP address of vIST, if vIST is configured on a Leaf. You cannot configure vIST cannot on a Controller.
VirtualIstLocalMask	Specifies the local subnet mask of vIST, if vIST is configured on a Leaf. You cannot configure vIST on a Controller.
VirtualIstPeerAddr	Specifies the peer IP address of vIST, if vIST is configured on a Leaf. You cannot configure vIST on a Controller.
VirtualIstClusterId	Specifies the cluster ID of vIST, if vIST is configured on a Leaf. You cannot configure vIST on a Controller. Set to 0 if vIST is not configured.
VirtualIstIsid	Specifies the I-SID if vIST is configured.
OperState	Specifies the operational state of the node.
VrrpEnable Note: Exception: not supported on 5320 Series or 5420 Series.	Enables Virtual Router Redundancy Protocol (VRRP) advertisements on DvR VLAN Service Instance Identifiers (I-SIDs). The default value is disabled.
VrrpPriority Note: Exception: not supported on 5320 Series or 5420 Series.	Configures the VRRP priority for the elected VLAN and associated I-SID. Switches advertise their priority to use in the primary router election process. When multiple Controllers elect the same DvR VLAN and I-SID, the Controller with the higher priority becomes the primary router. The elected Controller sends VRRP advertisements on all DvR VLANs. The default value is 100.
VrrpElectionVlanId Note: Exception: not supported on 5320 Series or 5420 Series.	Displays the election VLAN for VRRP on DvR I-SIDs. A value of 4095 indicates no elected VLAN exists.

Configure DvR-VRRP Coexistence



Note

This procedure does not apply to 5320 Series or 5420 Series.

When a Layer 2 Virtual Service Network (VSN) extends outside of the Distributed Virtual Routing (DvR) domain, DvR Controllers can route traffic that enters on the network-to-network interface (NNI) by using Virtual Router Redundancy Protocol (VRRP).

Before You Begin

The VLAN must be a DvR VLAN. A DvR VLAN is a VLAN configured on a DvR Controller with a VLAN IP address, a VLAN I-SID, the DvR gateway IP address, and DvR enabled.

About This Task

If you enable this functionality, the switch automatically creates a VRRPv2 interface with the virtual router ID (VRID) 37 and a *holddown-timer* configuration of 60 on the DvR VLAN that you identify for VRRP election. You can configure one VLAN for election.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. For **DvrVrrpElection**, select **true** to select the election VLAN for DvR VRRP.
5. Select **Apply**.
6. In the navigation pane, expand **Configuration > Fabric**.
7. Select **DVR**.
8. Select the **Globals** tab.
9. (Optional) For **VrrpPriority**, configure the DvR VRRP priority for the elected VLAN.
10. Select **VrrpEnable** to enable VRRP on all DvR Service Instance Identifiers (I-SIDs).
11. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Field	Descriptions
DomainId	Identifies the domain that the node belongs to. The range for a Controller or a Leaf is 1 to 255. Set to 0 if is not configured.
Role	Specifies the role of the node in the domain as one of the following: <ul style="list-style-type: none">• controller• leaf• isolatedController

Field	Descriptions
Enable	Specifies whether Distributed Virtual Routing (DvR) is enabled on the node. Configuring a Controller or Leaf sets this parameter to true.
DomainIsid	Uniquely identifies the domain I-SID that the node belongs to. 0 indicates that is not configured.
Backbonelsid	Uniquely identifies the backbone I-SID that the node belongs to. The valid backbone I-SID is 16678216; it is configured to 0 if is not configured.
GatewayMac	Specifies the Gateway MAC address used by all domains.
InbandMgmtIp	Specifies the In-band Management IP address configured under IS-IS. You can use this IP address to manage the node, irrespective of whether DvR is enabled on it.
InjectDefaultRouteDisable	Specifies whether injection of default routes is disabled on the Controller in the domain. By default, Controllers inject default routes into the domain so that all Leaf nodes in the domain learn these routes with the next hop as the Controller that advertised it. Select this parameter to disable this behavior.
VirtualIstLocalAddr	Specifies the local IP address of vIST, if vIST is configured on a Leaf. You cannot configure vIST cannot on a Controller.
VirtualIstLocalMask	Specifies the local subnet mask of vIST, if vIST is configured on a Leaf. You cannot configure vIST on a Controller.
VirtualIstPeerAddr	Specifies the peer IP address of vIST, if vIST is configured on a Leaf. You cannot configure vIST on a Controller.
VirtualIstClusterId	Specifies the cluster ID of vIST, if vIST is configured on a Leaf. You cannot configure vIST on a Controller. Set to 0 if vIST is not configured.
VirtualIstIsid	Specifies the I-SID if vIST is configured.
OperState	Specifies the operational state of the node.

Field	Descriptions
VrrpEnable Note: Exception: not supported on 5320 Series or 5420 Series.	Enables Virtual Router Redundancy Protocol (VRRP) advertisements on DvR VLAN Service Instance Identifiers (I-SIDs). The default value is disabled.
VrrpPriority Note: Exception: not supported on 5320 Series or 5420 Series.	Configures the VRRP priority for the elected VLAN and associated I-SID. Switches advertise their priority to use in the primary router election process. When multiple Controllers elect the same DvR VLAN and I-SID, the Controller with the higher priority becomes the primary router. The elected Controller sends VRRP advertisements on all DvR VLANs. The default value is 100.
VrrpElectionVlanId Note: Exception: not supported on 5320 Series or 5420 Series.	Displays the election VLAN for VRRP on DvR I-SIDs. A value of 4095 indicates no elected VLAN exists.

View DvR Routes

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the DvR routes (host routes and the IPv4 network routes) that are learned on a DvR Controller or a Leaf node.

Controllers display all the IP subnet routes configured for that DvR domain. Leaf nodes display the IP subnet routes learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. Leaf nodes also display any redistributed routes into the DvR Domain that are learned from the Controllers (direct routes, static routes and the default route).

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Routes** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Routes *Field Descriptions*

Use the data in the following table to use the **Routes** tab.

Name	Description
DestIpAddrType	Specifies the IPv4 destination address type of the DvR route.
DestIpAddr	Specifies the IPv4 destination address of the DvR route.
DestMask	Specifies the destination mask of the DvR route.
L3Isid	Specifies the Layer 3 I-SID of the DvR route.
EcmpIndex	Specifies the ECMP index for the ECMP routes of the DvR route.
NextHopMac	Specifies the MAC address of the next hop BEB in the DvR route.
L2Isid	Specifies the Layer 2 I-SID of the DvR route.
VrfId	Specifies the VRF ID.
Cost	Specifies the SPB cost of the DvR route.
NextHopName	Specifies the host name of the next hop BEB, in the DvR route.
Type	Specifies the route type of the DvR route.

View Members of a DvR Domain

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the members of all DvR domains namely the Controllers and Leaf nodes.

You can view this information on either a Controller or a Leaf node. Both the Controller and the Leaf node displays the members of the DvR domain to which it belongs.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Members** tab.
4. (Optional) To filter the rows based on specific criteria, select **Filter**.

Members *Field Descriptions*

Use the data in the following table to use the **Members** tab.

Name	Description
MacAddress	Specifies the system ID or the nodal MAC address of this DvR member.
SysName	Specifies the system name of this DvR member.
NickName	Specifies the nickname of this DvR member.
Role	Specifies the DvR role (Controller or Leaf) of this DvR member.
DomainId	Specifies the domain ID of the DvR domain that this member belongs to.

View DvR Backbone Members

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

DvR backbone members are either DvR Controllers or non-DvR BEBs that receive redistributed host routes from all other DvR Controllers in the SPB network.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Backbone Members** tab.
4. (Optional) To filter the rows based on specific criteria, select **Filter**.

Backbone Members *Field Descriptions*

Use the data in the following table to use the **Backbone Members** tab.

Name	Description
MacAddress	Specifies the system ID or the nodal MAC address of this DvR backbone member.
SysName	Specifies the system name of this DvR backbone member.
NickName	Specifies the nickname of this DvR backbone member.

Name	Description
Role	Specifies the role of this DvR backbone member. It is either a DvR Controller or a non-DvR BEB.
DomainId	Specifies the domain ID of the DvR domain that this backbone member belongs to. The domain ID is 0 for a non-DvR BEB.

View Remote DvR Backbone Members



Note

This procedure does not apply to 5320 Series and 5420 Series.

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

Perform this procedure to view the remote DvR backbone members. DvR backbone members are either DvR Controllers or non-DvR BEBs that receive redistributed host routes from all other DvR Controllers in the SPB network.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR Remote**.
3. Select the **Backbone Members** tab.
4. (Optional) To filter the rows based on specific criteria, select **Filter**.

Backbone Members *Field Descriptions*

Use the data in the following table to use the **Backbone Members** tab.

Name	Description
MacAddress	Specifies the system ID or the nodal MAC address of this DvR backbone member.
SysName	Specifies the system name of this DvR backbone member.
NickName	Specifies the nickname of this DvR backbone member.

Name	Description
Role	Specifies the role of this DvR backbone member. It is either a DvR Controller or a non-DvR BEB.
DomainId	Specifies the domain ID of the DvR domain that this backbone member belongs to. The domain ID is 0 for a non-DvR BEB.

View DvR Interfaces

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the DvR interfaces on either a Controller or a Leaf node.

On Controllers, DvR interfaces are created when you configure IP on a DvR enabled Layer 2 VSN (VLAN, I-SID). Only Controllers display the administrative state of the interfaces because this is where you enable or disable the interfaces. On a Leaf node, the DvR interface information that the Controllers push, for example, subnet routes and the gateway IP addresses for the Layer 2 VSNs, are displayed.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Interfaces** tab.
4. (Optional) To filter rows based on specific filter criteria, select **Filter**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
VlanIpAddrType	Specifies the VLAN IP address type of the DvR interface.
VlanIpAddr	Specifies the VLAN IP address (IPv4) of the DvR interface.
L3Isid	Specifies the Layer 3 I-SID of the DvR interface. The range is 1 to 16777215.
L2Isid	Specifies the Layer 2 I-SID of the DvR interface. The range is 1 to 16777215.

Name	Description
VlanIpMask	Specifies the VLAN IP address mask of the DvR interface.
VrfId	Specifies the VRF ID of the DvR interface. The VRF ID is 0 for the GRT.
VlanId	Specifies the VLAN ID of the DvR interface.
GwIpAddrType	Specifies the address type of the DvR gateway IP address (IPv4).
GwIpAddr	Specifies the DvR gateway IP address (IPv4).
AdminState	Specifies the administrative state of the DvR interface.
SpbmState	Specifies the state of IP Multicast over Fabric Connect, on the DvR interface.
IgmpVersion	Specifies the version of ICMP that runs on the DvR interface.

View DvR Host Entries

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View DvR host entries (IPv4 remote ARPs) on either a Controller or a Leaf node. The node displays the host entries learned either locally on its UNI port or dynamically from other nodes in the DvR domain.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Host Entries** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Host Entries Field Descriptions

Use the data in the following table to use the **Host Entries** tab.

Name	Description
IpAddrType	Specifies the address type of the DvR host entry (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR host entry.
Mask	Specifies the subnet mask of the DvR host entry.

Name	Description
L3lsid	Specifies the Layer 3 I-SID of the DvR host entry.
MacAddr	Specifies the MAC address of the DvR host entry.
L2lsid	Specifies the Layer 2 I-SID of the DvR host entry.
Vrfid	Specifies the VRF ID associated with the DvR host entry.
Port	Specifies the port of the DvR host entry.
DomainId	Specifies the DvR domain ID of the DvR host entry.
Type	Specifies the host type of the DvR host entry.
NextHopName	Specifies the next hop system name of the DvR host entry.
NextHopMac	Specifies the next hop system MAC address of the DvR host entry.
ClearEntry	Clears the entry if the configured value is true.

View Remote DvR Host Entries



Note

This procedure does not apply to 5320 Series and 5420 Series.

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

Perform this procedure to view the remote DvR host entries (IPv4 remote ARPs) on either a Controller or a Leaf node. The node displays the host entries learned either locally on its UNI port or dynamically from other nodes in the DvR domain.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR Remote**.
3. Select the **Host Entries** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Host Entries Field Descriptions

Use the data in the following table to use the **Host Entries** tab.

Name	Description
IpAddrType	Specifies the address type of the DvR host entry (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR host entry.
Mask	Specifies the subnet mask of the DvR host entry.
L3Isid	Specifies the Layer 3 I-SID of the DvR host entry.
MacAddr	Specifies the MAC address of the DvR host entry.
L2Isid	Specifies the Layer 2 I-SID of the DvR host entry.
VrfId	Specifies the VRF ID associated with the DvR host entry.
Port	Specifies the port of the DvR host entry.
DomainId	Specifies the DvR domain ID of the DvR host entry.
Type	Specifies the host type of the DvR host entry.
NextHopName	Specifies the next hop system name of the DvR host entry.
NextHopMac	Specifies the next hop system MAC address of the DvR host entry.
ClearEntry	Clears the entry if the configured value is true.

Clear DvR Host Entries

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

Clear DvR host entries (IPv4 remote host routes) on a Controller. The host entries are learned on the switch either locally on its UNI port or dynamically from other nodes in the DvR domain.



Note

You can clear DvR host entries only on a DvR Controller.

An error message displays if you attempt clearing of host entries on a DvR Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Clear Host Entries** tab.
4. Update the fields as necessary, and then select **Apply** to save your configuration.

Clear Host Entries Field Descriptions

Use the data in the following table to use the **Clear Host Entries** tab.

Name	Description
ClearAll	Select to clear all DvR host entries.
ClearIpv4	Specifies the IPv4 address of the DvR host entries to clear. The IPv4 address must not be the VLAN IP address on any Controller within the DvR domain.
ClearL2Isid	Specifies the Layer 2 VSN I-SID of the DvR host entries to clear. The range is 0 to 16777215.
ClearL3Isid	Specifies the Layer 3 VSN I-SID of the DvR host entries to clear. The range is 0 to 16777215.

View Layer 3 VSN Information

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View VRFs corresponding to Layer 3 (routed) VSN I-SIDs on either a Controller or a Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **L3-VSN** tab.
4. (Optional) To filter rows based on specific filter criteria, select **Filter**.

L3-VSN Field Descriptions

Use the data in the following table to use the **L3-VSN** tab.

Name	Description
VrfId	Specifies the VRF ID of the VRF corresponding to the Layer 3 VSN I-SID.
Isid	Specifies the Layer 3 VSN I-SID.
VrfName	Specifies the VRF name of the VRF corresponding to the Layer 3 VSN I-SID.
InjectDefaultRouteDisable	Specifies whether injection of default routes is disabled.

View the DvR Database

Before You Begin

Ensure that you enable DvR on the node.

About This Task

View all DvR routes on a Controller or a Leaf node.

The Controller node displays all the IP subnet routes configured for that DvR domain. A Leaf node displays all IP subnet routes learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. It also displays the Host Routes (ARPs) learned from other DvR enabled nodes.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **Database** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Database Field Descriptions

Use the data in the following table to use the **Database** tab.

Name	Description
DestIpAddrType	Specifies the address type of the IPv4 destination address of the DvR database entry.
DestIpAddr	Specifies the IPv4 destination address of the DvR database entry.
DestMask	Specifies the destination mask of the DvR database entry.
L3Isid	Specifies the Layer 3 I-SID of the DvR database entry.

Name	Description
EcmpIndex	Specifies the ECMP index for the DvR database entry.
NextHop	Specifies the MAC address of the next hop BEB, in the DvR database entry.
L2lsid	Specifies the Layer 2 I-SID of the DvR database entry.
Vrfid	Specifies the VRF ID for the DvR database entry.
OutgoingInterface	Specifies the outgoing interface (port or MLT) of the DvR database entry.
SpbCost	Specifies the SPB cost of the DvR database entry.
PrefixCost	Specifies the prefix cost of the DvR database entry.
NextHopName	Specifies the host name of the next hop BEB, in the DvR database table entry.
Age	Specifies the uptime since creation of the DvR database table entry.

View the Remote DvR Database



Note

This procedure does not apply to 5320 Series and 5420 Series.

Before You Begin

Ensure that you enable DvR on the node.

About This Task

Perform this procedure to view all remote DvR routes on a Controller or a Leaf node.

The Controller node displays all the IP subnet routes configured for that DvR domain. A Leaf node displays all IP subnet routes learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. It also displays the Host Routes (ARPs) learned from other DvR enabled nodes.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR Remote**.
3. Select the **Database** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Database Field Descriptions

Use the data in the following table to use the **Database** tab.

Name	Description
DestIpAddrType	Specifies the address type of the IPv4 destination address of the DvR database entry.
DestIpAddr	Specifies the IPv4 destination address of the DvR database entry.
DestMask	Specifies the destination mask of the DvR database entry.
L3Isid	Specifies the Layer 3 I-SID of the DvR database entry.
EcmpIndex	Specifies the ECMP index for the DvR database entry.
NextHop	Specifies the MAC address of the next hop BEB, in the DvR database entry.
L2Isid	Specifies the Layer 2 I-SID of the DvR database entry.
VrfId	Specifies the VRF ID for the DvR database entry.
OutgoingInterface	Specifies the outgoing interface (port or MLT) of the DvR database entry.
SpbCost	Specifies the SPB cost of the DvR database entry.
PrefixCost	Specifies the prefix cost of the DvR database entry.
NextHopName	Specifies the host name of the next hop BEB, in the DvR database table entry.
Age	Specifies the uptime since creation of the DvR database table entry.

View DvR Backbone Entries on a Controller

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

View the DvR backbone entries (redistributed host routes) learned from all Controllers in all DvR domains.



Note

You can view DvR backbone entries only on a Controller. Viewing backbone entries does not apply to a Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.

2. Select **DVR**.
3. Select the **Backbone Entries** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Backbone Entries Field Descriptions

Use the data in the following table to use the **Backbone Entries** tab.

Name	Description
IpAddrType	Specifies the address type of the DvR backbone host (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR backbone host.
L3lsid	Specifies the Layer 3 I-SID of the DvR backbone host.
DomainId	Specifies the domain ID of the DvR backbone host.
EcmpIndex	Specifies the ECMP index of the DvR backbone host.
HostMacAddr	Specifies the MAC address of DvR backbone host.
L2lsid	Specifies the Layer 2 I-SID of the DvR backbone host.
AdvControllerName	Specifies the host name of the advertising Controller.
AdvController	Specifies the host MAC address of the advertising Controller.
NextHopName	Specifies the host name of the next hop Backbone host in the DvR route.
NextHopMac	Specifies the MAC address of the next hop Backbone host in the DvR route.

View Remote DvR Backbone Entries on a Controller



Note

This procedure does not apply to 5320 Series and 5420 Series.

Before You Begin

Ensure that you enable DvR globally on the node.

About This Task

Perform this procedure to view the remote DvR backbone entries (redistributed host routes) learned from all Controllers in all DvR domains.



Note

You can view DvR backbone entries only on a Controller. Viewing backbone entries does not apply to a Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR Remote**.
3. Select the **Backbone Entries** tab.
4. (Optional) To filter the rows based on the specific criteria, select **Filter**.

Backbone Entries *Field Descriptions*

Use the data in the following table to use the **Backbone Entries** tab.

Name	Description
IpAddrType	Specifies the address type of the DvRbackbone host (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR backbone host.
L3Isid	Specifies the Layer 3 I-SID of the DvR backbone host.
DomainId	Specifies the domain ID of the DvR backbone host.
EcmpIndex	Specifies the ECMP index of the DvR backbone host.
HostMacAddr	Specifies the MAC address of DvR backbone host.
L2Isid	Specifies the Layer 2 I-SID of the DvR backbone host.
AdvControllerName	Specifies the host name of the advertising Controller.
AdvController	Specifies the host MAC address of the advertising Controller.
NextHopName	Specifies the host name of the next hop Backbone host in the DvR route.
NextHopMac	Specifies the MAC address of the next hop Backbone host in the DvR route.

View DvR Multi-area SPB Backbone Entries on a Controller



Note

This procedure does not apply to 5320 Series, 5420 Series, or 5520 Series.

Before You Begin

Enable Distributed Virtual Routing (DvR) globally on the node.

About This Task

Perform this procedure to view the DvR Multi-area SPB backbone entries that the system learns from all Controllers in all DvR domains.



Note

You can view the backbone entries only on a DvR Controller. Viewing backbone entries does not apply to a DvR Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR**.
3. Select the **MultiArea Backbone Entries** tab.
4. (Optional) Select the **Filter** button to filter the rows on basis of specific criteria.

MultiArea Backbone Entries Field Descriptions

Use the data in the following table to use the **MultiArea Backbone Entries** tab.

Name	Description
IpAddrType	Specifies the address type of the DvR backbone host (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR backbone host.
L3Isid	Specifies the Layer 3 I-SID of the DvR backbone host.
DomainId	Specifies the domain ID of the DvR backbone host.
EcmpIndex	Specifies the ECMP index of the DvR backbone host.
HostMacAddr	Specifies the MAC address of DvR backbone host.
L2Isid	Specifies the Layer 2 I-SID of the DvR backbone host.
AdvControllerName	Specifies the host name of the advertising Controller.
AdvController	Specifies the host MAC address of the advertising Controller.

Name	Description
NextHopName	Specifies the host name of the next hop Backbone host in the DvR route.
NextHopMac	Specifies the MAC address of the next hop Backbone host in the DvR route.
Area	Specifies the type of area as home or remote.

View Remote DvR Multi-area SPB Backbone Entries on a Controller



Note

This procedure does not apply to 5320 Series, 5420 Series, or 5520 Series.

Before You Begin

Enable Distributed Virtual Routing (DvR) globally on the node.

About This Task

Perform this procedure to view the remote DvR Multi-area SPB backbone entries that the system learns from all Controllers in all DvR domains.



Note

You can view the backbone entries only on a DvR Controller. Viewing backbone entries does not apply to a DvR Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **DVR Remote**.
3. Select the **MultiArea Backbone Entries** tab.
4. (Optional) Select the **Filter** button to filter the rows on basis of specific criteria.

MultiArea Backbone Entries *Field Descriptions*

Use the data in the following table to use the **MultiArea Backbone Entries** tab.

Name	Description
IpAddrType	Specifies the address type of the DvR backbone host (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR backbone host.
L3Isid	Specifies the Layer 3 I-SID of the DvR backbone host.
DomainId	Specifies the domain ID of the DvR backbone host.

Name	Description
EcmpIndex	Specifies the ECMP index of the DvR backbone host.
HostMacAddr	Specifies the MAC address of DvR backbone host.
L2Isid	Specifies the Layer 2 I-SID of the DvR backbone host.
AdvControllerName	Specifies the host name of the advertising Controller.
AdvController	Specifies the host MAC address of the advertising Controller.
NextHopName	Specifies the host name of the next hop Backbone host in the DvR route.
NextHopMac	Specifies the MAC address of the next hop Backbone host in the DvR route.
Area	Specifies the type of area as home or remote.



Extensible Authentication Protocol over LAN

[EAPoL](#) on page 1011

[EAPoL Configuration Using CLI](#) on page 1037

[EAP Configuration Using Enterprise Device Manager](#) on page 1067

Table 100: Extensible Authentication Protocol over LAN product support

Feature	Product	Release introduced
Extensible Authentication Protocol (EAP) and EAP over LAN (EAPoL)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
EAPoL MHMA-MV	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 100: Extensible Authentication Protocol over LAN product support (continued)

Feature	Product	Release introduced
EAPoL enhancements: Enhanced MHMV, Fail Open VLAN, Guest VLAN	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
EAP enhancements: EAP on Flex UNI ports, Auto-sense ports, auto-isid-offset	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3
EAP enhancements: Wake on LAN, Guest I-SID, Fail Open I-SID	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

Table 100: Extensible Authentication Protocol over LAN product support (continued)

Feature	Product	Release introduced
Non EAPoL MAC RADIUS authentication	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
QoS Priority Assignment	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
RADIUS Dynamic User-Based Policies	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

Table 100: Extensible Authentication Protocol over LAN product support (continued)

Feature	Product	Release introduced
RADIUS Port and VLAN based Attributes	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.4
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.4
	VSP 7400 Series	VOSS 8.4
Continuity Mode for Fail Open VLAN and Fail Open I-SID	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	Fabric Engine 8.6
	5520 Series	Fabric Engine 8.6
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.6
	VSP 7400 Series	VOSS 8.6
NEAP fallback if EAP authentication fails	4220 Series	Fabric Engine 9.3
	5320 Series	Fabric Engine 9.3
	5420 Series	Fabric Engine 9.3
	5520 Series	Fabric Engine 9.3
	5720 Series	Fabric Engine 9.3
	7520 Series	Fabric Engine 9.3
	7720 Series	Fabric Engine 9.3
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.3
	VSP 7400 Series	VOSS 9.3

EAPoL

Extensible Authentication Protocol over LAN (EAPoL or EAP) is a port-based network access control protocol. EAP provides security by preventing users from accessing

network resources before they are authenticated. The EAP authentication feature prevents users from accessing a network to assume a valid identity and access confidential material or launch denial-of-service attacks.

You can use EAP to set up network access control on internal LANs and to exchange authentication information between an end station or server that connects to a switch and an authentication server (such as a RADIUS server). This security feature extends the benefits of remote authentication to internal LAN clients. For example, if a new client PC fails the authentication process, EAP prevents the new client PC from accessing the network.

EAP terminology

This section lists some components and terms used with EAP-based security.

- Supplicant—a device, such as a PC, that applies for access to the network.
- Authenticator—software on a switch that authorizes or rejects a Supplicant attached to the other end of a LAN segment.
 - Port Access Entity (PAE)—software that controls each port on the device. The PAE, which resides on the switch, supports the Authenticator functionality.
 - Controlled Port—any port on the device with EAP enabled.
- Authentication Server—a RADIUS server that provides AAA services to the authenticator.

EAP Configuration Considerations

This section lists EAP configuration considerations.

- You must configure at least one EAP RADIUS server and shared secret fields.
- You cannot configure EAP on ports that are currently configured for the following:
 - Shared segments
 - MultiLink Trunking
- Change the authentication status to auto for each port that you want to control. The *auto* setting automatically authenticates the port according to the results of the RADIUS server. The default authentication setting for each port is *authorized*.
- When multiple clients are authenticated on the same port, the priority of the latest incoming client is applied on the port, and this priority is retained until all the clients log out on that port.

Configuration Process

The Authenticator facilitates the authentication exchanges that occur between the Supplicant and the Authentication Server. The Authenticator PORT ACCESS ENTITY (PAE) encapsulates the EAP message into a RADIUS packet, and then sends the packet to the Authentication Server to begin the EAP authentication process.

The Authenticator manages the access to controlled port. At system initialization, or when a Supplicant initially connects to one of the controlled ports on the device, the system blocks data traffic of the Supplicant until gets authenticated. After the

Authentication Server notifies the Authenticator PAE about the success or failure of the authentication, the Authenticator decides whether to permit/deny the traffic of client on controlled port.

The following figure illustrates how the switch, configured with EAP, reacts to a new network connection.

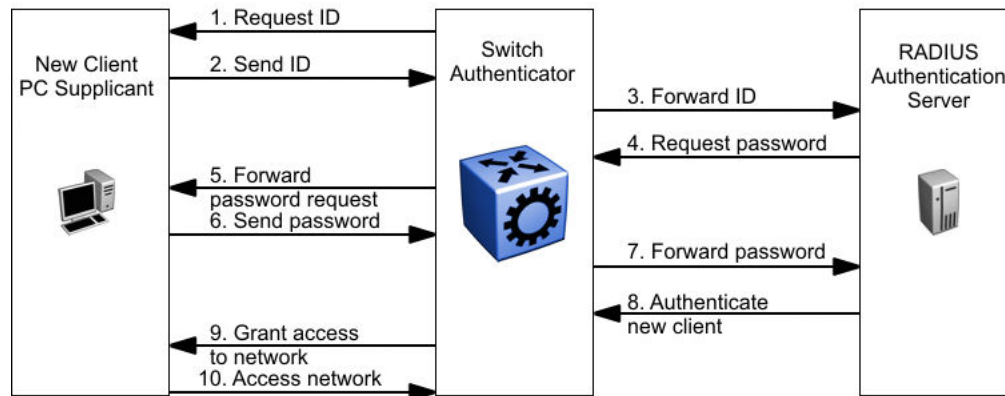


Figure 55: EAP configuration example

In the preceding figure, the switch uses the following steps to authenticate a new client:

1. The switch detects a new connection on one of its EAP-enabled ports and requests a user ID from the new client PC.
2. The new client sends its user ID to the switch.
3. The switch uses RADIUS to forward the user ID to the RADIUS server.
4. The RADIUS server responds with a request for the password of the user.
5. The switch forwards the request from the RADIUS server to the new client.
6. The new client sends an encrypted password to the switch, within the EAP packet.
7. The switch forwards the EAP packet to the RADIUS server.
8. The RADIUS server authenticates the password.
9. The switch grants the new client access to the network.
10. The new client accesses the network.

The following figure shows the flow diagram for EAP on a switch.




Figure 56: EAP flow diagram

EAP System Requirements

The following are the minimum system requirements for EAP:

- RADIUS server

- Client software that supports EAP

You must specify the RADIUS server that supports EAP as the primary RADIUS server for the switch. You must configure your switch for VLANs and EAP security.

If you configure EAP on a port, the following limitations apply:

- You cannot enable EAP on ports that belong to an MLT group or add EAP-enabled ports to an MLT group.
- You cannot configure EAP on MLT/LACP interfaces or add EAP-enabled ports to an LACP group.
- You cannot enable the following features on EAP-enabled ports:
 - MACsec
 - VLACP
- Manual VLAN changes on a EAP enabled port is restricted.
- You cannot change the VLAN port tagging on EAP enabled ports.
- You cannot configure the default VLAN ID. Use the Guest VLAN configuration to access unauthenticated devices.
- You can configure a total of 32 MAC clients, EAP and NEAP hosts, on an EAP-enabled port. Two MAC clients per port is a typical configuration.
- You cannot enable EAP on a network-to-network interface (NNI).
- You cannot egress mirror an EAP PDU.
- Do not use EAP with a brouter port.
- Ping to and from services between nodes over the NNI will work even when it contains only EAP enabled ports with no authenticated clients on it.
- MHSA and Fail Open VLAN are mutually exclusive.
- Fail-Open I-SID is not supported in MHSA mode.
- You cannot change the EAP operation mode on EAP enabled ports.
- You cannot configure private VLANs or SPBM B-VLAN as Fail Open VLAN or Guest VLAN.
- You cannot delete a VLAN if the VLAN is configured as Fail Open VLAN or Guest VLAN.

EAP Dynamic VLAN Assignment

If you configure a RADIUS server to send a VLAN ID in the Access-Accept response, the EAP feature dynamically changes the VLAN configuration of the port by adding the port to the specified VLAN.

EAP dynamic VLAN assignment affects the following VLAN configuration values:

- Port membership
- Port priority
- Default VLAN ID

When you disable EAP on a port that was previously authorized, VLAN configuration values for that port are restored directly from the nonvolatile random access memory (NVRAM) of the device.

You can set up your Authentication Server (RADIUS server) for EAP dynamic VLAN assignments. You can use the Authentication Server to configure user-specific settings for VLAN memberships and port priority.

When you log on to a system that is configured for EAP authentication, the Authentication Server recognizes your user ID and notifies the device to assign preconfigured (user-specific) VLAN membership and port priorities to the device. The configuration settings are based on configuration parameters that were customized for your user ID and previously stored on the Authentication Server.

**Note**

Static entries like IGMP, ARP, FDB configured on a port of a VLAN interface, will not be retained if the port is assigned a same VLAN by the RADIUS server and the client authenticated on the port gets disconnected or unauthenticated.

Multiple Host Multiple VLAN

With the MHMV feature, you can assign multiple authenticated devices to different VLANs on the same EAP-enabled port using device MAC addresses. Using RADIUS VLAN attributes, different clients can access different VLANs. This separates traffic for different MAC clients.

In MHMV mode, the port-priority assigned by the RADIUS server is configured by MAC address for each authenticated client. After configuration, the QoS level on the port does not change.

Use MHMV to assign multiple authenticated devices to different VLANs on the same port. Clients can access different VLANs access using the MAC address of the devices. Different clients with different level of access (unauthorized to authorized) in different VLANs and with different QoS priorities, can exist on the same port.

With MHMV, EAP Multihost VLAN supports tagged and untagged ports. A port can be a member of multiple tagged and untagged VLANs.

In MHMV mode, MAC based VLANs support traffic separation between different authenticated MAC clients. MAC based VLAN traffic separation applies only to untagged VLAN traffic. If the data traffic is tagged and if VLAN is configured on the port, then the traffic is forwarded to the VLAN associated with the tag.

Multiple Host Multiple VLAN Usage

The following example illustrates the usage scenario for a MHMV port with unauthenticated clients:

- Clients (n) connect to a switch port. For the maximum number of clients (EAP + NEAP) allowed on a port, see *Fabric Engine Release Notes*.
- EAP is enabled and the default operation mode is MHMV.
- Modify client counters to authenticate n clients.
- Initial VLANs are the VLANs which are manually set up before EAP is enabled.
- Port default VLAN ID is equal to one of the initial VLAN ID.
- All clients are unauthenticated, hence the clients cannot access the network.

The following figure represents the functionality when clients are not authenticated.

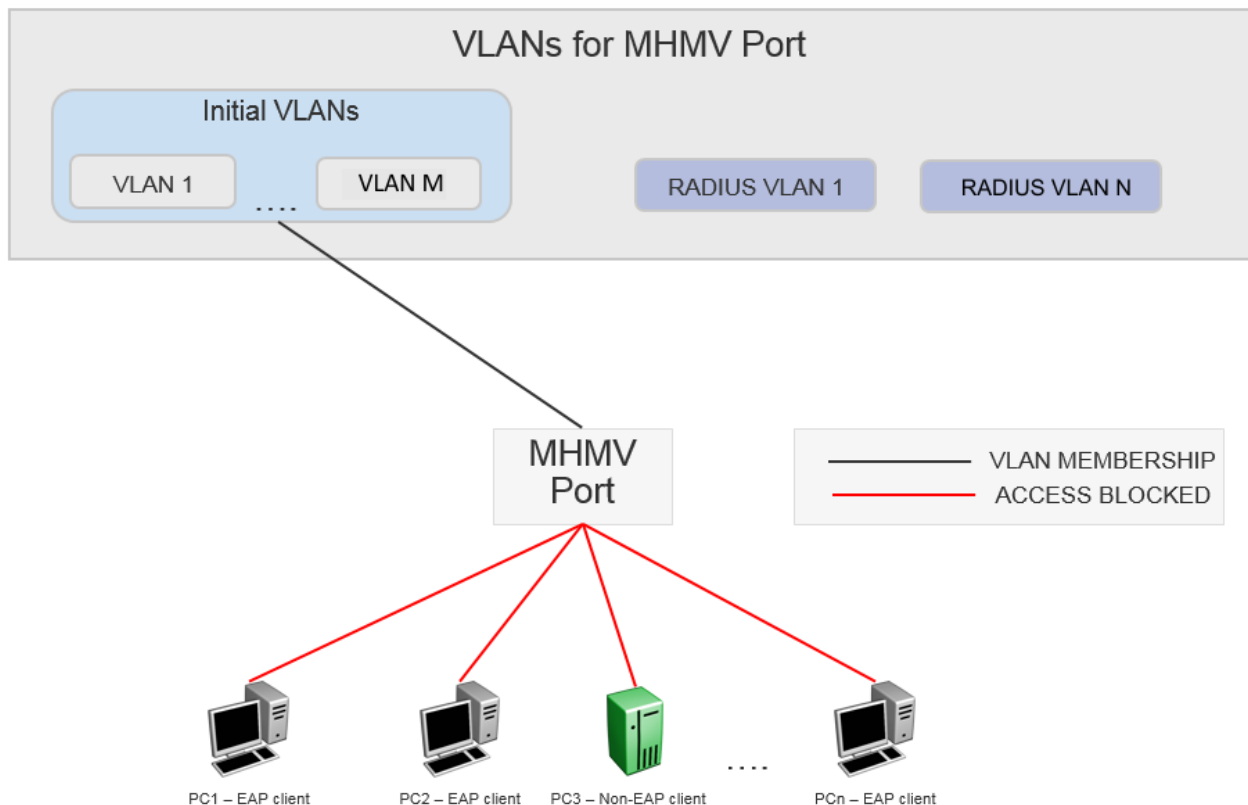


Figure 57: MHMV - clients not authenticated



Note

The clients cannot access the network because they are not authenticated.

When client PC1 authenticates, there are two scenarios:

1. Client PC1 does not receive RADIUS VLAN attribute:
 - There are no changes to the port membership and port default VLAN ID.
 - PC1 is the only client that is allowed access to the initial VLANs.
 - A VLAN MAC rule is added that associates the MAC with the default VLAN ID.
 - If the VLAN is configured on the port, then the tagged traffic from PC1 is forwarded to the VLAN associated with the tag.
 - Untagged traffic from PC1 is forwarded to the port default VLAN.
2. Client PC1 receives RADIUS VLAN attribute:
 - The port is left in all initial VLANs and added to the VLAN corresponding to the RADIUS VLAN attribute.
 - Port default VLAN remains unchanged.
 - A VLAN MAC based rule is configured for client PC1.
 - Using the VLAN MAC based capabilities, the untagged traffic from PC1 goes to the RADIUS assigned VLAN 1 as shown in [Figure 58](#).
 - Client PC1 can access all initial VLANs using tagged frames.
 - The remaining clients stay unauthenticated and cannot access any VLANs.

The following figure represents the functionality when client PC1 authenticates.

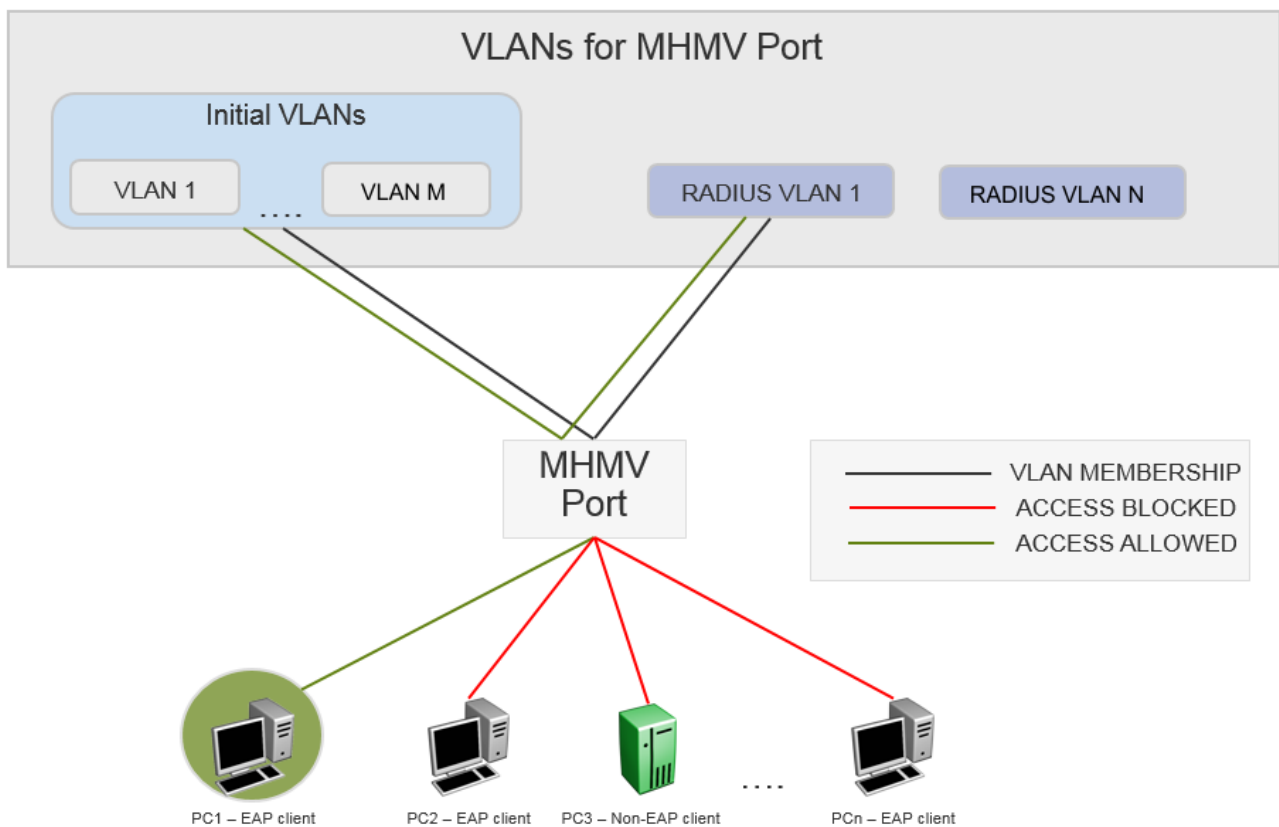


Figure 58: MHMV - authenticated client



Note

PC1 is authenticated with RADIUS VLAN 1. The other clients cannot access the network as they are unauthenticated.

When a client disconnects the following happens:

- The MAC VLAN rule is removed from the switch.
- If the RADIUS VLAN attribute was used with the client was authenticated and no other clients are authenticated on that RADIUS VLAN, then the port is removed from the VLAN.
- The RADIUS accounting attribute Acct-Terminate-Cause indicates how a session was terminated.
- The RADIUS accounting attribute Event-Timestamp indicates the time that an event occurred on the Network Access Server (NAS).

EAP Functionality on Flex UNI Ports

After you enable EAP on a port, all client MACs (MHMV mode) must be authenticated by the RADIUS server in order to have network access. In Multiple Host Multiple VLAN (MHMV) mode, the RADIUS server allocates a MAC to I-SID binding to each client that connects to the switch and uses it to transmit traffic. The binding is not between the

MAC and the VLAN. Untagged S-UNIs generated from the RADIUS server for a MAC or MACs are considered as MAC-based S-UNIs.

The RADIUS server also provides the VLAN:ISID binding for the MAC, which results in the addition of an untagged Switched UNI (S-UNI) for that particular I-SID. Only the MAC or MACs that receive the I-SID from the RADIUS server can transmit traffic to Extensible Authentication Protocol (EAP)-enabled Flex UNI ports.

The switch uses MAC-based S-UNIs with EAP-enabled Flex UNI ports in MHMV mode only.

The MAC-based S-UNI model does not apply to Multiple Host Single Authentication (MHSA) mode. MHSA mode uses the untagged S-UNI model that exists on VOSS switches. S-UNIs generated from the information obtained from the RADIUS server are considered as classic or default untagged S-UNIs.

**Note**

EAP is not supported on MLT/SMULTs. Only the EAP I-SIDs are synchronized between one vIST peer and another vIST peer. S-UNIs are not synchronized with the vIST peer.

You can configure an untagged Switched UNI (S-UNI) on an EAP-enabled interface if one is not already configured. You must associate the untagged S-UNI with a platform VLAN.

EAP with Flex UNI and DvR

EAP with Flex UNI is supported on Distributed Virtual Routing (DvR) Leafs. An untagged S-UNI (where the system learns MACs based on the I-SID to MAC binding) must have a platform VLAN associated with it. If a default untagged S-UNI is used, the corresponding S-UNI must be received from the DvR Controllers.

EAP and Fabric Attach

With Extensible Authentication Protocol (EAP) and Fabric Attach (FA), FA-capable switches can forward traffic from EAP/NEAP clients over the SPB cloud. The traffic for authenticated clients is mapped to I-SIDs received from RADIUS server.

You must configure the desired bindings for EAP/NEAP clients on the RADIUS server. When confirming the authentication request, the RADIUS server also sends the corresponding binding for the EAP/NEAP client.

The FA Proxy sends to the FA Server the binding received from the RADIUS server. If the FA Server rejects all the bindings, the client is disconnected. EAP clients are moved from AUTHENTICATED state to HELD state.

On an FA Server, when an EAP/NEAP device is authenticated and an FA binding is received from the RADIUS server, a Switched UNI (S-UNI) is created.

After an EAP/NEAP client is disconnected, the switch cleans-up the binding associated with the client, if no other EAP/NEAP client on that port uses it.

EAP and FA can be enabled in any order; however, EAP must have Flex UNI enabled in order to function on an FA-enabled port.

FA clients that generate S-UNI bindings must be used with EAP MHSA mode, while FA clients that do not generate S-UNI bindings should be used with EAP MHMV mode.

MAC Move Detection on EAP Ports

The MAC moves mechanism detects when a MAC address migrates from one port to another port, such as:

- EAPoL-enabled port to EAPoL-enabled port
- EAPoL-enabled port to non-EAPoL (NEAP)-enabled port
- non EAPoL-enabled port to EAPoL-enabled port

When a MAC address migrates from one port to another port, the new EAPoL-enabled port triggers a new RADIUS authentication. New bindings are applied on the new EAPoL-enabled port. The old port detects the MAC is moved and automatically deletes the old binding or bindings.

The Mac moves mechanism works between vIST peers when a MAC address on one peer migrates to the other peer but only if the I-SID in which the MAC address is learned exists on the new and the old peer.

Auto-sense Ports

EAP and NEAP is integrated with Auto-sense infrastructure. If a RADIUS server is configured on the switch, Auto-sense-enabled ports activate EAP and NEAP authentication automatically.

For more information about Auto-sense functionality, see [Auto-sense](#) on page 21.

RADIUS-Assigned VLAN or VLAN:ISID Bindings

RADIUS-assigned VLAN and VLAN:ISID bindings provide greater flexibility and a more centralized assignment. The RADIUS server can dynamically assign VLANs or VLAN:ISID bindings to a port.

Use VLAN attributes for RADIUS assignments in VLAN mode. This mode applies when the EAP-enabled port does not have Flex-UNI enabled.

Use VLAN:ISID attributes for RADIUS assignments in I-SID mode. This mode applies when the EAP-enabled port does have Flex-UNI enabled.

For more information, see [RADIUS Attributes](#) on page 3108.

RADIUS Configuration Prerequisites for EAP

Connect the RADIUS server to a force-authorized port. This ensures that the port is always available and not tied to whether or not the device is EAP-enabled.

RADIUS Accounting for EAP

The switch provides the ability to account EAP and NEAP sessions using the RADIUS accounting protocol. A user session is defined as the time frame between when a user is authenticated until the user is unauthenticated.

The following table summarizes the accounting events and information logged.

Table 101: Summary of accounting events and information logged

Event	RADIUS attributes	Description
User is authenticated by EAP	Acct-Status-Type	Start
	Nas-IP-Address	IP address to represent the switch
	Nas-Port	Port number on which the user is EAP or NEAP authorized
	Acct-Session-ID	Unique string representing the session
	User-Name	EAP user name or NEAP MAC
User logs off	Acct-Status-Type	Stop
	Nas-IP-Address	IP address to represent the switch
	Nas-Port	Port number on which the user is EAP or NEAP unauthorized
	Acct-Session-ID	Unique string representing the session
	User-Name	EAP user name
	Acct-Input-Octets	Number of octets input to the port during the session
	Acct-Output-Octets	Number of octets output to the port during the session
	Acct-Terminate-Cause	Reason for terminating user session. For more information about the mapping of 802.1x session termination cause to RADIUS accounting attribute, see the following table.
	Acct-Session-Time	Session interval

The following table describes the mapping of the causes of 802.1x session terminations to the corresponding RADIUS accounting attributes.

Table 102: 802.1x session termination mapping

IEEE 802.1Xdot1xAuthSessionTerminateCause Value	RADIUSAcct-Terminate-Cause Value
supplicantLogoff(1)	User Request (1)
portFailure(2)	Lost Carrier (2)
supplicantRestart(3)	Supplicant Restart (19)
reauthFailed(4)	Reauthentication Failure (20)
authControlForceUnauth(5)	Admin Reset (6)
portReInit(6)	Port Reinitialized (21)
portAdminDisabled(7)	Port Administratively Disabled (22)
notTerminatedYet(999)	

RADIUS Dynamic User-Based Policies

RADIUS Dynamic User-Based Policies is a security feature to control access services on user devices that connect to the network. Before enabling any services on the user device, the RADIUS server authenticates each device that connects to the switch port and assigns that port to a VLAN or a VLAN to I-SID binding. RADIUS Dynamic User-Based Policies implement a dynamic method to apply filter Access Control List (ACL) rules to Extensible Authentication Protocol (EAP) and Non-EAP (NEAP) authenticated user traffic. The RADIUS server authenticates the user device for switch access and sends rules for that user device to the switch.

The system clears the rules when the following events occur:

- You disable EAPoL globally on the switch.
- EAP and NEAP sessions are cleared.
- You shutdown the port.



Note

- You must enable RADIUS and EAP over LAN (EAPoL) on the switch. For more information, see [Enable RADIUS Authentication](#) on page 3061 and [Globally Enable EAP on the Device](#) on page 1037.
- You must configure an EAP-enabled RADIUS server. For more information, see [Configure an EAP-Enabled RADIUS Server](#) on page 1040.

RADIUS Dynamic User-Based Policies support one time configuration of policy attributes on the RADIUS server and dynamically creates the policies on multiple switches within the network. This process of automatically creating policies enhances the speed of network access for authenticated users and also facilitates faster network synchronization in the event of network-wide policy changes.

Extreme Networks Vendor ID 1916 supports the following RADIUS Vendor Specific Attribute (VSA) for RADIUS Dynamic User-Based Policies:

- Extreme-Dynamic-ACL (ID 251)

For more information, see [RADIUS Attributes](#) on page 3108.

The RADIUS server contains the RADIUS VSAs in a configuration file for each EAP or NEAP client that the switch authenticates. Following is an example of a RADIUS VSA configured on the RADIUS server:

```
000000000000a Cleartext-Password := "000000000000a"
Service-Type = Framed-User,
Framed-Protocol = PPP,
Auth-Type := Accept,
Fabric-Attach-ISID = 10:100,
Extreme-Dynamic-ACL = CLIENT RadiusGuest
Extreme-Dynamic-ACL = acl inPort
Extreme-Dynamic-ACL = ace 1 sec name ACE-A1 ethernet ether-type eq 0x800 & action deny & ip ipprotocol-
type eq 17 & protocol dst-port eq 4000
Extreme-Dynamic-ACL = ace 2 sec name ACE-A2 ethernet ether-type eq ip & ip dst-ip eq 10.10.10.1 &
action deny
Extreme-Dynamic-ACL = acl set default-action deny
```

When the switch receives a new VSA with ACL and Access Control Entries (ACE) rules from the RADIUS server, the switch dynamically creates the ACL infrastructure based on the following:

- Dynamic ACLs - the switch allocates one dynamic ACL for each EAP enabled port. You cannot manually configure the dynamic ACL. The dynamic behavior of the ACL depends on the EAP port state (MHMV or MHSA). RADIUS Dynamic User-Based Policies support the inPort and outPort ACL types. You can display the filter ACL configuration on the switch using the **show filter acl** command, to identify the source of ACL configuration (static or dynamic).
- Dynamic ACEs - after the switch configures an ACL as dynamic, the system automatically considers the ACEs in that ACL as dynamic. You cannot manually configure the ACEs in a dynamic ACL. When the switch receives an ACE rule from the RADIUS server, the system allocates an ACE ID to it. Each ACE rule carries a relative order that helps the switch to set priority for the ACE rules that the switch receives. For handling of Radius ACL rules, the switch parses the rules first. Based on the actions, the system classifies the rules as security ACEs or QoS ACEs. If the switch is unable to recognize the qualifiers or actions in a rule, then the switch ignores that rule.
- Multiple Host Multiple VLAN (MHMV) operating mode - the system authenticates each MAC that the switch receives on the EAP-enabled port and assigns the MAC to a specific VLAN or VLAN to I-SID binding. The system uses the VLAN to I-SID binding when Flex UNI is enabled on a port. The system processes the ACE rules that the switch receives from the RADIUS server on a per MAC basis, the system translates the default-action into an ACE rule with actions, deny or permit. When the switch processes the RADIUS VSAs, the system adds the MAC as a qualifier for each ACE rule.
- Multiple Host Single Authentication (MHSA) operating mode - the system processes the ACE rules that it receives from the RADIUS server on a per port basis.

NEAP host

The following section provides information about NEAP hosts on EAP-enabled ports and RADIUS authentication.

NEAP Hosts on EAP Enabled Ports

For an EAP-enabled port configured for NEAP host support, devices with MAC addresses getting authenticated are allowed access to the port.

The switch allows NEAP hosts whose MAC addresses are authenticated by RADIUS.

Support for NEAP hosts on EAP-enabled ports is primarily intended to accommodate printers and other passive devices sharing a hub with EAP clients.

Support for NEAP hosts on EAP-enabled ports includes the following features:

- Authenticated NEAP clients are hosts that satisfy one of the following criteria:
 - Host MAC address is authenticated by RADIUS.
- NEAP hosts are allowed even if no authenticated EAP hosts exist on the port.
- When a new host is seen on the port, NEAP authentication is performed as follows:
 - The switch generates a <user name, password> pair, which it forwards to the network RADIUS server for authentication.

NEAP MAC RADIUS Authentication

For RADIUS authentication of a NEAP host MAC address, the switch generates a <user name, password> pair as described in the following list:

- The user name is the NEAP MAC address in string format.
- The password is a string that combines the switch IP address, MAC address, port number and user-configurable key string. If padding option is enabled, the system will specify a dot(.) for every missing parameter. IP address is represented by three decimal characters per octet.



Important

Use the following Global Configuration examples to select a password format that combines one or more of these three elements:

- Padding enabled , password = 010010011253..05. (when the switch IP address and port are used).
- Padding enabled, password = 010010011253... (when only the switch IP address is used).
- No padding (default option). Password = 000011220001 (when only the user's MAC address is used).

The following example illustrates the <user name, password> pair format with no padding enabled and using the IP address, MAC address, and key-string as the password.

```
switch IP address = 192.0.2.5
non-EAP host MAC address = 00 C0 C1 C2 C3 C4
```

```
port = 25
Key-String = abcdef
```

- user name = 00C0C1C2C3C4
- password = 010010011253.00C0C1C2C3C4.25.abcdef

Use the command **show eapol system** to verify the formatting.

```
Switch:1>show eapol system

=====
                        Eapol System
=====
                        eap : enabled
                        Eapol Version : 3
                        non-eap-pwd-fmt : mac-addr
                        non-eap-pwd-fmt key : *****
                        non-eap-pwd-fmt padding : disabled
                        auto-isid-offset status : disabled
                        auto-isid-offset value : 15980000
```

NEAP client

The following section provides information for NEAP client.

NEAP Client Re-Authentication

The NEAP client re-authentication feature supports the re-authentication of NEAP clients at defined intervals.

Without re-authentication, client MACs that age out are removed from the client list and are no longer authenticated. If you enable re-authentication, the client MACs remain as long as the RADIUS server is reachable.

When you enable NEAP client re-authentication, an authenticated NEAP client is only removed from the authenticated client list if you remove the client account from the RADIUS server, or if you clear the NEAP authenticated client from the switch.

If you enable NEAP client re-authentication and the RADIUS server that the switch connects to becomes unavailable, the system clears all authenticated NEAP and removes those clients from the switch NEAP client list.

You cannot authenticate one NEAP client on more than one switch port simultaneously. If you connect NEAP clients to a switch port through a hub, those clients are authenticated on that switch port. If you disconnect a NEAP client from the hub and connect it directly to another switch port, the client is authenticated on the new port and its authentication is removed from the port to which the hub is connected.

Silent Devices

For silent devices, you can configure the re-authentication period to 0, which prevents EAP and NEAP sessions on the port from aging-out. If you use this configuration, the session remains active even if the RADIUS server is unreachable.



Caution

Preventing re-authentication can introduce a security risk. As an alternative to using a zero value, configure the maximum re-authentication period on ports that connect to silent devices.

NEAP MAC Learning and Authentication

The system learns the MACs based on the I-SID to MAC binding. When a packet ingresses on a port, which is associated with Switched UNI (S-UNI) I-SID, the system performs MAC look up based on the I-SID. The RADIUS server provides the VLAN:I-SID assignment.



Important

If the default untagged S-UNI is used, you must have a platform VLAN associated with it. This is required to properly transmit traffic and to generate MAC learning events for traffic sent to MAC-based S-UNIs.

When an untagged S-UNI is present on the port, the untagged MAC is initially learned on that S-UNI. When an untagged S-UNI does not exist on the port, the untagged MAC is learned on a special (internal) VLAN. The RADIUS server provides the VLAN:I-SID assignment.

MAC learning for tagged traffic occurs only if there is a tagged S-UNI with the corresponding C-VID on that port. The RADIUS server reconfirms the S-UNI that performed MAC learning.

NEAP Fallback

With NEAP fallback functionality, a single client session can support both EAP and NEAP authentication methods. If a client fails to authenticate using EAP and NEAP authentication is already enabled on the port, the switch automatically attempts NEAP fallback to authenticate the client.

When the EAP process begins, NEAP fallback is triggered after a four-second delay—but only if the Authentication Server detects a new MAC event and you enable NEAP on the port. This ensures NEAP acts as a backup mechanism.

EAP and NEAP Limitations

The EAP and NEAP MAC clients on port limits the maximum number of all EAP and NEAP clients per port. EAP and NEAP MAC clients on port enhancements independently limits the EAP and NEAP clients per port.

The following enhancements are added:

- EAP-MAC-MAX : Limits the total number of EAP clients

- NON-EAP-MAC-MAX: Limits the total number of NEAP clients

**Note**

Do not connect more than 100 EAP and 100 NEAP devices on the switch.

EAP and NEAP mac-max Settings

The total number of EAP clients can be set between 0 and 32, while the total number of NEAP clients can be set between 0 and 8192.

**Note**

EAP-MAC-MAX is overwritten by MAC-MAX. Even if EAP-MAC-MAX is set to a higher limit, then MAC-MAX must not exceed and you must not authenticate more than MAC-MAX clients.

**Note**

NON-EAP-MAC-MAX is overwritten by MAC-MAX. Even if NON-EAP-MAC-MAX is set to a higher limit, then MAC-MAX must not exceed and you must not authenticate more than MAC-MAX clients.

Example Scenarios

1. Scenario 1:

- EAP-MAC-MAX 32
- NON-EAP-MAC-MAX 32
- MAC-MAX 10

In this scenario, there are ten EAP and NEAP authenticated clients, in the order of authentication.

2. Scenario 2:

- EAP-MAC-MAX 1
- NON-EAP-MAC-MAX 1
- MAC-MAX 1

In this scenario, only one EAP or one NEAP client is authenticated, in the order of authentication.

3. Scenario 3:

- EAP-MAC-MAX 5
- NON-EAP-MAC-MAX 10
- MAC-MAX 32

In this scenario, up to five EAP clients and ten NEAP clients are allowed.

4. Scenario 4:

- EAP-MAC-MAX 5
- NON-EAP-MAC-MAX 8
- MAC-MAX 7

In this scenario, up to five EAP clients and seven NEAP clients are allowed. The total number of EAP or NEAP clients is limited to seven.

Multiple Host Single Authentication

Multiple Host Single Authentication (MHSA) allows MACs to access the network without EAP and NEAP authentication. Unauthenticated devices can access the network only after an EAP or NEAP client is successfully authenticated on a port. The VLAN to which the devices are allowed is the client authenticated VLAN. Unless Guest VLAN is configured, there is no authenticated client on the port, and no MAC is allowed to access the network.

MHSA is primarily intended to accommodate printers and other passive devices sharing a hub with EAP and NEAP clients.

MHSA support is on a port-by-port basis for EAP enabled ports.

MHSA supports the following functionality:

- The port remains unauthorized when no authenticated hosts exist on the port. Before the first successful authentication occurs, both EAP and NEAP clients are allowed to negotiate access on that port but only one host is allowed to perform authentication.
- In MHSA mode, QoS level is configured when processing the Port-Priority attribute, because there can only be one authenticated client. The devices behind the authenticated client use the port priority established by the main client.
- In MHSA mode, the Guest VLAN applies only when no authenticated client is present on the port.
- After the first EAP or NEAP client successfully authenticates on a port, other clients cannot negotiate authentication on that port.
- After the first successful authentication, MACs that are already learned on that port is flushed.
- NEAP clients are not removed at age event in MHSA mode.
- There is no limit to the number of MACs that are allowed after first successful authentication.

EAP and NEAP MAC Clients on a Port with MHSA

EAP and NEAP client counters, such as MAC-MAX, EAP-MAC-MAX, and NON-EAP-MAC-MAX do not apply when the port operates in MHSA mode. In MHSA mode, there can be only one authenticated client (EAP or NEAP). Subsequent MACs seen on the port are allowed automatically without authentication.

Guest VLAN

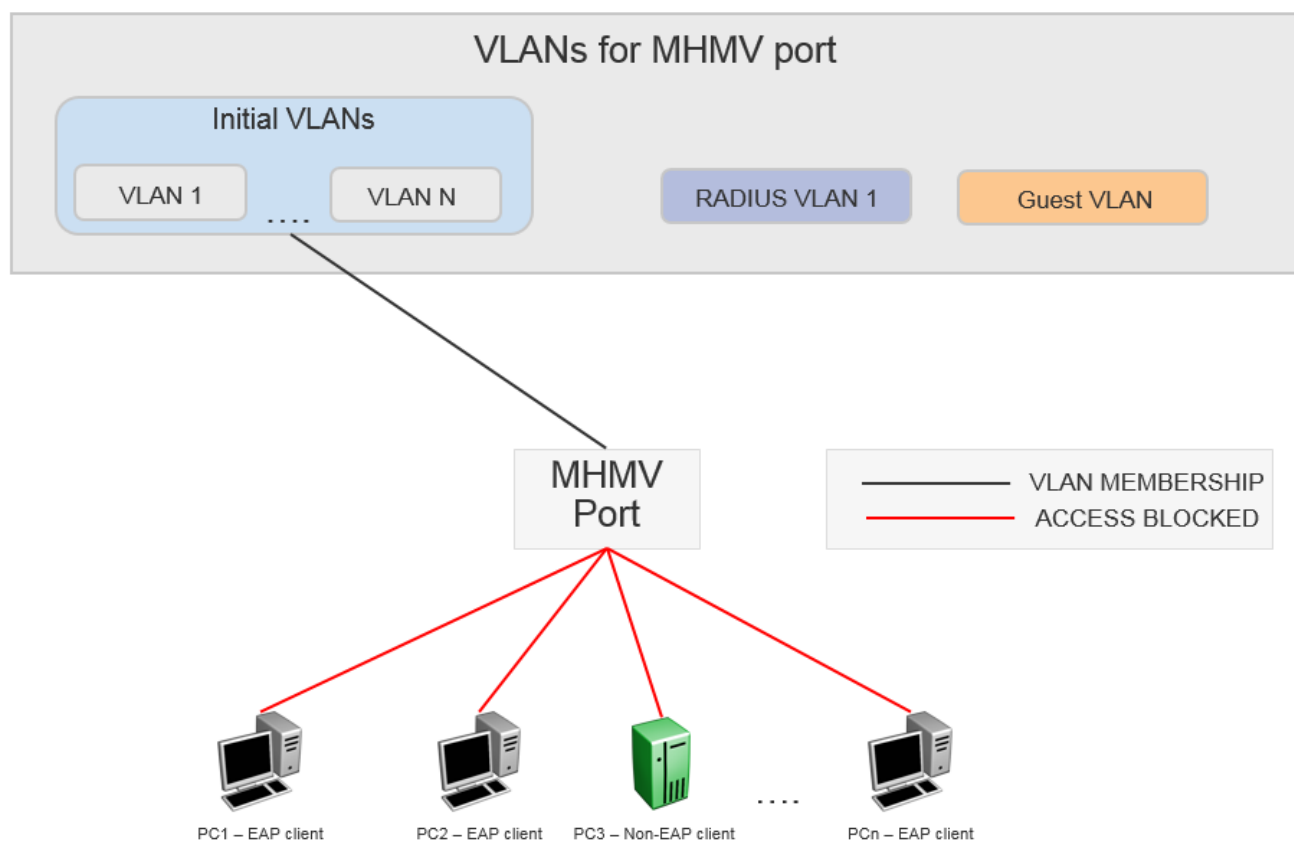
Guest VLAN support provides limited network access until the client is authenticated. Guest VLAN is configured irrespective of the number of authenticated clients present on the port. Guest VLAN is available for each port. Only port based VLANs are used as Guest VLANs. When the Guest VLAN is active, port is added to the VLAN ID, and port default VLAN ID changes to Guest VLAN ID.

Guest VLAN on a MHMV Port Usage Scenario

The following example illustrates the configuration of Guest VLAN support with an EAP MHMV port:

- Clients connect to a switch port through a hub.
- The initial VLANs are the VLANs on which the ports resides after a switch reboot.
- EAP is enabled.
- The port is a member of initial VLANs. The clients cannot access the VLANs since the VLANs are not authenticated. The port default VLAN ID corresponds to one of the initial VLAN IDs.
- Guest VLAN support is not activated.

The following figure represents the functionality when clients are not authenticated.




Note

The clients cannot access the network as they are not authenticated and Guest VLAN is not configured.

- Guest VLAN support is activated.
- The MHMV port is in the initial VLAN stage but gets added to the Guest VLAN ID. The default VLAN ID is updated to correspond to the Guest VLAN ID.
- All Clients behind the port can access the Guest VLAN.

The following figure represents the functionality when Guest VLAN is activated.

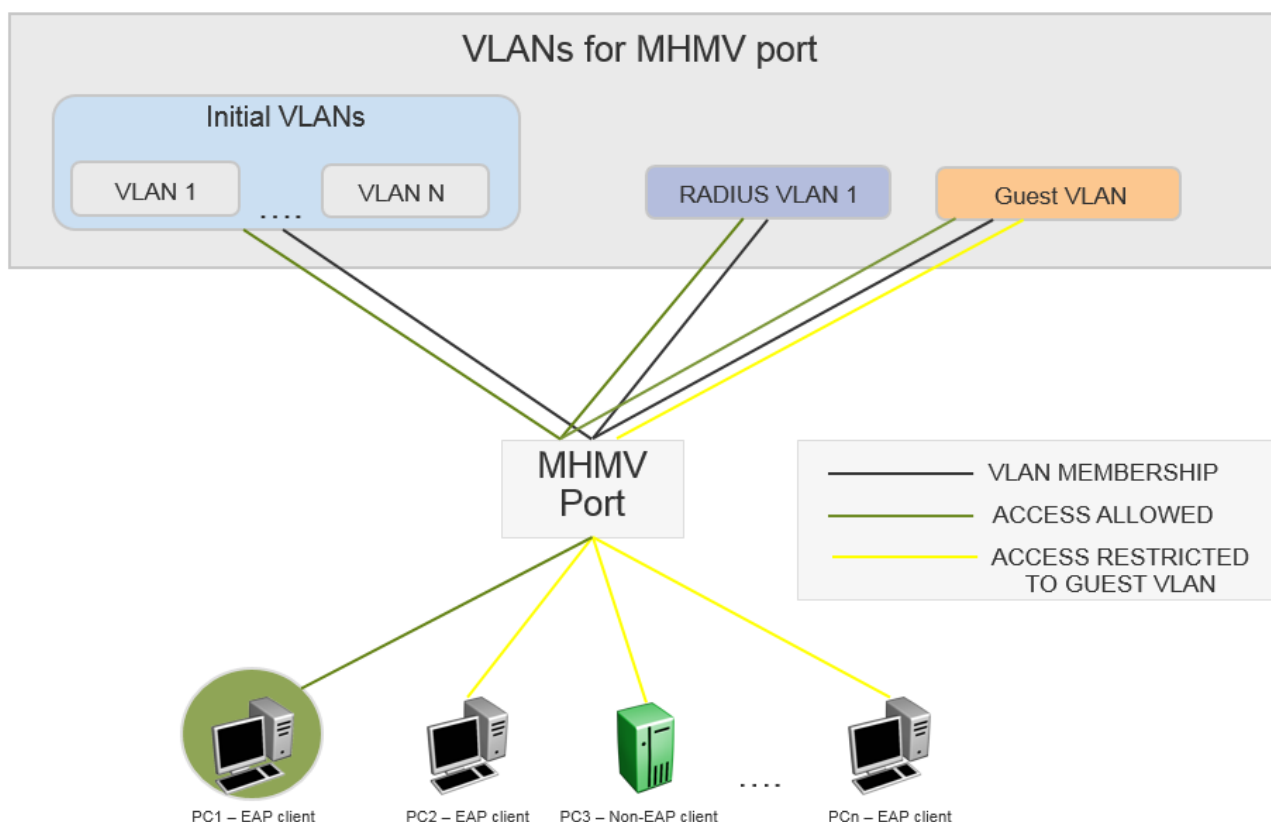


Note

All clients have Guest VLAN access.

- A client behind the MHMV port gets authenticated. For this usage scenario let us consider PC1 as the authenticated client.
- The port default VLAN ID is equal to the Guest VLAN ID and remains unchanged.
- The port is copied into the RADIUS assigned VLAN (if any).
- The untagged traffic that originates from PC1 (identified by MAC address) can access only the RADIUS assigned VLAN or the initial port default VLAN ID, if the RADIUS VLAN attribute is missing.
- The remaining clients that send untagged traffic are unauthenticated devices. The unauthenticated devices can access only the Guest VLAN because the port VLAN ID is equal to the Guest VLAN ID.
- The initial VLANs are accessed by the following devices:
 - Authenticated devices that are missing RADIUS VLAN attributes.
 - Authenticated devices that send corresponding tagged packets.
- When another client gets authenticated, the authenticated client undergoes the same process as PC1.

The following figure represents the functionality when a client gets authenticated:



Note

PC1 is authenticated with RADIUS VLAN 1. The remain clients have guest VLAN access.

When a client disconnects the following happens:

- The MAC VLAN rule is removed from the switch.
- If the RADIUS VLAN attribute was used with the client was authenticated and no other clients are authenticated on that RADIUS VLAN, then the port is removed from the VLAN. If other clients are authenticated on that RADIUS VLAN, then the VLAN MAC rule is deleted.
- If RADIUS VLAN attribute is not used when the client is authenticated, then only the VLAN MAC rule is deleted.

Guest VLAN on a MHSa Port Usage Scenario

The following is a usage example when Guest VLAN is configured with an EAP MHSa port:

- There are no authenticated EAP or NEAP clients on a port.
- The port is removed from the initial VLANs and moved to Guest VLAN ID.
- The default port VLAN ID changes to Guest VLAN ID.
- All MACs seen on the port have Guest VLAN access.

- Port is removed from the Guest VLAN ID.
- If no RADIUS assigned VLAN is present, then the VLAN membership and the default port VLAN ID is restored to default settings.
- If the RADIUS assigned VLAN is present, then the VLAN membership and the default port VLAN ID is changed according to its value.
- Guest VLAN loses its purpose because all MACs are allowed automatically without authentication

In MHSA mode, the Guest VLAN applies only when no authenticated client is present on the port.

Guest I-SID

Guest I-SID support provides limited network access until the client is authenticated. The switches use the Guest I-SID to forward traffic until the client authenticates and receives other VLAN:I-SID bindings from the RADIUS server.

Guest I-SID is a per-port option. You must configure an I-SID either as a C-VLAN or as an ELAN with an associated platform VLAN before you can configure it as the Guest I-SID. After you configure the Guest I-SID and you enable EAP, an untagged S-UNI is created based on the supplied I-SID. When you change the Guest I-SID while EAP is enabled, the untagged S-UNI is replaced on the port.

In MHSA mode, only one untagged S-UNI can exist on a port at one time. Consider the following:

- If there is a manually configured untagged S-UNI on the port, the untagged S-UNI, which uses the Guest I-SID replaces it.
- If the RADIUS server provides an untagged S-UNI after the client is authenticated, it replaces the untagged S-UNI, which was created based on the Guest I-SID.
- If the Guest I-SID is removed, the previous manually configured untagged S-UNI is automatically restored.
- If the RADIUS-assigned untagged S-UNI is no longer present, EAP recreates the untagged S-UNI created based on the Guest I-SID.

In MHMV mode, the untagged S-UNIs provided by the RADIUS server are treated as MAC-based untagged S-UNIs, which are different from the untagged S-UNI on the port. Consider the following factors:

- If there is a manually configured untagged S-UNI on the port, the untagged S-UNI, which uses the Guest I-SID, replaces it.
- If the Fail-Open I-SID and the Guest I-SID are both configured, the Guest I-SID is applied, as long as a RADIUS server is reachable.
- If the RADIUS server becomes unreachable, the untagged S-UNI based on the Fail-Open I-SID is removed and the untagged S-UNI is created based on the Guest I-SID.

EAP and NEAP separation

EAP and NEAP separation provide the ability to have only NEAP clients allowed on one port. This is done by allowing eap-mac-max to be set to 0. This enhancement gives you the ability to disable EAP clients authentication without disabling NEAP clients. There are no additional configuration commands. For more information, see [Configure Maximum EAP Clients](#) on page 1052 and [Configure Maximum NEAP Clients](#) on page 1053.

EAP and NEAP VLAN names

VLAN names configures VLAN membership of EAP and NEAP clients. You do not have to configure this feature as this mode is always enabled by default.

Fail Open VLAN with Continuity Mode

Fail Open VLAN provides network connectivity when the switch cannot connect to a RADIUS server. If an authentication failure occurs that is based on a RADIUS timeout, the port immediately transitions to the Fail Open VLAN.



Note

Prior to releases that support Continuity Mode, transition to the Fail Open VLAN is based on interval-based RADIUS server reachability checks. If the RADIUS server is reachable, the switch continues to check the reachability at a default interval of three minutes. This interval-based check can lead to a transition delay of up to three minutes, from the moment when the RADIUS Server becomes unreachable until the port moves to the Fail Open VLAN.

If the switch cannot connect to the primary and secondary RADIUS servers, then after a specified number of attempts to restore connectivity, the switch declares the RADIUS servers unreachable.

Fail Open VLAN provides the following functionality:

- When the EAP RADIUS servers are not reachable, Fail Open VLAN provides restricted access to devices, which is separate from the Guest VLAN.
- The EAP and NEAP clients are not affected when the RADIUS servers are not reachable.

To use Fail Open VLAN:

- Fail Open VLAN is a per-port configuration.
- Enable Fail Open VLAN by configuring a valid Fail Open VLAN ID and configure the selected VLAN ID on the switch.
- Use only port-based VLANs as Fail Open VLANs.

When you configure Fail Open VLAN on a port and the RADIUS servers are not reachable, then the Fail Open VLAN provides the following functionality:

- The port is removed from Guest VLAN, if configured, but all other VLAN membership is kept and the port is added to the Fail Open VLAN.
- The default VLAN ID is changed to the Fail Open VLAN ID.
- Traffic from the authenticated EAP and NEAP clients are forwarded as before.
- If re-authentication is enabled in Fail Open VLAN mode, then EAP and NEAP clients stop performing re-authentication.
- All new MACs seen on the port are considered as potential EAP and NEAP clients and are granted Fail Open VLAN access.

When at least one RADIUS server recovers, all EAP-enabled ports are removed from the Fail Open VLAN. All unauthenticated MACs are flushed to give the MACs an opportunity to authenticate.

Fail Open VLAN with Guest VLAN scenarios

When an EAP port is configured with both Fail Open VLAN and Guest VLAN, consider the following scenarios:

1. EAP port operating in MHMV mode:
 - If the EAP RADIUS servers are reachable, then all the authenticated clients have Guest VLAN ID access.
 - If the EAP RADIUS servers are not reachable, then Guest VLAN must be removed from the port completely. The Fail Open VLAN is the new default VLAN. All unauthenticated MACs have Fail Open VLAN access.
2. EAP port operating in MHSA mode:
 - Fail Open VLAN has no impact on the Guest VLAN functionality in MHSA mode.

Fail Open I-SID with Continuity Mode

Fail Open I-SID provides network connectivity with restricted access to devices when the switch cannot connect to a RADIUS server. If a failure occurs that is based on a RADIUS timeout, the port immediately transitions to the Fail Open I-SID.



Note

Prior to releases that support Continuity Mode, transition to the Fail Open I-SID is based on interval-based RADIUS server reachability checks. If the RADIUS server is reachable, the switch continues to check the reachability at a default interval of three minutes. This interval-based check can lead to a transition delay of up to three minutes, from the moment when the RADIUS Server becomes unreachable until the port moves to the Fail Open I-SID.



Note

EAP and NEAP clients are not affected when the RADIUS servers are unreachable.

To use Fail Open I-SID:

- Fail Open I-SID is a per-port configuration.
- You must configure an I-SID either as a C-VLAN or as an ELAN with an associated platform VLAN before you can configure it as the Fail-Open I-SID.
- After you configure the Fail Open I-SID and you enable EAP, an untagged S-UNI is created based on the supplied I-SID. When you change the Fail Open I-SID while EAP is enabled, the untagged S-UNI is replaced on the port.

**Note**

Fail Open I-SID is not supported in MHSA mode.

In MHMV mode, the untagged S-UNIs provided by the RADIUS server are treated as MAC-based untagged S-UNIs, which are different from the untagged S-UNIs on the port. Consider the following factors:

- If there is a manually configured untagged S-UNI on the port, the untagged S-UNI, which uses the Fail Open I-SID, replaces it.
- Caution is advised when both Fail-Open I-SID and Guest I-SID are configured. In this scenario, if a RADIUS server becomes reachable, the untagged S-UNI created based on the Fail-Open I-SID is removed and another untagged S-UNI based on the Guest I-SID is created.

EAP Auto-ISID-Offset

EAP auto-isid-offset functionality is used for MACs that do not receive an I-SID attribute from the RADIUS server. The configured I-SID offset value is used to calculate an I-SID value for a Switched UNI (S-UNI) when the switch receives only the VLAN attribute from the RADIUS server (not the FA VLAN:I-SID binding). In that case, the I-SID value is calculated as follows: $I-SID = VLAN\ ID + \text{configured I-SID offset value}$.

Wake On LAN

Wake On LAN (WoL) networking standard enables remotely powering-up a shutdown computer from a sleeping state. In this process, the computer is shutdown with power reserved for the network card. A packet known as Magic Packet is broadcast on the local LAN or subnet. The network card on, receiving the Magic Packet, verifies the information. If the information is valid, the network card powers-up the shutdown computer.

The WoL Magic Packet is a broadcast frame sent over a variety of connectionless protocols, such as UDP. The most commonly used connectionless protocol is UDP. The Magic Packet contains data that is a defined constant represented in hexadecimal as FF:FF:FF:FF:FF:FF, followed by 16 repetitions of the target computer MAC address and possibly by a four or six byte password.

If you implement enhanced network security using 802.1X, the transmission of Magic Packets to sleeping or unauthorized network devices is blocked. You can use an interface specific 802.1X feature known as traffic-control to address this requirement of supporting both WoL and 802.1X Authentication simultaneously. The default mode is

in-out. This mode blocks both ingress and egress unauthenticated traffic on an 802.1X port. Configuring the traffic control mode to *in* enables the transmission of Magic Packets to sleeping or unauthenticated devices. This mode allows any network control traffic, such as a WoL Magic Packet, to be sent to a workstation irrespective of the authentication or sleep status.

**Important**

If a PC client is assigned to a VLAN based on a previous RADIUS Assigned VLAN, when the client goes into sleep or hibernation mode it reverts to either the default port-based VLAN or Guest VLAN configured for that port. Therefore, the WoL Magic Packet must be sent to the default VLAN or Guest VLAN.

EAPoL Configuration Using CLI

EAPoL (EAP) uses RADIUS protocol for EAP-authorized logons. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration.

Before configuring your device, you must configure at least one EAP RADIUS server and shared secret fields.

You cannot configure EAP on ports that are currently configured for:

- Shared segments
- MultiLink Trunking (MLT)

Change the status of each port that you want to be controlled to auto. The auto setting automatically authenticates the port according to the results of the RADIUS server. The default authentication setting for each port is authorized.

Globally Enable EAP on the Device

Enable EAP globally on the switch before you enable it on a port or interface.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Globally configure EAP:

```
eapol enable
```

Configure EAP on an Interface

Configure EAP on an interface.

Before You Begin

- EAP must be globally enabled.

About This Task

When you configure a port with the EAP status of auto (Authorization depends on result of EAP authentication), only one supplicant is allowed on this port. Multiple EAP supplicants are not allowed on the same physical switch port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable EAP on an interface:

```
eapol status {authorized|auto}
```

3. Disable EAP on an interface:

```
no eapol status
```

Examples

Enable EAP on an interface:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface GigabitEthernet 1/2  
Switch:1(config-if)# eapol status auto
```

Disable EAP on an interface:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface GigabitEthernet 1/2  
Switch:1(config-if)# no eapol status
```

Variable Definitions

The following table defines parameters for the **eapol status** command.

Variable	Value
<i>authorized</i>	Specifies that the port is always authorized. The default value is authorized.
<i>auto</i>	Specifies that port authorization depends on the results of the EAP authentication by the RADIUS server. The default value is authorized.

Configure EAP on a Port

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the maximum EAP requests sent to the supplicant before timing out the session:

```
eapol [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/  
all][all]] max-request <1-10>
```

3. Configure the time interval between authentication failure and the start of a new authentication:

```
eapol [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/  
all][all]] quiet-interval <1-65535>
```

4. Enable reauthentication:

```
eapol [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/  
all][all]] re-authentication enable
```

5. Configure the time interval between successive authentications:

```
eapol [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/  
all][all]] re-authentication-period <0, 60-65535>
```

6. Configure the EAP authentication status:

```
eapol [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/  
all][all]] status {authorized|auto}
```

Example

Configure the maximum EAP requests sent to the supplicant before timing out the session:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface GigabitEthernet 1/2  
Switch:1(config-if)#eapol max-request 10  
Switch:1(config-if)#eapol port 1/2 quiet-interval 500
```

Variable Definitions

The following table defines parameters for the **eapol port** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Specifies the port or list of ports used by EAP. Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>max-request <1-10></code>	Specifies the maximum EAP requests sent to the supplicant before timing out the session. The default is 2.
<code>quiet-interval <1-65535></code>	Specifies the time interval in seconds between the authentication failure and start of a new authentication. The default is 60.
<code>re-authentication enable</code>	Enables reauthentication of an existing supplicant at a specified time interval.
<code>re-authentication-period <0, 60-65535></code>	<p>Specifies the time interval, in seconds, between successive authentication. Configure the value to 0 to prevent EAP or NEAP sessions from aging out.</p> <p>Caution: Preventing re-authentication can introduce a security risk.</p> <p>The default is 3600 (1 hour).</p>
<code>status {authorized auto}</code>	Specifies the desired EAP authentication status for this port.

Configure an EAP-Enabled RADIUS Server

Before You Begin

- You must enable EAP globally.

About This Task

The RADIUS server uses the secret key to validate users.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a secret key:

```
radius server host WORD<0-113> used-by eapol key WORD<0-32>
```

3. (Optional) Configure accounting parameters for the RADIUS server:
 - a. Configure the UDP port:

```
radius server host WORD<0-113> used-by eapol acct-port <1-65536>
```
 - b. Enable accounting:

```
radius server host WORD<0-113> used-by eapol acct-enable
```
4. (Optional) Configure secure options for the RADIUS server:
 - a. Enable RADIUS Security (RADSec):

```
radius server host WORD<0-113> used-by eapol secure-enable
```
 - b. Select the protocol for secure mode:

```
radius server host WORD<0-113> used-by eapol secure-mode <dtls |  
tls>
```
 - c. Specify a secure profile for the server:

```
radius server host WORD<0-113> used-by eapol secure-profile  
WORD<1-16>
```
 - d. Configure the log level:

```
radius server host WORD<0-113> used-by eapol secure-log-level  
<critical | debug | error | info | warning>
```
5. (Optional) Configure other optional parameters for the RADIUS server:
 - a. Configure the UDP port:

```
radius server host WORD<0-113> used-by eapol port <1-65536>
```
 - b. Configure the server priority:

```
radius server host WORD<0-113> used-by eapol priority <1-10>
```
 - c. Configure the maximum number of retries:

```
radius server host WORD<0-113> used-by eapol retry <0-6>
```
 - d. Configure the server timeout:

```
radius server host WORD<0-113> used-by eapol timeout <1-180>
```
6. Enable the server:

```
radius server host WORD<0-113> used-by eapol enable
```

Example

Add an EAP RADIUS server:

```
Switch:1>enable  
Switch:1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#radius server host fe80:0:0:0:21b:4fff:fe5e:73fd key radiustest used-  
by eapol enable
```

Variable Definitions

The following table defines parameters for the **radius server host** command.

Variable	Value
<code>acct-enable</code>	Enables accounting for the server.
<code>acct-port <1-65536></code>	Specifies the accounting port. The default is 1813.
<code>enable</code>	Enables the RADIUS server host.
<code>host WORD<0-113></code>	Configures a host server. <i>WORD<0-113></i> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.
<code>key WORD<0-32></code>	Specifies the secret key.
<code>port <1-65535></code>	Specifies the port ID number. The default is 1812.
<code>priority <1-10></code>	Specifies the priority. The lowest number is the highest priority. The default is 10.
<code>retry <0-6></code>	Specifies the retry count. The default is 1.
<code>secure-enable</code>	Enables secure mode on the server. The default is disabled.
<code>secure-log-level <critical debug error info warning></code>	Specifies the RADIUS secure server log severity level. The default is error.
<code>secure-mode <dtls tls></code>	<p>Specifies the protocol for establishing the secure connection with the server. The possible values are:</p> <ul style="list-style-type: none"> • <code>dtls</code> - Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP) • <code>tls</code> - Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP) <p>The default is TLS.</p> <p>Important: To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the <code>radsecproxy.conf</code> configuration file:</p> <pre> tls default{ ... TlsVersion TLS1_2 } </pre>
<code>secure-profile WORD<1-16></code>	Specifies the secure profile name.
<code>timeout <1-180></code>	Specifies the timeout of the server. The default is 8.

Configure the Switch for EAP and RADIUS

Perform the following procedure to configure the switch for EAP and RADIUS.

About This Task

You must configure the switch, through which user-based-policy (UBP) users connect to communicate with the RADIUS server to exchange EAP authentication information, as well as user role information. You must specify the IP address of the RADIUS server, as well as the shared secret (a password that authenticates the device with the RADIUS server as an EAP access point). You must enable EAP globally on each device, and you must configure EAP authentication on each device port, through which EAP/UBP users connect.

For more information about EPM and UBP, see the user documentation for your Enterprise Policy Manager (EPM) application.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create a RADIUS server that is used by EAP:
`radius server host WORD<0-113> key WORD<0-20> used-by eapol enable`
3. Log on to the Interface Configuration mode:
`interface vlan <1-4059>`
4. Enable the device to communicate through EAP:
`eapol enable`
5. Exit from VLAN interface mode:
`exit`
6. Enter Interface Configuration mode:
`interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}`
7. Enable device ports for EAP authentication:
`eapol port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} status auto`
8. Enable periodic supplicant re-authenticating:
`eapol port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} re-authentication enable`

Example

Create a RADIUS server that is used by EAP and enable the device to communicate through EAP:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#radius server host fe90:0:0:0:21b:4eee:fe5e:75fd key radiustest used-by
eapol enable
Switch:1(config)#interface vlan 2
Switch:1(config-if)#eapol enable
```

Variable Definitions

The following table defines parameters for the **radius server host WORD<0-113>** used by **eapol** command.

Variable	Value
<i>host WORD<0-113></i>	Configures a host server. <i>WORD<0-113></i> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS. This address tells the device where to find the RADIUS server, from which it obtains EAP authentication and user role information.
<i>key WORD<0-20></i>	Specifies the shared secret key that you use for RADIUS authentication. The shared secret is held in common by the RADIUS server and all EAP-enabled devices in your network. It authenticates each device with the RADIUS server as an EAP access point. When you configure your RADIUS server, you must configure the same shared secret value as you specify here.

Change the Authentication Status of a Port

The switch authorizes ports by default, which means that the ports are always authorized and are not authenticated by the RADIUS server.

You can also make the ports controlled so that they are dependent on being authorized by the Radius Server when you globally enable EAP (auto).

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the authorization status of a port:

```
eapol status {authorized|auto}
```

Example

Configure the authorization status of a port:

```
Switch:1>enable  
Switch:1#configure terminal
```

```
Switch:1(config)#interface GigabitEthernet 1/1
Switch:1(config-if)#eapol status auto
```

Variable Definitions

The following table defines parameters for the **eapol status** command.

Variable	Value
<i>authorized</i>	Specifies that the port is always authorized. The default value is authorized.
<i>auto</i>	Specifies that port authorization depends on the results of the EAP authentication by the RADIUS server. The default value is authorized.

Delete an EAP-Enabled RADIUS Server

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Delete an EAP-enabled RADIUS server:

```
no radius server host WORD<0-113> used-by eapol
```

Variable Definitions

The following table defines parameters for the **radius server host WORD<0-113> usedby eapol** command.

Variable	Value
<i>host WORD<0-113></i>	Configures a host server. <i>WORD<0-113></i> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS. This address tells the device where to find the RADIUS server, from which it obtains EAP authentication and user role information.
<i>key WORD<0-20></i>	Specifies the shared secret key that you use for RADIUS authentication. The shared secret is held in common by the RADIUS server and all EAP-enabled devices in your network. It authenticates each device with the RADIUS server as an EAP access point. When you configure your RADIUS server, you must configure the same shared secret value as you specify here.

Configure Fail Open VLAN

About This Task

Use this procedure to configure Fail Open VLAN.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Fail Open VLAN:

```
eapol fail-open-vlan <1-4059>
```

Example

Configure the Fail Open VLAN.

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/1  
Switch:1(config)#eapol fail-open-vlan 10
```

Variable Definitions

The following table defines parameters for the **eapol fail-open-vlan** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Display the Current EAP-Based Security Status

Use the following procedure to display the status of the EAP-based security.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the current EAP-based security status:

- show eapol auth-stats interface [gigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
- show eapol port {interface [gigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] | {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}
- show eapol session-stats interface [gigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
- show eapol sessions {eap | neap} [vlan <1-4059>] [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] [verbose]
- show eapol summary port [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
- show eapol system
- show eapol sessions eap verbose
- show eapol sessions neap verbose

Examples

```
Switch:>enable
Switch:1#config terminal
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#show eapol port 1/2
```

```
=====
Eapol Configuration
=====
```

PORT NUM	STATUS	OPER MODE	DYN MHSA	Flex-UNI ENABLE	MAX REQ	QUIET INTVL	NON-EAP ENABLE	LLDP-AUTH ENABLE	MAX MAC	MAX EAP	MAX NEAP	ADMIN TRAFFIC CONTROL	OPER TRAFFIC CONTROL	TRAFFIC CONTROL ORIGIN	ORIGIN
1/2	Auth	MHVM	false	true	2	60	false	false	2	2	2	in-out	in-out	CONFIG	CONFIG

```
=====
Eapol Configuration
=====
```

PORT NUM	REAUTH ENABLE	REAUTH ORIGIN	REAUTH PERIOD	REAUTH PERIOD ORIGIN	GST VLAN	GST I-SID	FAIL VLAN	FAIL I-SID	COA ENABLE	ORIGIN
1/2	false	CONFIG	3600	CONFIG	N/A	N/A	N/A	N/A	false	CONFIG

```
Switch:>enable
Switch:1#config terminal
Switch:1(config)#show eapol sessions eap verbose
```

```
=====
Eap Oper Status Verbose
=====
```

PORT NUM	MAC	PAE STATUS	VLAN ID	PRI	Flex-UNI Enable	I-SID SOURCE	VLAN:I-SID	ACL	ACEs	DYN MHSA	RADIUS DYNAMIC SETTINGS
1/10	00:00:00:00:00:11	authenticated	20	0	true	radius 0:20				FALSE	IGMPSNOOP
1/14	00:00:00:00:00:01	authenticated	10	0	true	radius 0:10				FALSE	

```
Total Number of NEAP Sessions: 2
```

```
=====
Eapol Configuration
=====
```

PORT NUM	MAC	DYNAMIC CREATE	VLAN PV	ATTRIBUTES SV	I-SID	EV	VLAN NAME	I-SID NAME	MVPN	I-SID	IGMP SNOOP QUERIER ADDR
1/10	00:00:00:00:00:11	vlan	20	N/A	20	0	N/A	N/A	N/A		192.0.2.1
1/14	00:00:00:00:00:01	vlan	10	N/A	10	0	N/A	N/A	10010		192.0.2.25

```
Switch:1(config)#show eapol sessions neap verbose
```

```
=====
Non-Eap Oper Status Verbose
=====
```

PORT	MAC	STATE	VLAN	PRI	Flex-UNI	I-SID	NON-EAP	VLAN:I-SID	ACL	ACEs	DYN	RADIUS DYNAMIC
------	-----	-------	------	-----	----------	-------	---------	------------	-----	------	-----	----------------

NUM				ID	Enable	SOURCE	AUTH		MHSA	SETTINGS
1/10	00:00:00:00:00:11	authenticated	20	0	true	radius	radius	0:20	FALSE	IGMPSNOOP
1/14	00:00:15:7c:90:b6	authenticated	10	0	true	radius	radius	0:10	FALSE	
Total Number of NEAP Sessions: 2										
=====										
=====										
PORT	MAC	DYNAMIC VLAN ATTRIBUTES								
NUM		CREATE	PV	SV	I-SID	EV	VLAN NAME	I-SID NAME	MVPN I-SID	IGMP SNOOP QUERIER ADDR
=====										
1/10	00:00:00:00:00:11	vlan	20	N/A	20	0	N/A	N/A	N/A	192.0.2.1
1/14	00:00:15:7c:90:b6	vlan	10	N/A	10	0	N/A	N/A	10010	192.0.2.25
=====										
Switch:1>show eapol system										
=====										
Eapol System										
=====										
eap : disabled										
Eapol Version : 3										
non-eap-pwd-fmt : mac-addr										
non-eap-pwd-fmt key : *****										
non-eap-pwd-fmt padding : disabled										
auto-isid-offset status : disabled										
auto-isid-offset value : 1000										

Variable Definitions

The following table defines parameters for the **show eapol** command.

Variable	Value
<i>auth-stats</i> <i>[gigabitEthernet {slot/ port[/sub-port] [-slot/ port[/sub-port]] [, ...] [slot/all] [all]}]</i>	Displays the authentication statistics interface. Note: <i>auth-stats</i> <i>[gigabitEthernet {slot/port[/sub-port] [-slot/ port[/sub-port]] [, ...] [slot/all] [all]}]</i> is useful only for EAP supplicants. The command output changes only when the EAP supplicant tries to access the network.
<i>port {interface</i> <i>[gigabitEthernet {slot/ port[/sub-port] [-slot/ port[/sub-port]] [, ...] [slot/all] [all]}]}</i> <i>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	Specifies the ports to display. If no port is entered, all ports are displayed.
<i>session-stats interface</i> <i>[gigabitEthernet {slot/ port[/sub-port] [-slot/ port[/sub-port]] [, ...] [slot/all] [all]}]</i>	Displays the authentication session statistics interface.

Variable	Value
<i>sessions {eap neap} [vlan<1-4059>] [{slot/ port[/sub-port] [-slot/ port[/sub-port]] [,...] [slot/all] [all]}] [verbose]</i>	Displays EAP and non-EAP authentication sessions on the port.
<i>summary port[{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}]</i>	Displays EAP and NEAP clients.
<i>system</i>	Displays EAP settings.

Display Port VLAN Information

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the port VLAN information:
show interfaces gigabitEthernet vlan <1-4059>

Example

```
Switch:1>enable
Switch:1#show interfaces gigabitEthernet vlan
=====
Port Vlans
=====
```

PORT NUM	TAGGING	DISCARD TAGFRAM	DISCARD UNTAGFRAM	DEFAULT VLANID	VLAN IDS	PORT TYPE	UNTAG DEFVLAN	DYNAMIC VLANS	UNTAG VLANS	FA/ Flex UNI
1/1	disable	false	false	1	1	normal	disable	P	1	disable
1/2	enable	false	false	1	1,3,10	normal	disable	P	1,10	disable
1/3	enable	false	false	1	1,10,20	normal	disable	P		disable

```
-----
DYNAMIC VLAN Legend:
P=Protocol enabled.
```

Variable Definitions

The following table defines parameters for the **show interfaces** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/ port[/sub-port]] [, ...] [slot/ all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/ port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Configuring the Format of the RADIUS Password Attribute when Authenticating NEAP MAC Addresses using RADIUS

Use the following procedure to configure the format of the RADIUS password when authenticating NEAP MAC addresses using RADIUS.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Configure the RADUIS password format:

```
eapol multihost non-eap-pwd-fmt {[ip-addr] [key WORD<1-32>] [mac-addr]  
[padding] [port-number]}
```

Variable Definitions

The following table defines parameters for the **eapol multihost non-eap-pwd-fmt** command.

Variable	Value
<code>ip-addr</code>	Management ip-address of the switch.
<code>key WORD<1-32></code>	Key value used for non-eap password format.

Variable	Value
<i>mac-addr</i>	Mac-Address of the client.
<i>padding</i>	Includes a dot in the RADIUS password for every missing parameter.
<i>port-number</i>	Index of the port on which MAC is received.

**Note**

To derive the port number for an interface, use the command **show interfaces gigabitEthernet {slot/port[-slot/port][,...][slot/all][all]}** .

If you configure interface 1/6 on the product, to derive the port-number for this interface, use the command **show interfaces gigabitEthernet 1/6**. From this command, you can ascertain that port number used in the NEAP password is 197.

```
Switch:1(config)# show interfaces gigabitEthernet 1/6
```

Port Interface								
PORT NUM	INDEX	DESCRIPTION	LINK TRAP	PORT LOCK	MTU	PHYSICAL ADDRESS	STATUS ADMIN OPERATE	
1/6	197	1000BaseTX	true	false	1950	f8:15:47:e1:dd:05	up	up

Enable RADIUS Authentication of NEAP Hosts on EAP Enabled Ports

For RADIUS authentication of NEAP hosts on EAP-enabled ports, you must enable EAP globally on the switch and then enable NEAP hosts on the local interface.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable RADIUS authentication of NEAP hosts on the local interface:

```
eapol multihost radius-non-eap-enable
```

Configure the Maximum MAC Clients

Use this procedure to configure the maximum EAP and NEAP MAC clients supported on a port.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Set the maximum limit of allowed EAP and NEAP MAC clients supported on the port:

```
eapol multihost mac-max <1-8192>
```

Example

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)# interface GigabitEthernet 1/16  
Switch:1(config-if)# eapol multihost mac-max <1-8192>
```

Variable Definitions

The following table defines parameters for the **eapol multihost mac-max** command.

Variable	Value
<i>mac-max</i> <1-8192>	Specifies the maximum number of EAP and NEAP MAC addresses allowed on the port. The maximum limit is 8192 MAC addresses.

Configure Maximum EAP Clients

About This Task

Use this procedure to configure the maximum EAP clients allowed on the port at one time.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the maximum EAP clients:

```
eapol multihost eap-mac-max <0-32>
```



Note

eap-mac-max is also used to provide EAP and NEAP separation functionality. By default the EAP clients are enabled per port and *eap-mac-max* limit is 2. If *eap-mac-max* is set to 0 then EAP client authentication is disabled.

Example

Configure the maximum EAP clients allowed on the port at one time.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config)#eapol multihost eap-mac-max 10
```

Variable Definitions

The following table defines parameters for the **eapol multihost eap-mac-max** command.

Variable	Value
<0-32>	Specifies the maximum EAP clients allowed on the port at one time. The default is 2.

Configure Maximum NEAP Clients

About This Task

Use this procedure to configure the maximum NEAP clients allowed on the port at one time.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the maximum NEAP clients:

```
eapol multihost non-eap-mac-max <0-8192>
```



Note

non-eap-mac-max is also used to provide EAP and NEAP separation functionality. By default the NEAP clients are enabled per port and *non-eap-mac-max* limit is 2. If *non-eap-mac-max* is set to 0 then NEAP client authentication is disabled.

Example

Configure the maximum NEAP clients allowed on the port at one time.

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/1  
Switch:1(config)#eapol multihost non-eap-mac-max 10
```

Variable Definitions

The following table defines parameters for the **eapol multihost non-eap-mac-max** command.

Variable	Value
<0-8192>	Specifies the maximum NEAP clients allowed on the port at one time. The default is 2.

Configure the Guest VLAN ID

About This Task

Use this procedure to configure the Guest VLAN ID.

Procedure

- 1. Enter GigabitEthernet Interface Configuration mode:
`enable`

`configure terminal`

`interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- 2. Configure the Guest VLAN ID:
`eapol guest-vlan <1-4059>`

Example

Configure the Guest VLAN ID.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config)#eapol guest-vlan 10
```

Variable Definitions

The following table defines parameters for the **eapol guest-vlan** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Clear an EAP Client Session

About This Task

Perform this procedure to clear an EAP client session on the switch.

Procedure

- 1. Enter Privileged EXEC mode:
`enable`

2. Clear the EAP client session.

Use one of the following options:

a. Clear by MAC address:

```
clear eapol eap [<0x00:0x00:0x00:0x00:0x00:0x00>]
```

b. Clear by port:

```
clear eapol eap {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]  
[slot/all] [all]} [<0x00:0x00:0x00:0x00:0x00:0x00>]
```

Example

Clear an EAP client session by MAC address:

```
Switch:1>enable  
Switch:1#clear eapol eap 00:00:00:00:00:01
```

Clear an EAP client session by port:

```
Switch:1>enable  
Switch:1#clear eapol eap 1/14 00:00:00:00:00:01
```

Clear the NEAP Session

Use this procedure to clear the NEAP session that is learned on the switch.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear the NEAP client session.

Use one of the following options:

a. Clear by MAC address:

```
clear eapol non-eap [<0x00:0x00:0x00:0x00:0x00:0x00>]
```

b. Clear by port:

```
clear eapol non-eap {slot/port[/sub-port] [-slot/port[/sub-port]]  
[, ...] [slot/all] [all]} [<0x00:0x00:0x00:0x00:0x00:0x00>]
```

Example

Clear a non-EAP client session by MAC address:

```
Switch:1> enable  
Switch:1#clear eapol non-eap 00:1b:63:84:45:e6
```

Clear a non-EAP client session by port:

```
Switch:1> enable  
Switch:1#clear eapol non-eap 1/16 00:1b:63:84:45:e6
```

Variable Definitions

The following table defines parameters for the **clear eapol non-eap** command.

Variable	Value
<code>{slot/port[-slot/port] [,...][slot/all][all]}</code>	Specifies the port list on which the NEAP MAC is learnt.
<code>0x00:0x00:0x00:0x00:0x00:0x00</code>	Specifies the MAC-Address on the NEAP session.

Configure EAP Operational Mode

About This Task

Use this procedure to configure the EAP operational mode.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. configure the EAP operational mode:

```
eapol multihost eap-oper-mode {mhm | mhsa}
```



Note

The default EAP operational mode is MHMV.

Example

Configure the EAP operational mode:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/1  
Switch:1(config)#eapol eap-oper-mode mhsa
```

Variable Definitions

The following table defines parameters for the **eapol multihost eap-oper-mode** command.

Variable	Value
<i>mhm</i>	Specifies the EAP operational mode as Multiple Host Multiple VLAN.
<i>mhs</i>	Specifies the EAP operational mode as Multiple Host Single Authentication.

Enable Dynamic Changes to EAP Sessions on a Port

Before You Begin

You must enable EAP globally and at the port level.

About This Task

Configure a port to allow dynamic changes to EAP sessions. The default is enable.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable RADIUS dynamic authorization server processing requests.

```
eapol radius-dynamic-server enable
```

Example

```
Switch:1>enable  
Switch:1#config terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#interface gigabitEthernet 1/4  
Switch:1(config-if)#eapol radius-dynamic-server enable
```

Configure EAP auto-isid-offset

Before You Begin

- Enable EAP globally or on the port.

Procedure

- 1. Enter Global Configuration mode:
`enable`
`configure terminal`
- 2. Configure an I-SID offset value:
`eapol auto-isid-offset <0-15995903>`
- 3. Enable EAP globally:
`eapol enable`
- 4. Confirm that your configuration is correct:
`show eapol system`

Examples

Configure an I-SID offset value and enable I-SID offset globally on the switch:

```
Switch:1> enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#eapol auto-isid-offset 1000
Switch:1(config)#eapol enable
```

View the current device configuration:

```
Switch:1>show eapol system
=====
                        Eapol System
=====
                eap : disabled
        Eapol Version : 3
        non-eap-pwd-fmt : mac-addr
        non-eap-pwd-fmt key : *****
        non-eap-pwd-fmt padding : disabled
        auto-isid-offset status : disabled
        auto-isid-offset value : 1000
```

Variable Definitions

The following table defines parameters for the **eapol auto-isid-offset** command.

Variable	Value
<0-15995903>	Specifies the auto I-SID offset value. The default is 15995903.
enable	Enables auto I-SID offset. The default is disabled.

Configure the Guest I-SID

Before You Begin

Configure a platform VLAN and associate the Guest I-SID. 0 indicates that Guest I-SID is not enabled for this port.

About This Task

Use this procedure to configure the Guest I-SID.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the Guest I-SID:

```
eapol guest-isid <1-16000000>
```

Example

Configure the Guest I-SID.

```
Switch:1>enable  
Switch:1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#interface gigabitEthernet 1/10  
Switch:1(config-if)#eapol guest-isid 1000
```

Variable Definitions

The following table defines parameters for the **eapol guest-isid** command.

Variable	Value
<0-16000000>	Specifies the Guest I-SID value. 0 indicates that Guest I-SID is not enabled for this port.

Configure Fail Open I-SID

Before You Begin

Configure a platform VLAN and associate the Fail Open I-SID.

About This Task

Use this procedure to configure Fail Open I-SID. If the switch declares the RADIUS servers unreachable, then all new devices gain access into the configured Fail Open I-SID. 0 indicates that Fail Open I-SID is not enabled for this port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Fail Open I-SID:

```
eapol fail-open-isid <0-16000000>
```

Example

Configure the Fail Open I-SID.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config-if)#interface gigabitEthernet 1/10
Switch:1(config-if)#eapol fail-open-isid 1000
```

Variable Definitions

The following table defines parameters for the **eapol fail-open-isid** command.

Variable	Value
<0-16000000>	Specifies the Fail Open I-SID value. 0 indicates that Fail Open I-SID is not enabled for this port.

Configure Wake-On-LAN

Use the following procedure to configure Wake-On-LAN functionality.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Wake-On-LAN:

```
eapol traffic-control <in | in-out>
```

Variable Definitions

The following table defines parameters for the **eapol traffic-control** command.

Variable	Value
<i>in</i>	Specifies incoming traffic is blocked when there is no authenticated device.
<i>in-out</i>	Specifies incoming and outgoing traffic is blocked when there is no authenticated device. The default value is in-out.

Show the EAPoL Status of the Device

Display the current device configuration.



Note

Use the **clear-stats** command to clear EAP or NEAP statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the current device configuration by using the following command:

```
show eapol system
```

Example

```
Switch:1>show eapol system
=====
                        Eapol System
=====
                        eap : disabled
                        Eapol Version : 3
                        non-eap-pwd-fmt : mac-addr
                        non-eap-pwd-fmt key : *****
                        non-eap-pwd-fmt padding : disabled
```

```
auto-isid-offset status : disabled
auto-isid-offset value  : 1000
```

Show EAPoL Authenticator Statistics

Display the authenticator statistics to manage network performance.



Note

Use the **clear-stats** command to clear EAP or NEAP statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the authenticator statistics:

```
show eapol auth-stats interface [gigabitEthernet [{slot/port[/sub-
port]}[-slot/port[/sub-port]][,...]]]
```

Example

```
Switch:1#show eapol auth-stats interface
```

```
=====
                        Eap Authenticator Statistics
=====
```

PORT	EAP RCVD	AUTH-EAP TX	START RCVD	LOGOFF RCVD	INVALID FRAMES	LENGTH ERROR	LAST-RX VER	LAST-RX SRC
1/1	716	1074	0	0	0	0	1	18:a9:05:b1:04:ce
1/2	0	0	0	0	0	0	0	00:00:00:00:00:00
1/3	0	0	0	0	0	0	0	00:00:00:00:00:00
1/4	0	5	0	0	0	0	0	00:00:00:00:00:00
1/5	0	0	0	0	0	0	0	00:00:00:00:00:00
1/6	0	0	0	0	0	0	0	00:00:00:00:00:00
1/7	0	0	0	0	0	0	0	00:00:00:00:00:00
1/8	0	0	0	0	0	0	0	00:00:00:00:00:00
1/9	0	0	0	0	0	0	0	00:00:00:00:00:00
1/10	0	0	0	0	0	0	0	00:00:00:00:00:00

```
--More-- (q = quit)
```

Variable Definitions

Use the data in the following table to use the **show eapol auth-stats interface** command.

Variable	Value
<code>{slot/port[/sub-port]}[-slot/port[/sub-port]][,...] [slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View EAPoL Session Statistics

View EAPoL session statistics to manage network performance.



Note

Use the **clear-stats** command to clear EAP/NEAP statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the session statistics:

```
show eapol session-stats interface [gigabitEthernet [{slot/port[/sub-
port] [-slot/port[/sub-port]] [,...]]]
```

Example

```
Switch:1#show eapol session-stats interface
=====
Eap Authenticator Session Statistics
=====
```

PORT NUM	MAC	SESSION ID	AUTHENTIC METHOD	SESSION TIME	TERMINATE CAUSE	USER NAME
1/1	18:a9:05:b1:04:ce	cb000000	remote-server	0 day(s), 05:58:16	not-terminated	sachin
1/4	00:00:00:00:00:01	cb000002	remote-server	0 day(s), 05:48:01	not-terminated	000000000001

```
=====
```

Variable Definitions

Use the data in the following table to use the **show eapol session-stats interface** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View Non-EAPoL MAC Information

Use this procedure to view non-EAPoL client MAC information on a port.



Note

Use the **clear-stats** command to clear EAP and NEAP statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the non-EAPoL MAC information:

```
show eapol multihost non-eap-mac status [vlan <1-4059>] [verbose]
[{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}]
```

Example

```
Switch:1(config)#show eapol multihost non-eap-mac status
```

```
=====
Eap Oper Status
=====
PORT  MAC              STATE      VLAN  PRI  Flex-UNI  I-SID      NON-EAP  VLAN:I-SID
NUM   ID                ID         Enable   SOURCE   AUTH      AUTH
-----
1/10  00:00:00:00:00:0e authenticated 10    2    true      autoconfig radius  0:1015
-----
```

```
Total Number of EAP Sessions: 1
```

```
Switch(config-if)#show eapol multihost non-eap-mac status verbose
```

```
=====
Non-Eap Oper Status Verbose
=====
PORT  MAC              STATE      VLAN  PRI  Flex-UNI  I-SID      NON-EAP  VLAN:I-SID  ACL  ACES
NUM   ID                ID         Enable   SOURCE   AUTH      AUTH
-----
1/11  00:00:00:00:01:02 authenticated  N/A  0    true      radius      radius  0:10555
-----
```

```
Total Number of NEAP Sessions: 1
```

Variable Definitions

Use the data in the following table to use the **show eapol multihost non-eap-mac status** command.

Variable	Value
<i>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i><1-4059></i>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p>
<i>verbose</i>	Displays non-EAPoL client MAC information.

View Port EAPoL Operation Statistics

Use this procedure to view port EAPoL operation statistics.



Note

Use the **clear-stats** command to clear EAP/NEAP statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the port EAPoL operation statistics information:

```
show eapol status interface [gigabitEthernet [{slot/port[/sub-port]][-
slot/port[/sub-port]][,...][slot/all][all]] [vlan <1-4059>] [verbose]
```

Examples

```
Switch:1(config)#show eapol status interface
```

```
=====
                        Eap Oper Status
=====
PORT  MAC                STATE          VLAN  PRI  Flex-UNI  I-SID  VLAN:I-SID
NUM   NUM                  ID            ID      Enable  SOURCE  AUTH
-----
1/1   18:a9:05:b1:04:ce  authenticated  10     2    true      radius  0:1015
-----
Total Number of EAP Sessions: 1
```

```
Switch:1(config-if)#show eapol status interface gigabitEthernet 2/1,2/2 verbose
```

```
=====
                        Eap Oper Status Verbose
=====
PORT  MAC                PAE    VLAN  PRI  Flex-UNI  I-SID  VLAN:I-SID  ACL  ACEs
NUM   NUM                  STATUS ID      Enable  SOURCE
-----
2/1   00:00:02:eb:34:e9  authenticated  N/A  0    true      radius  111:10111  2    2,3
2/1   00:00:02:eb:34:eb  authenticated  N/A  0    true      radius  111:10111  2    1
2/1   00:00:02:eb:34:ed  authenticated  333  0    true      n/a     111:10111  2    4
2/1   00:00:02:eb:34:ef  authenticated  444  0    true      n/a     111:10111  2    5
2/2   00:00:00:06:a9:df  authenticated  11   0    true      n/a     111:10111  1    1
2/2   00:00:00:06:a9:e1  authenticated  N/A  0    true      radius  111:10111  1    2
2/2   00:00:00:06:a9:e3  authenticated  33   0    true      n/a     111:10111  1    2
2/2   00:00:00:06:a9:e5  authenticated  44   0    true      n/a     111:10111  1    3
-----
Total Number of EAP sessions : 8
```

Variable Definitions

Use the data in the following table to use the **show eapol status** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code><1-4059></code>	Specifies the VLAN ID for which to show the statistics.
<code>verbose</code>	Displays detailed statistics information.

EAP Configuration Using Enterprise Device Manager

EAPoL (EAP) uses RADIUS protocol for EAP-authorized logons. RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all but the following case. When adding a RADIUS server in Enterprise Device Manager (EDM) or modifying a RADIUS configuration in EDM, you must specify if the address type is an IPv4 or an IPv6 address.

Before You Begin

- Before configuring your device, you must configure at least one EAP RADIUS server and shared secret fields.
- You cannot configure EAP on ports that are currently configured for:
 - Shared segments
 - MultiLink Trunking (MLT)
- Change the status of each port that you want to be controlled to auto. For more information on changing the status, see [Configure EAP on a Port](#) on page 1068. The auto setting automatically authenticates the port according to the results of the RADIUS server. The default authentication setting for each port is force-authorized.

Globally Configure EAP on the Server

About This Task

Globally enable or disable EAP on the switch. By default, EAP is disabled.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **802.1X - EAPOL**.
3. Select the **Global** tab.
4. From the AccessControl options, select **enable**.
5. (Optional) Select the appropriate **NonEapRadiusPwdAttrFmt** check boxes to configure the format of the RADIUS password when authenticating non-EAP MAC addresses using RADIUS.

6. (Optional) Enter the key string in **NonNonEapRadiusPwdAttrkeystring**.
7. (Optional) Select **ClearNonEap** to clear the NEAP session that is learned on the switch.
8. (Optional) Type an I-SID offset number in **AutolsidOffset**.
9. (Optional) Select **AutolsidOffsetEnable** to enable Auto I-SID offset on the switch.
10. Select **Apply**.

Global Field Descriptions

Use the data in the following table to use the **Global** tab.

Name	Description
EapolVersion	Displays the EAP version on the switch.
AccessControl	Enables system authentication control. EAP is disabled by default.
NonEapRadiusPwdAttrFmt	Specifies the password attribute format for non EAP RADIUS authentication. <ul style="list-style-type: none"> ipAdd: Specifies IP address. macAddr: Specifies MAC address. portNumber: Specifies port number padding: Specifies padding.
NonEapRadiusPwdAttrKeyString	Specifies the attribute key string for non EAP RADIUS password.
ClearNonEap	Clears the NEAP session that is learned on the switch.
AutolsidOffset	Specifies the Auto I-SID Offset value.
AutolsidOffsetEnable	Enables or disables the Auto I-SID Offset feature.

Configure EAP on a Port

About This Task

Configure EAP or change the authentication status on one or more ports.

Ports are force-authorized by default. Force-authorized ports are always authorized and are not authenticated by the RADIUS server. You can change this setting so that the ports are always unauthorized.

Procedure

1. In the Device Physical View tab, select the port you need to configure.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **EAPOL** tab.
5. (Optional) Select **AllowNonEapHost** to enable hosts that do not participate in 802.1X authentication to get network access.

6. Select the **Status** option as **auto** or **forceAuthorized**.
7. In **MultiHostMaxClients**, type the maximum limit of allowed EAP and NEAP clients supported on this port.
8. In **GuestVlanId**, type the VLAN ID to be used as a Guest VLAN ID.
9. In **FailOpenVlanId**, type the Fail Open VLAN ID.
10. In **NonEapMaxClients**, type the maximum number NEAP authentication MAC addresses allowed on this port.
11. In **EapMaxClients**, type the maximum number of EAP authentication MAC addresses allowed on this port.
12. Select **MultiHostSingleAuthEnabled** to automatically authenticate NEAP MAC addresses on this port.
13. In **PortGuestIsid**, type the I-SID to be used as a Guest I-SID.
14. In **FailOpenIsid**, type the Fail Open I-SID.
15. Select the **AdminTrafficControl** option as **inOut** or **in**.
16. (Optional) Select **LldpAuthEnabled** to enable LLDP authentication for network access.
17. Select **ReAuthEnabled**.
18. In **QuietPeriod**, type the time interval.
19. In **ReauthPeriod**, type the time between re-authentication.
20. In **RetryMax**, type the number of times.
21. Select **Apply**.

EAPoL Field Descriptions

Use the data in the following table to use the **EAPoL** tab.

Name	Description
PortCapabilities	Displays the capabilities of the Port Access Entity (PAE) associated with the port. This parameter indicates whether Authenticator functionality, supplicant functionality, both, or neither, is supported by the PAE of the port. The following capabilities are supported by the PAE of the port: <ul style="list-style-type: none"> • authImplemented: A Port Access Controller Protocol (PACP) Extensible Authentication Protocol (EAP) authenticator functions are implemented. • virtualPortsImplemented: Virtual Port functions are implemented.
PortVirtualPortsEnable	Displays the status of the Virtual Ports function for the real port as True or False.
PortCurrentVirtualPorts	Displays the current number of virtual ports running in the port
PortAuthenticatorEnable	Displays the status of the Authenticator function in the Port Access Entity (PAE) as True or False.

Name	Description
PortSupplicantEnable	Displays the Supplicant function in the Port Access Entity (PAE) as True or False.
AllowNonEapHost	Enables network access to hosts that do not participate in 802.1X authentication. The default is disabled.
Status	Configures the authentication status for this port. The default is forceAuthorized. <ul style="list-style-type: none"> • auto: enables the EAP authentication process by sending the EAP request messages to the RADIUS server. • forceAuthorized: disables the EAP authentication and puts the port into force-full authorized mode.
MultiHostMaxClients	Specifies the value representing the maximum number of supplicants allowed to get authenticated on the port.
GuestVlanId	Specifies the VLAN to be used as a Guest VLAN. Access to unauthenticated hosts connected to this port is provided through this VLAN. 0 indicates that Guest VLAN is not enabled for this port.
FailOpenVlanId	Specifies the Fail Open VLAN ID for this port. If the switch declares the RADIUS servers unreachable, then all new devices are allowed access into the configured Fail Open VLAN. 0 indicates that Fail Open VLAN is not enabled for this port.
NonEapMaxClients	Specifies the maximum number of NEAP authentication MAC addresses allowed on this port. Zero indicates that NEAP authentication is disabled for this port.
EAPMaxClients	Specifies the maximum number of EAP authentication MAC addresses allowed on this port. Zero indicates that EAP authentication is disabled for this port.
MultiHostSingleAuthEnabled	Indicates that the unauthenticated devices can access the network only after an EAP or NEAP client is successfully authenticated on the port. The VLAN to which the devices are allowed access is the authenticated client's VLAN. The default is false.
PortGuestIsid	Specifies the I-SID to be used as a Guest I-SID. Access to unauthenticated hosts connected to this port is provided through this I-SID. 0 indicates that Guest I-SID is not enabled for this port.
FailOpenIsid	Specifies the Fail Open I-SID for this port. If the switch declares the RADIUS servers unreachable, then all new devices are allowed access into the configured Fail Open I-SID. 0 indicates that Fail Open I-SID is not enabled for this port.
FlexUniStatus	Displays the current Flex-UNI status for this port.

Name	Description
AdminTrafficControl	Configures the Administrative Traffic Control. The default is inOut. <ul style="list-style-type: none"> inOut: enables the Admin Traffic Control for input and output traffic. in: enables the Admin Traffic Control for input traffic only.
OperTrafficControl	Displays the current Operational Traffic Control status.
LldpAuthEnabled	Enables LLDP authentication for this port. The default is disabled.
PortOrigin	Specifies the source of EAP configuration on the port: <ul style="list-style-type: none"> config - through CLI or EDM autoSense - through Zero Touch Fabric Configuration
DynamicMHSAEnabled	Displays the Dynamic MHSa configuration status.
ReauthOrigin	Specifies the origin of EAPOL reauthentication configuration on the port, either manually configured through CLI or dynamically configured through RADIUS.
ReauthPeriodOrigin	Specifies the origin of EAPOL reauthentication period configuration on the port, either manually configured through CLI or dynamically configured through RADIUS.
TrafficControlOrigin	Specifies the origin of Traffic Control configuration on the port. The supported values are: <ul style="list-style-type: none"> config - Traffic Control is enabled by the user. radius - Traffic Control is enabled by Extensible Authentication Protocol (EAP) through Remote Authentication Dial-In User Service (RADIUS) response.
Authenticator configuration	Displays the current Authenticator Port Access Entity (PAE) state. The states are: <ul style="list-style-type: none"> authenticate authenticated Failed
ReAuthEnabled	Reauthenticates an existing supplicant at the time interval specified in ReAuthPeriod. The default is disabled.
QuietPeriod	Configures the time interval (in seconds) between authentication failure and the start of a new authentication.

Name	Description
ReAuthPeriod	Specifies the time interval, in seconds, between successive authentication. Configure the value to 0 to prevent EAP or NEAP sessions from aging out. Caution: Preventing re-authentication can introduce a security risk. The default is 3600 (1 hour).
RetryMax	Specifies the maximum Extensible Authentication Protocol (EAP) requests sent to the supplicant before timing out the session. The default is 2.
RetryCount	Specifies the maximum number of retries attempted.

Configure EAP on an Extreme Integrated Application Hosting Port



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to configure EAP or change the authentication status on Extreme Integrated Application Hosting (IAH) ports. IAH ports are force-authorized by default and are not authenticated by the RADIUS server. You can change this setting so that the IAH ports stay unauthorized.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **EAPOL** tab.
4. (Optional) Select **AllowNonEapHost**.
5. In **Status**, select the required option.
6. In **MultiHostMaxClients**, type a value.
7. In **GuestVlanId**, type a VLAN ID.
8. In **FailOpenVlanId**, type a VLAN ID.
9. In **NonEapMaxClients**, type a value.
10. In **EapMaxClients**, type a value.
11. Select **MultiHostSingleAuthEnabled**.
12. In **PortGuestIsid**, type the I-SID to be used as a Guest I-SID.
13. In **FailOpenIsid**, type the Fail Open I-SID.
14. Select the **AdminTrafficControl** option as **inOut** or **in**.
15. Select **LldpAuthEnabled** to enable LLDP authentication for network access.
16. Select **ReAuthEnabled**.
17. In **QuietPeriod**, type a time interval.

18. In **ReAuthPeriod**, type a time interval.

19. In **RetryMax**, type a value.

20. Select **Apply**.

EAPOL Field Descriptions

Use data in the following table to use the **EAPOL** tab.

Name	Description
PortCapabilities	Shows the capabilities of the Port Access Entity (PAE) associated with the Extreme Integrated Application Hosting (IAH) port. This parameter indicates whether Authenticator functionality, supplicant functionality, both, or neither, is supported by the PAE of the IAH port. The following capabilities are supported by the PAE of the IAH port: <ul style="list-style-type: none"> • authImplemented: A Port Access Controller Protocol (PACP) Extensible Authentication Protocol (EAP) authenticator functions are implemented. • virtualPortsImplemented: Virtual Port functions are implemented.
PortVirtualPortsEnable	Shows the status of the Virtual Ports function for the IAH port.
PortCurrentVirtualPorts	Shows the current number of virtual ports running on the IAH port.
PortAuthenticatorEnable	Shows the status of the Authenticator function in the PAE.
PortSupplicantEnable	Shows the Supplicant function in the PAE.
AllowNonEapHost	Enables network access to hosts that do not participate in 802.1X authentication. The default is disabled.
Status	Specifies the authentication status for the IAH port. <ul style="list-style-type: none"> • auto - enables EAP authentication process by sending the EAP request messages to the RADIUS server. • forceAuthorized - disables EAP authentication and puts the IAH port into force-full authorized mode. The default is forceAuthorized.
MultiHostMaxClients	Specifies the maximum number of supplicants authenticated on the IAH port.

Name	Description
GuestVlanId	Specifies the VLAN ID to be used as a Guest. Access to unauthenticated hosts connected to the IAH port is provided through this VLAN. 0 indicates that Guest VLAN is not enabled.
FailOpenVlanId	Specifies the Fail Open VLAN ID for the specific IAH port. If RADIUS server is not reachable on the switch, then all new devices are allowed access to the configured Fail Open VLAN ID. 0 indicates that Fail Open VLAN ID is not enabled.
NonEapMaxClients	Specifies the maximum number of NEAP authentication MAC addresses allowed on the specific IAH port. 0 indicates that NEAP authentication is disabled.
EAPMaxClients	Specifies the maximum number of EAP authentication MAC addresses allowed on the specific IAH port. 0 indicates that EAP authentication is disabled.
MultiHostSingleAuthEnabled	Enables the functionality for network access to the unauthenticated devices only after an EAP or NEAP client is successfully authenticated on the IAH port. The VLAN ID to which the devices are allowed access is the authenticated client's VLAN. The default is disabled.
PortGuestIsid	Specifies the I-SID to be used as a Guest I-SID. Access to unauthenticated hosts connected to the IAH port is provided through this I-SID. 0 indicates that Guest I-SID is not enabled for this port.
FailOpenIsid	Specifies the Fail Open I-SID for the IAH port. If the switch declares the RADIUS servers unreachable, then all new devices are allowed access into the configured Fail Open I-SID. 0 indicates that Fail Open I-SID is not enabled for this port.
FlexUniStatus	Displays the current Flex-UNI status for this IAH port.
AdminTrafficControl	Configures the Administrative Traffic Control. The default is inOut. <ul style="list-style-type: none"> • inOut: enables the Admin Traffic Control for input and output traffic. • in: enables the Admin Traffic Control for input traffic only.
OperTrafficControl	Displays the current Operational Traffic Control status.
LldpAuthEnabled	Enables LLDP authentication for this IAH port. The default is disabled.

Name	Description
PortOrigin	Specifies the source of EAP configuration on the IAH port: <ul style="list-style-type: none"> • config - through CLI or EDM • autoSense - through Zero Touch Fabric Configuration
DynamicMHSAEnabled	Displays the Dynamic MHSA configuration status.
TrafficControlOrigin	Indicates the origin of Traffic Control configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Traffic Control is enabled by the user. • radius - Traffic Control is enabled by Extensible Authentication Protocol (EAP) through Remote Authentication Dial-In User Service (RADIUS) response.
Authenticate	Shows the current Authenticator Port Access Entity (PAE) authenticate status.
Authenticated	Shows the current Authenticator Port Access Entity (PAE) authenticated status.
Failed	Shows the current Authenticator Port Access Entity (PAE) failure status.
ReAuthEnabled	Enables reauthentication of an existing supplicant based on the specified reauthentication time interval. The default is disabled.
QuietPeriod	Specifies the time interval (in seconds) between authentication failure and start of authentication.
ReauthPeriod	Specifies the time interval (in seconds) between successive reauthentications. The default is 3600 (1 hour).
RetryMax	Specifies the maximum Extensible Authentication Protocol (EAP) requests sent to the supplicant before timing out the session. The default is 2.
RetryCount	Specifies the maximum number of retries attempted.

Show the Port Access Entity Port Table

About This Task

Use the Port Access Entity (PAE) Port Table to display system-level information for each port the PAE supports. An entry display in this table for each port of this system.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **802.1X - EAPOL**.
3. Select the **EAP Security** tab.

EAP Security Field Descriptions

Use the data in the following table to use the **EAP Security** tab.

Name	Description
PortNumber	Indicates the port number associated with this port.
PortCapabilities	Indicates the capabilities of this PAE port. <ul style="list-style-type: none"> • authImplemented—PACP EAP authenticator functions are implemented in this PAE. • virtualPortsImplemented—Virtual Port functions are implemented in this PAE.
PortVirtualPortsEnable	Displays the status of the Virtual Ports function for the real port as True or False.
PortCurrentVirtualPorts	Displays the current number of virtual ports running in the port
PortAuthenticatorEnable	Displays the status of the Authenticator function in the Port Access Entity (PAE) as True or False.
PortSupplicantEnable	Displays the Supplicant function in the Port Access Entity (PAE) as True or False.
AllowNonEapHost	Displays the status if the system is enabled to allow hosts that do not participate in 802.1X authentication to get network access.
Status	Displays the authentication status for this port. The default is forceAuthorized.
MultiHostMaxClients	Indicates the value representing the maximum number of supplicants allowed to get authenticated on the port.
GuestVlanId	Specifies the VLAN to be used as a Guest VLAN. Access to unauthenticated hosts connected to this port is provided through this VLAN. 0 indicates that Guest VLAN is not enabled for this port.
FailOpenVlanId	Specifies the Fail Open VLAN ID for the port. If the switch declares the RADIUS servers unreachable, then all new devices are allowed access into the configured Fail Open VLAN. 0 indicates that Fail Open VLAN is not enabled for this port.
NonEapMaxClients	Indicates the maximum number of non-EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port.

Name	Description
EapMaxClients	Indicates the maximum number of EAPoL authentication MAC addresses allowed on this port. Zero indicates that EAPoL authentication is disabled for this port.
MultiHostSingleAuthEnabled	Indicates that the unauthenticated devices can access the network only after an EAP or NEAP client is successfully authenticated on the port. The VLAN to which the devices are allowed access is the authenticated client's VLAN. The default is false.
ProcessRadiusCOAPackets	Specifies whether to process any RADIUS requests-server packets that are received on this port.
PortGuestIsid	Specifies the I-SID to be used as a Guest I-SID. Access to unauthenticated hosts connected to this port is provided through this I-SID. 0 indicates that Guest I-SID is not enabled for this port.
FailOpenIsid	Specifies the Fail Open I-SID for the port. If the switch declares the RADIUS servers unreachable, then all new devices are allowed access into the configured Fail Open I-SID. 0 indicates that Fail Open I-SID is not enabled for this port.
FlexUniStatus	Displays the Flex-UNI status for the port.
AdminTrafficControl	Specifies the Administrative Traffic Control for the port. The default is inOut.
OperTrafficControl	Displays the Operating Traffic Control for the port.
LldpAuthEnabled	Specifies if LLDP Authentication is enabled. The default is 0 (disabled).
PortOrigin	Displays the Port Origin configuration status for the port.
DynamicMHSAEnabled	Displays the Dynamic MHSA status for the port.
TrafficControlOrigin	Specifies the origin of Traffic Control configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Traffic Control is enabled by the user. • radius - Traffic Control is enabled by Extensible Authentication Protocol (EAP) through Remote Authentication Dial-In User Service (RADIUS) response.
ReauthOrigin	Specifies the origin of EAPoL reauthentication configuration on the port, either manually configured through CLI or dynamically configured through RADIUS.
ReauthPeriodOrigin	Specifies the origin of EAPoL reauthentication period configuration on the port, either manually configured through CLI or dynamically configured through RADIUS.

Show EAP Authentication

About This Task

Use the Authenticator Configuration table to display configuration objects for the Authenticator PAE associated with each port.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **802.1X - EAPOL**.
3. Select the **Authentication** tab.

Authentication Field Descriptions

Use the data in the following table to use the **Authentication** tab.

Name	Description
PortNumber	Indicates the number associated with this port.
ReAuthEnabled	Indicates the re-authentication status of an existing supplicant at the time interval specified in ReAuthPeriod. The default is false.
QuietPeriod	Indicates the time interval (in seconds) between authentication failure and the start of a new authentication. The default is 60.
ReAuthPeriod	Indicates the time interval in seconds between successive re-authentications. The default is 3600 (1 hour).
RetryMax	Indicates the maximum Extensible Authentication Protocol (EAP) requests sent to the supplicant before timing out the session. The default is 2.
RetryCount	Indicates the count of the number of authentication attempts.

View Multihost Status Information

Use the following procedure to display multiple host status for a port.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **802.1X – EAPOL**.
3. Select the **MultiHost Status** tab.
4. (Optional) To clear an EAP client session, double-click a value in the **MacClear** column.
 - a. Select **true**.

- b. Select **Apply**.

MultiHost Status *Field Descriptions*

The following table describes values on the **MultiHost Status** tab.

Name	Description
PortNumber	Indicates the port number associated with this client.
ClientMACAddr	Indicates the MAC address of the client.
PaeState	Indicates the current state of the authenticator PAE state machine.
VlanId	Indicates the VLAN assigned to the client.
Priority	Specifies the priority associated with this client MAC. This priority could be the Radius assigned priority or the port QOS level.
SwUniBindings	Indicates the Extensible Authentication Protocol (EAP) VLAN:ISID bindings that the switch represents as an hexadecimal value.
IsidSource	Indicates the origin of I-SID value: <ul style="list-style-type: none"> radius - received from the RADIUS server. autoconfig - calculated using the auto-isid-offset command, that the user configures on the switch. config - configured statically. notAvaliable - does not use EAP with FlexUNI, hence there is no I-SID to use.
AcId	Indicates the dynamic Access Control List (ACL) on the specific port.
AcIdList	Indicates the list of dynamic Access Control Entries (ACE) on the specific port.
DynamicSettings	Displays the Dynamic settings received from the Remote Authentication Dial-In User Service (RADIUS) server.
MhsaDynamic	Specifies whether or not the switch receives the MHSA VSA settings from RADIUS.
MacClear	Clears the EAP MAC entry associated with a specific port.

View EAP Session Statistics

Use the following procedure to display multiple host session information for a port.

Procedure

1. In the navigation pane, expand **Configuration --> Security --> Data Path**.
2. Select **802.1X – EAPOL**.
3. Select the **MultiHost Session** tab.

MultiHost Session Field Descriptions

The following table describes values on the **MultiHost Session** tab.

Name	Description
StatsPortNumber	Indicates the port number associated with this port.
StatsClientMACAddr	Indicates the MAC address of the client.
Id	Indicates the unique identifier for the session.
AuthenticMethod	Indicates the authentication method used to establish the session.
Time	Indicates the elapsed time of the session.
TerminateCause	Indicates the cause of the session termination.
UserName	Indicates the user name that represents the identity of the supplicant PAE.

View EAPoL Authenticator Statistics

Use the following procedure to display the Authenticator Port Access Entity (PAE) statistics for each selected port.

Procedure

1. On the Device Physical View, select the port you want to graph.
The system displays a yellow outline around the selected ports.

To select multiple ports, press and hold down the Ctrl key while you select the ports you want to configure. The system displays a yellow outline around the selected ports.

2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select **EAPOL Stats**.
5. If you select multiple ports from the **Graph port EAPoL Stats** tab Show list, select: Absolute Value, Cumulative, Average/sec, Minimum/sec, Maximum/sec, or LastVal/sec.

EAPOL Stats *Field Descriptions*

The following table describes values on the **EAPOL Stats** tab.

Name	Description
InvalidFramesRx	Displays the number of EAPoL frames received by this Authenticator in which the frame type is not recognized.
EapLengthErrorFramesRx	Displays the number of EAPoL frames received by this Authenticator in which the Packet Body Length field is invalid.
StartFramesRx	Displays the number of EAPoL start frames received by this Authenticator.
EapFramesRx	Displays the number of EAPoL-EAP frames received by this Authenticator.
LogoffFramesRx	Displays the number of EAPoL Logoff frames received by this Authenticator.
LastRxFrameVersion	Displays the last received version of the EAPoL frame by this Authenticator.
LastRxFrameSource	Displays the source MAC address of the last received EAPoL frame by this Authenticator.
AuthEapFramesTx	Displays the number of EAPoL-EAP frames transmitted by the Authenticator.

View NEAP MAC Information

Use this procedure to view NEAP client MAC information on a port.

Procedure

1. In the navigation pane, expand **Configuration --> Security --> Data Path**.
2. Select **802.1X – EAPOL**.
3. Select **NEAP Radius** tab.

NEAP Radius *Field Descriptions*

The following table describes values on the **NEAP Radius** tab.

Name	Description
MacPort	Indicates the port number associated with this port.
MacAddr	Indicates the MAC address of the client.
MacStatus	Indicates the authentication status of the NEAP host that is authenticated using the RADIUS server.
VlanId	Indicates the VLAN assigned to the client.
MacClear	Clears the non EAP MAC entry associated with a specific index.

Name	Description
MacPriority	Indicates the priority associated with this Non-EAP client MAC. This priority could be the Radius assigned priority or the port QOS level.
SwUniBindings	Indicates the VLAN and I-SID bindings. VLAN is represented with 2 bytes and I-SID is represented with 4 bytes. The output is a continuous hexadecimal representation of the VLAN that is followed by the corresponding I-SID.
IsidSource	Indicates the source of the I-SID value. An I-SID value is generated in one of the following ways: <ul style="list-style-type: none"> • radius—Indicates that that I-SID value is learned from the RADIUS server. • autoconfig—Indicates that I-SID value is calculated by using the auto-isid-offset that you configured. • config—Indicates that the I-SID value is statically configured. • notAvailable—Indicates that no I-SID value is available because EAP with FlexUNI is not used.
NonEapAuthType	Indicates the authentication type of the Non-EAP client: <ul style="list-style-type: none"> • radius - received from RADIUS server. • lldp - received from Link Layer Discovery Protocol (LLDP).
AcId	Indicates the dynamic Access Control List (ACL) on the specific port.
AcIdList	Indicates the list of dynamic Access Control Entries (ACE) on the specific port.
DynamicSettings	Displays the dynamic settings received from the Remote Authentication Dial-In User Service (RADIUS) server and successfully applied on the port or VLAN.
MhsaDynamic	Specifies whether or not the switch receives the MHSa VSA settings from RADIUS.

View EAPoL Vendor Specific Attributes Information

About This Task

Use this procedure to view Extensible Authentication Protocol over LAN (EAPoL) Vendor Specific Attributes (VSA) information for a port.


Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**
2. Select **802.1x - EAPOL**.
3. Select the **EAP/NEAP VSA** tab.

EAP/NEAP VSA Field Descriptions

Use the data in the following table to use the **EAP/NEAP VSA** tab.

Name	Description
PortNumber	Specifies the port number associated with the RADIUS client.
Client MACAddr	Specifies the MAC address of the RADIUS client.
ID	Specifies the VSA ID value as received from the RADIUS server.
Index	Specifies the VSA index received from the RADIUS server.
Content	Specifies the VSA-specific string format.



ExtremeCloud IQ Support

[ExtremeCloud IQ Agent Configuration Considerations](#) on page 1085

[Zero Touch Deployment](#) on page 1086

[Unified Metrics and Events Reporting](#) on page 1087

[ExtremeCloud IQ Support Configuration using CLI](#) on page 1087

[ExtremeCloud IQ Support Configuration using EDM](#) on page 1093

Table 103: IQ Agent product support

Feature	Product	Release introduced
ExtremeCloud IQ Agent	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2 Supported on: <ul style="list-style-type: none">• VSP4900-48P• VIM5-4YE• VIM5-4X• VIM5-4XE VOSS 8.2.5 Supported on: <ul style="list-style-type: none">• VSP4900-24S• VSP4900-24XE• VSP4900-12MXU-12XE• VIM5-2Q
	VSP 7400 Series	VOSS 8.2

Table 103: IQ Agent product support (continued)

Feature	Product	Release introduced
Unified Metrics and Events Reporting	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	Not Supported
	VSP 7400 Series	Not Supported

For the most current information on switches supported by ExtremeCloud™ IQ, see [ExtremeCloud™ IQ Learning What's New](#).

ExtremeCloud IQ provides cloud-managed networking, and delivers unified, full-stack management of wireless access points, switches, and routers. It enables onboarding, configuration, monitoring, troubleshooting, reporting, and more. Using innovative machine learning and artificial intelligence technologies, ExtremeCloud IQ analyzes and interprets millions of network and user data points, from the network edge to the data center, to power actionable business and IT insights, and to deliver new levels of network automation and intelligence.

The switch supports a zero touch connection to ExtremeCloud IQ. Zero touch deployment is used to deploy and configure a switch using ExtremeCloud IQ.

The switch integrates with ExtremeCloud IQ using ExtremeCloud IQ Agent. When you enable IQAgent, you can configure and monitor your devices using ExtremeCloud IQ.

For more information about ExtremeCloud IQ, see <https://www.extremenetworks.com/support/documentation/extremecloud-iq/>.

ExtremeCloud IQ Agent Configuration Considerations

The following configuration considerations apply to ExtremeCloud IQ Agent:

- SSH and SSH password authentication is required.

boot config flag ssh is enabled when ExtremeCloud IQ Agent is enabled. **boot config flag ssh** cannot be disabled while ExtremeCloud IQ Agent is enabled.

- SNMP is required.

boot config flag block-snmp is disabled when ExtremeCloud IQ Agent is enabled. **boot config flag block-snmp** cannot be enabled while ExtremeCloud IQ Agent is enabled.

- High Secure mode disables ExtremeCloud IQ Agent automatically. ExtremeCloud IQ Agent must be enabled manually when this mode is enabled.
- ExtremeCloud IQ Agent is not supported in Enhanced Secure mode.
- An IP address that corresponds to the ExtremeCloud IQ pool and can display it in the NTP list. The IP does not try to synchronize if NTP is globally disabled on the switch. If NTP is enabled, you can see synchronization failure messages if the IP for the pool is blocked or is unreachable. As a best practice, if you have issues connecting to the cloud, check the clock on the switch and if it is incorrect, resolve this by either configuring an NTP server or manually configuring the correct time.

**Note**

You must configure a Segmented Management Instance to use ExtremeCloud IQ Agent. For more information, see [Segmented Management](#) on page 106.

For information about onboarding switches, see <https://www.extremenetworks.com/support/documentation/extremecloud-iq/>.

Zero Touch Deployment

Zero Touch Deployment enables a switch to be deployed automatically with ExtremeCloud IQ but you still must onboard the switch on the ExtremeCloud IQ side. When the switch powers on, the DHCP Client obtains the IP address and gateway from the DHCP Server, and discovers the Domain Name Server, connecting the switch automatically to ExtremeCloud IQ Site Engine or to ExtremeCloud IQ cloud management applications.

With Zero Touch Deployment, ExtremeCloud IQ Agent is enabled by default.

To use zero touch functionality, your switch must be in a Zero Touch Deployment-ready configuration mode, which means the switch cannot have existing primary or secondary configuration files loaded. Factory shipped switches are Zero Touch Deployment ready because they deploy without configuration files. However, existing switches require manual preparation before Zero Touch Deployment can function.

For more information about preparing your switch for Zero Touch Deployment, see [Zero Touch Deployment](#) on page 91.

DHCP Option 43 Support

With the support of DHCP option 43, DHCP can dynamically configure the IP address of a private/non-public ExtremeCloud IQ server for zero touch deployments when the default ExtremeCloud IQ server (nac.extremecloudiq.com) is not desired.

To use this functionality, DHCP Client must be enabled. For information about DHCP Client for a Segmented Management Instance, see [DHCP Client for Segmented Management Instance](#) on page 119.

Considerations

The following considerations apply with DHCP option 43:

- A dynamic IP address overwrites the default value (hac.extremecloudiq.com) or 0.0.0.0.
- A static server IP address overwrites a dynamic server IP address.
- A dynamic server IP address does not overwrite an existing static server IP address.

If a static server IP address is already configured and a new value is received from the DHCP server, the following warning displays on the console: `WARNING Dynamic Cloud IQ Server Address x.x.x.x provided by DHCP option 43 could not be set. Static configured server address y.y.y.y cannot be overwritten by a dynamic address.`

- The default value (hac.extremecloudiq.com) replaces the dynamic server IP address if the DHCP Client is disabled on the switch.
- The dynamic server IP address is not saved in the running configuration.

Unified Metrics and Events Reporting

The Unified Metrics and Events Reporting feature collects data from multiple standard input devices and streams it dynamically and directly to ExtremeCloud IQ, instead of using the MIBs and the SNMP traps through the ExtremeCloud IQ Agent.

This feature collects and reports statistics and events on the switches, such as:

- CPU utilization
- DHCP Snooping VLAN attributes
- Fabric Adjacencies
- Fan rotations per minute, pulse width modulation, failure, and recovery
- Port operational state changes
- Power over Ethernet measurements and counters
- Power Supply Units failure and recovery

ExtremeCloud IQ Support Configuration using CLI

After your device is onboarded (that is, the serial number for the device is associated with your ExtremeCloud IQ account), you are only required to enable ExtremeCloud IQ Agent. Other feature configuration, such as configuring proxy parameters and configuring access to ExtremeCloud IQ is optional.



Note

You must configure a Segmented Management Instance to use ExtremeCloud IQ Agent. For more information, see [Segmented Management Instance Configuration using the CLI](#) on page 123.

For information about onboarding switches, see <https://www.extremenetworks.com/support/documentation/extremecloud-iq/>.

Configure ExtremeCloud IQ Agent

You must first onboard the device. When zero touch connection establishes, IQ Agent is enabled, by default. Before IQ Agent is operational, you must first disable IQ Agent, configure the ExtremeCloud IQ IPv4 address or DNS name, and then reenable IQ Agent.

Before You Begin

You must first onboard the device.

Procedure

1. Enter Application Configuration mode:

```
enable

configure terminal

application
```
2. Disable IQ Agent:

```
no iqagent enable
```
3. Configure the ExtremeCloud IQ IPv4 address or DNS name:

```
iqagent server address WORD<1-255>
```
4. Enable IQ Agent:

```
iqagent enable
```

Example

Configure IQ Agent:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#no iqagent enable
Switch:1(config-app)#iqagent server address hac.extremecloudiq.com
Switch:1(config-app)#iqagent enable
```

Display default IQ Agent configuration:

```
Switch:1>show application iqagent

=====
IQAgent Info
=====
Agent Admin State      : true
Agent Version          : 0.4.3
Agent Oper State       : disconnected
Server Address         : hac.extremecloudiq.com
Server Address Origin  : None
Proxy Address          : 0.0.0.0
Proxy TCP Port         : 0
Proxy Username         :
```

Configure Access to ExtremeCloud IQ

Use this task to configure IQ Agent parameters to access ExtremeCloud IQ.

Before You Begin

You must onboard the device and configure any optional IQ Agent parameters on the supported device before you enable IQ Agent.

You can configure the IQ Agent parameters on the supported devices first, and then onboard the devices (that is, add the serial numbers for the devices in the ExtremeCloud IQ GUI) or vice versa.

For information about onboarding switches, see <https://www.extremenetworks.com/support>.

Procedure

1. Enter Application Configuration mode:
`enable`
`configure terminal`
`application`
2. Configure the ExtremeCloud IQ IPv4 address or DNS name:
`iqagent server address WORD<1-255>`

Examples

Configure access to ExtremeCloud IQ using an IPv4 address:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#iqagent server address 192.0.2.1
```

Configure access to ExtremeCloud IQ using a DNS name:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#iqagent server address extremecloudiq.com
```

Variable Definitions

The following table defines parameters for the **iqagent server** command.

Variable	Value
<code>address <WORD 1-255></code>	Specifies the ExtremeCloud IQ IPv4 address or DNS name.

Configure Proxy Parameters

If you use a proxy https server in your network, you must configure proxy parameters so that the IQ Agent on the device can communicate with ExtremeCloud IQ through the proxy.

Use this task to configure the proxy parameters for ExtremeCloud IQ on the IQ Agent.



Note

You must onboard the device and configure any optional IQ Agent parameters on the supported device before you enable IQAgent.

You can configure the IQ Agent parameters on the supported devices first, and then onboard the devices (that is, add the serial numbers for the devices in the ExtremeCloud IQ GUI) or vice versa.

For information about onboarding switches, see <https://www.extremenetworks.com/support>.

Procedure

1. Enter Application Configuration mode:

```
enable
```

```
configure terminal
```

```
application
```

2. Configure the proxy IPv4 address or DNS name:

```
iqagent proxy address <WORD 1-255> tcp-port <1-49151>
```

3. Configure the proxy user name and password for the ExtremeCloud IQ account:

```
iqagent proxy username <WORD 1-64> password <WORD 1-128>
```

Examples

Configure proxy parameters using an IPv4 address:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#iqagent proxy address 192.0.2.254 tcp-port 21
Switch:1(config-app)#iqagent proxy username admin password ****
```

Configure proxy parameters using a DNS name:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#iqagent proxy address hac.extremecouldiq.com tcp-port 21
Switch:1(config-app)#iqagent proxy username admin password ****
```


Variable Definitions

The following table defines parameters for the **iqagent proxy** command.

Variable	Value
<i>address</i> <WORD 1-255>	Specifies the proxy IPv4 address or DNS name.
<i>tcp-port</i> <1-49151>	Specifies the TCP port.
<i>username</i> <WORD 1-64>	Specifies the proxy server user name.
<i>password</i> <WORD 1-128>	Specifies the proxy server password.

Display ExtremeCloud IQ Agent Information

About This Task

Use this task to display ExtremeCloud IQ Agent configuration information and status.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IQ Agent configuration information and status:

```
show application iqagent
```

Example

Display IQ Agent configuration information and status when IQ Agent is enabled using the default ExtremeCloud IQ server:

```
Switch:1>show application iqagent

=====
                        IQAgent Info
=====
Agent Admin State      : true
Agent Version          : 0.2.7
Agent Oper State       : connected
Server Address         : hac.extremecloudiq.com
Server Address Origin  : None
Proxy Address          : extremeiq.com
Proxy TCP Port         : 21
Proxy Username         : admin
```

Display IQ Agent disabled state:

```
Switch:1>show application iqagent

=====
                        IQAgent Info
=====
Agent Admin State      : false
Agent Version          : 0.2.7
Agent Oper State       : disconnected
Server Address         : 0.0.0.0
Server Address Origin  : None
Proxy Address          : 0.0.0.0
Proxy TCP Port         : 0
Proxy Username         :
```

Display IQ Agent configuration information and status when DHCP provides a dynamic server IP address:

```
Switch:1>show application iqagent
```

```
=====
                        IQAgent Info
=====
```

```
Agent Admin State      : true
Agent Version          : 0.2.7
Agent Oper State       : disconnected
Server Address         : 192.0.2.1
Server Address Origin  : DHCP
Proxy Address          : 0.0.0.0
Proxy TCP Port         : 0
Proxy Username         :
```

Display IQ Agent configuration information and status when DHCP Client is disabled on the switch:

```
Switch:1>show application iqagent
```

```
=====
                        IQAgent Info
=====
```

```
Agent Admin State      : false
Agent Version          : 0.2.7
Agent Oper State       : disconnected
Server Address         : hac.extremecloudiq.com
Server Address Origin  : None
Proxy Address          : 0.0.0.0
Proxy TCP Port         : 0
Proxy Username         :
```

Display ExtremeCloud IQ Agent Status

About This Task

Use this task to display IQ Agent status information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IQ Agent status information:

```
show application iqagent status
```

Examples

The following output shows an example of an IQ Agent that is enabled with a successful connection to ExtremeCloud IQ.

```
Switch:1>show application iqagent status
```

```
=====
                        IQAgent Status
=====
```

```
Connection Status      : Connected
Last Onboard Time      : 18:54:23 11 27 2019 UTC
Agent Version          : 0.2.7
Association URL         : https://10.16.231.98/hac-webapp/rest/v1/association
Poll URL               : https://10.16.231.98/hac-webapp/rest/v1/poll/1904Q-20028
Monitor Frequency      : 600
```

```

Poll Frequency           : 30
Last Poll Status         : SUCCESS
Last Poll Success Time   : 14:39:16 11 28 2019 UTC
Last Health Status       : SUCCESS
Last Health Success Time : 14:38:35 11 28 2019 UTC
Last Monitor Status      : SUCCESS
Last Monitor Success Time: 14:38:35 11 28 2019 UTC

```

The following output shows an example of an IQ Agent that is enabled but not connected to ExtremeCloud IQ.

```

Switch:1>show application iqagent status
=====
                        IQAgent Status
=====
Connection Status       : Disconnected
Time Disconnected       : 19:22:33 08 25 2022 UTC
Association Frequency    : Slow
Previous Association Attempt : 19:04:06 08 29 2022 UTC
Next Association Attempt  : 20:04:43 08 29 2022 UTC

```

Reinstall ExtremeCloud IQ Agent Firmware

Perform this procedure to return the ExtremeCloud IQ Agent firmware version on the switch to the version bundled with the OS image currently installed on the switch, for example, if you downgrade the OS image version and do not reconnect to ExtremeCloud IQ automatically.

Procedure

1. Enter Application Configuration mode:

```
enable
```

```
configure terminal
```

```
application
```
2. Disable ExtremeCloud IQ Agent:

```
no iqagent enable
```
3. Reinstall the ExtremeCloud IQ Agent firmware:

```
software iqagent reinstall
```
4. Enable ExtremeCloud IQ Agent:

```
iqagent enable
```

ExtremeCloud IQ Support Configuration using EDM

Perform the procedures in this section to configure ExtremeCloud IQ Agent on the switch using the Enterprise Device Manager (EDM).

Configure ExtremeCloud IQ Agent

Before You Begin

You must first onboard the device and configure any optional IQ Agent parameters before you enable IQ Agent.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **IQAgent**.
3. Select the **Globals** tab.
4. Configure optional parameters as required.
5. Select **Apply**.
6. Select **Enable**.
7. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Global** tab to configure the ExtremeCloud IQ Agent.

Name	Description
Enable	Specifies whether IQ Agent is enabled. The default is enabled.
Version	Displays the ExtremeCloud IQ Agent firmware version running on the switch.
OperStatus	Displays the operational status of ExtremeCloud IQ Agent on the switch.
ServerAddressType	Specifies the address type of the ExtremeCloud IQ Agent server address.
ServerAddress	Specifies the ExtremeCloud IQ IPv4 address or DNS name. The default is hac.extremcloudiq.com.
ServerAddressOrigin	Specifies the origin for the ExtremeCloud IQ server address: <ul style="list-style-type: none"> • none—not configured • configured—manual configuration • dhcp—obtained through DHCP
Address	Specifies the proxy IPv4 address or DNS name.
TcpPort	Specifies the TCP connection port.
UserName	Specifies the proxy server user name.
Password	Specifies the proxy server password.
AssociationUrl	Displays the association URL of ExtremeCloud IQ.
PollUrl	Displays the poll URL of ExtremeCloud IQ.
MonitorFreq	Displays the monitoring frequency, in seconds, of ExtremeCloud IQ.
PollFreq	Displays the polling frequency, in seconds, of ExtremeCloud IQ.

Name	Description
LastOnboardTime	Displays the last onboard time of ExtremeCloud IQ.
LastPollStatus	Displays the last poll status of ExtremeCloud IQ.
LastPollTime	Displays the last poll time for a successful attempt.
LastMonitorStatus	Displays the last monitoring status of ExtremeCloud IQ.
LastMonitorTime	Displays the last monitor time for a successful attempt.
LastHealthStatus	Displays the last health status of ExtremeCloud IQ.
LastHealthTime	Displays the last health time for a successful attempt.
LastDisconnectedTime	Shows when IQ Agent disconnected from ExtremeCloud IQ. This status is not valid unless IQ Agent is enabled and in a disconnected state.
LastAttemptedAssociationTime	Shows the last time IQ Agent attempted to associate with ExtremeCloud IQ. This status is not valid unless IQ Agent is enabled and in a disconnected state.
NextAssociationAttempt	Shows the next time IQ Agent will attempt to associate with ExtremeCloud IQ. This field displays <i>In progress</i> if IQ Agent is currently attempting to associate. This status is not valid unless IQ Agent is enabled and in a disconnected state.
AssociationFrequencyMode	Shows the frequency at which IQ Agent attempts to associate with ExtremeCloud IQ. Fast mode indicates an attempt every 40 seconds. Slow mode indicates an attempt every hour. For the first 24 hours of disconnection, the mode is Fast mode. The mode then changes to slow and remains in slow mode until successfully connected. This status is not valid unless IQ Agent is enabled and in a disconnected state. When you first onboard a device to ExtremeCloud IQ, it does not connect to the device immediately; it waits until the device attempts association again, which can take up to 1 hour.



Extreme Integrated Application Hosting

[Extreme Integrated Application Hosting](#) on page 1097

[Operational Considerations and Restrictions](#) on page 1101

[Third Party Virtual Machine](#) on page 1101

[Fabric IPsec Gateway Fundamentals](#) on page 1103

[Virtual Services Configuration using CLI](#) on page 1112

[Virtual Services Configuration using EDM](#) on page 1126

[Fabric IPsec Gateway Configuration using CLI](#) on page 1135

Table 104: Extreme Integrated Application Hosting product support

Feature	Product	Release introduced
Extreme Integrated Application Hosting	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1.5 Supported on VSP4900-12MXU-12XE and VSP4900-24XE
	VSP 7400 Series	VOSS 8.0

Table 104: Extreme Integrated Application Hosting product support (continued)

Feature	Product	Release introduced
Fabric IPsec Gateway virtual machine	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3 Supported on VSP4900-12MXU-12XE and VSP4900-24XE
	VSP 7400 Series	VOSS 8.2

Extreme Integrated Application Hosting

Extreme Integrated Application Hosting (IAH) architecture provides a flexible and open solution that enables organizations to deploy high-performance and flexible visibility applications pervasively throughout their network for improved monitoring and troubleshooting. Enabled by the Network Operating System (NOS), this preconfigured Quick Emulator (QEMU) Kernel-based Virtual Machine (KVM) environment leverages high performance x86 CPUs to host these applications, extending visibility customized to the business and operational needs of the organization across the entire network.

The QEMU KVM environment supports several pretested and well-known packet capture applications in a Linux virtual machine, including Wireshark and tcpdump. There are a wide variety of additional applications, tools, and utilities that organizations are able to run in this environment, such as data analytics applications, packet generators, monitoring tools, troubleshooting utilities, and many others. While the QEMU KVM environment is open and can host any application, it is designed and ideally suited for networking applications, tools, and utilities.

IAH architecture supports the creation and use of virtualization domains, such as virtual machines, and Docker containers. This design creates a common-use host, which coordinates and automates multiple guest-networking functions into chains. The hardware boots into the virtual Linux OS, providing the ability to run additional applications or services within a specific virtual machine or a Docker container, and simultaneously supporting the regular functionality of the switch.

Yet Another Next Generation (YANG) model is used to manage configuration and retrieve operational data. You access the YANG model through Representational State Transfer Configuration Protocol (RESTCONF) using a northbound interface, namely

ExtremeCloud IQ Site Engine, that provides an additional way to configure and monitor the switch. For more information on RESTCONF, see [Representational State Transfer Configuration Protocol \(RESTCONF\) Fundamentals](#) on page 3136.

Virtual Services Resources

The virtual services resources are isolated from each other, as well as from the Network Operating System (NOS) running the switch.

The resources available for all virtual services on 7520 Series and 7720 Series switches are as follows:

- Six Central Processing Unit (CPU) cores
- 12 GB Random Access Memory (RAM)
- 100 GB Solid State Drive (SSD) flash memory

The resources available for all virtual services on 5720-24MXW and 5720-48MXW switches are as follows:



Note

You must install a modular SSD unit to use virtual services on 5720-24MXW and 5720-48MXW switches.

- Two CPU cores
- 4 GB RAM
- 120 GB SSD flash memory (separately available modular SSD unit), with 104 GB dedicated for IAH storage.

The switch OS uses the following resources on 5720-24MXW and 5720-48MXW:

- Two CPU cores
- 4 GB RAM
- 8 GB internal flash memory storage

Extreme Integrated Application Hosting Ports

Extreme Integrated Application Hosting (IAH) ports are labeled as Insight ports, which are internal ports used to support Ethernet connectivity by the virtual services configured on the switch. IAH ports operate at 10 Gigabits per second (Gbps). The following features support IAH ports on the switch:

- VLANs
- Filters
- Port Statistics
- Basic Interface Configuration
- Mirroring

- Switched UNI
- Transparent Port UNI



Note

Network-to-network interface (NNI) support is not available for IAH ports. IS-IS adjacencies cannot be established on IAH ports.

For information about how to configure IAH ports, see the following tasks:

- [Configure a Virtual Service](#) on page 1113
- [Configure Virtual Ports](#) on page 1130

Connection Types

The VM and Docker virtual ports map to a physical Extreme Integrated Application Hosting port using the following connection types:

- Open vSwitch (OVS)
- Single Root I/O Virtualization (SR-IOV)
- Virtualization Technology for Directed I/O (VT-d)



Note

You must enable trunking on the Extreme Integrated Application Hosting port when you use SR-IOV and OVS connection types. For more information about enabling trunking, see [MultiLink Trunking and Split MultiLink Trunking](#) on page 2657.

You can configure Extreme Integrated Application Hosting ports 1/s1 and 1/s2 to accommodate different connect types. Extreme Integrated Application Hosting ports 1/s1 and 1/s2 can accommodate virtual ports of SR-IOV, OVS, or VT-d connect types as shown in the following table. Using the **virtual-service** command, you can specify which Extreme Integrated Application Hosting port is associated with the configured connect type. You can also configure the Network Interface Card (NIC) type of the virtual port using the **virtual-service** command.

The following table lists the compatible Extreme Integrated Application Hosting port connect type configurations.

Extreme Integrated Application Hosting port 1/s1	Extreme Integrated Application Hosting port 1/s2
SR-IOV	OVS
SR-IOV	SR-IOV
SR-IOV	VT-d
OVS	SR-IOV
OVS	OVS
OVS	VT-d
VT-d	VT-d

Extreme Integrated Application Hosting port 1/s1	Extreme Integrated Application Hosting port 1/s2
VT-d	SR-IOV
VT-d	OVS

Link Flapping

When the switch initializes, the Extreme Integrated Application Hosting ports connect to the underlying Linux hypervisor. When a virtual port of connection type OVS or SR-IOV is configured on the switch, the Linux hypervisor saves this connection, and the link state of the Extreme Integrated Application Hosting port does not change. However, when a virtual port of connection type VT-d is configured on the switch, control of the Extreme Integrated Application Hosting port is passed from the Linux hypervisor to the configured Virtual Machine (VM). The Extreme Integrated Application Hosting port flaps due to this transition, and the switch reports it in the system log. The Extreme Integrated Application Hosting port flaps twice during the transition:

1. when the Extreme Integrated Application Hosting port is removed from the Linux hypervisor.
2. when the Extreme Integrated Application Hosting port is added to the VM.

A similar link flap sequence takes place on the Extreme Integrated Application Hosting port when the associated VM is disabled on the switch, and the control of the Extreme Integrated Application Hosting port is passed from the VM back to the Linux hypervisor.

Configuration Requirements

- To use an Extreme Integrated Application Hosting port as an analyzer port on a monitoring BEB for Fabric RSPAN (Mirror to I-SID), you must associate outer-tag 4091 to egress port 1/s1 or 1/s2 if the connect type is OVS or SR-IOV. Use the **monitor-by-isid <1-1000> map-to-vid <1-4093>** command to configure VLAN 4091 for Fabric RSPAN.
- To use an Extreme Integrated Application Hosting port with a connect type as OVS or SR-IOV for Port Mirroring, associate VLAN 4091 to the virtual machine (VM) vport to send the mirrored packets to the VM.
- To enable Flex UNI on an Extreme Integrated Application Hosting port with a connect type of VT-d, enable dot1q encapsulation on the VM interfaces. Flex UNI enables tagging on these ports by default; you must tag the VM ports with the VLANs that these ports use.

Operational Considerations and Restrictions

Consider the following when deploying Extreme Integrated Application Hosting (IAH) on various switches:

Table 105: Operational Considerations

	5720 Series	7520 Series	7720 Series
Number of IAH ports	5720-24MXW: 2 5720-48MXW: 2	7520-48XT-6C: 1 7520-48Y-8C: 1 7520-48YE-8CE: 1	7720-32C: 2
Multiple simultaneous VMs	Supported	Supported	Supported
Pre-installed VM	Third Party Virtual Machine Fabric IPsec Gateway	Third Party Virtual Machine Fabric IPsec Gateway	Third Party Virtual Machine Fabric IPsec Gateway
Additional components required	None	None	None

Third Party Virtual Machine

The Extreme Integrated Application Hosting (IAH) feature supports the pre-installed Third Party Virtual Machine (TPVM). For switches that use a modular Solid State Drive (SSD) for IAH, the virtual machine is pre-installed on the modular SSD. For switches that do not use a modular Solid State Drive (SSD) for IAH, the virtual machine is pre-installed on the switch.



Note

The Third Party Virtual Machine (TPVM) version is based on Ubuntu 24.04.

You can use the **show virtual-service config** command to view the information about the pre-installed virtual machine on the switch. For more information, see [Display Virtual Service Configuration](#) on page 1120.



Important

You must upgrade virtual services independently of switch software upgrade; separate images for virtual services are available. For more information, see [Upgrade a Virtual Service](#) on page 1125.

For more information about how to configure virtual services, see [Virtual Services Configuration using CLI](#) on page 1112 and [Virtual Services Configuration using EDM](#) on page 1126.



Note

You cannot use SFTP to transfer files larger than 2 GB to the switch. You need to use SCP or FTP.

Third Party Virtual Machine (TPVM) provides a set of troubleshooting tools on the switch. The following installed packages are available on TPVM:

- build-essential
- checkinstall
- iperf
- mtools
- netperf
- qemu-guest-agent
- tshark
- valgrind
- vim-gnome with vim-gtk3
- wireshark
- xterm
- isc-dhcp-client with dhcpcd
- isc-dhcp-server with kea
- iperf3
- libpcap
- rpcapd
- resolvconf



Important

TPVM includes an administrator account with a default user name and password. To ensure security, you must change the default password when you access TPVM for the first time, before enabling the IAH ports using the **no shutdown** command. The software automatically prompts you to change this password at first boot; no action can be taken with the VM until you change the password.

The following user applications are available on TPVM:

- Dynamic Host Configuration Protocol (DHCP) server
- Domain Name Server (DNS)
- Authentication, authorization, and accounting (AAA) server for Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control Service Plus (TACACS+).
- Syslog server
- Simple Network Management Protocol (SNMP) trap receiver
- Suricata - a free and open-source robust network threat detection engine that provides real time intrusion detection (IDS), inline intrusion prevention (IPS), network security monitoring (NSM), and offline packet capture (pcap) processing.
- Wireshark – a protocol analyzer that provides packet capturing and analysis.
- Ostinato - provides packet crafting, network traffic generation, and analysis with a user-friendly Graphical User Interface (GUI).

The following are the virtual services resources for TPVM:

- Two CPU cores
- 4 GB RAM
- One virtual port of VT-d connection type
- 1.8 GB up to 32 GB SSD

**Note**

To use this feature on the applicable models of 5720 Series, you must install an SSD module in the switch.

Fabric IPsec Gateway Fundamentals

The Fabric IPsec Gateway feature introduces a Virtual Machine (VM) that supports aggregation of Fabric Extend Tunnels with fragmentation, reassembly, and Internet Protocol Security (IPsec) encryption functions.

The minimum configuration requirements for the Fabric IPsec Gateway VM are as follows:

- Two Central Processing Unit (CPU) cores
- 4 GB Random Access Memory (RAM)
- One Virtualization Technology for Directed I/O (VT-d) vport (eth0)
- Minimum 10 GB SSD

**Note**

To use this feature on the applicable models of 5720 Series, you must install an SSD module in the switch

To configure IPsec on a switch through the Fabric IPsec Gateway VM, see [Fabric IPsec Gateway Configuration using CLI](#) on page 1135.

Fabric IPsec Gateway supports the following services through the VM:

- IPsec with fragmentation and reassembly - for the VXLAN traffic that needs IPsec, the network routes the packets through the Fabric IPsec Gateway VM that provides IPsec encryption and decryption for VXLAN packets. The system also supports fragmentation and reassembly for IPsec tunnels that you configure on the VM, and a minimum of 1300 bytes of Maximum Transmission Unit (MTU) value. You can configure fragmentation to occur before the packets are encrypted.
- Fragmentation and reassembly - the Fabric IPsec Gateway VM performs fragmentation and reassembly for VXLAN and IPsec tunnels, for which the network routes the packets through the VM. The system supports a minimum of 750 bytes of Maximum Transmission Unit (MTU) value.

Fabric Extend Over IPsec Limitations

Review the following limitations associated with Fabric Extend:

- Only pre-shared authentication key IPsec parameters are user configurable. Other, third-party solutions are not configurable.
- Supports IKEv2 protocol key exchange only.
- IPsec support is only added for Fabric Extend tunnels.
- IPsec is not supported for regular Layer 3 routed packets.

Egress Shaping for Fabric Extend Tunnels on Fabric IPsec Gateway

Table 106: Egress Tunnel Shaping product support

Feature	Product	Release introduced
Egress Shaper for Fabric Extend tunnels on Fabric IPsec Gateway virtual machine	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3.1 Supported on VSP4900-12MXU-12XE and VSP4900-24XE
	VSP 7400 Series	VOSS 8.3.1

You can configure the egress shaping rate to limit egress bandwidth for tunnels on the Fabric IPsec Gateway Virtual Machine (VM).

Egress Tunnel Shaping shapes unicast, multicast, and unknown unicast egress traffic on the tunnel.

Eight traffic queues are created for the tunnel traffic when a new shaper is configured. Packets are mapped to any of these queues based on the internal Class of Service (CoS) of the packets. When multiple tunnels use the common shaper value, an internally

unique shaper is associated for each tunnel. The following table shows the default minimum weight for each queue.

Table 107: Default minimum weight

Queue	Default minimum weight
0	5
1	20
2	30
3	40
4	50
5	50
6	Rate limited to 50% of configured shaper rate
7	Rate limited to 5% of configured shaper rate

Queues 6 and 7 are rate limited and have a higher priority. The maximum traffic allowed out of queue 6 and queue 7 would be shape rate/2 (50% bandwidth) and shape rate/20 (5% bandwidth).

Considerations

Consider the following when you configure the egress shaper rate:

- If the ingress data traffic receives excessive packets with the following DiffServ Code Point (DSCP) or 802.1p values (high priority control packets) and egress shaping is configured, a IS-IS flap can be seen.

DSCP value	802.1p value
0x28	6
0x2E	6
0x2F	6
0x30	7
0x38	7

- The egress tunnel shaping rate is impacted if the incoming packet size is greater than the Fabric Extend tunnel MTU. This is due to the additional packet header required for fragmentation.

Digital Certificates for Fabric IPsec Gateway

Table 108: Digital Certificates for IPsec Authentication product support

Feature	Product	Release introduced
Digital Certificates for IPsec Authentication	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW For Fabric Extend over IPsec
	7520 Series	Fabric Engine 8.10 For Fabric Extend over IPsec
	7720 Series	Fabric Engine 8.10 For Fabric Extend over IPsec
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3 Supported on VSP4900-12MXU-12XE and VSP4900-24XE For Fabric Extend over IPsec
	VSP 7400 Series	VOSS 8.3 For Fabric Extend over IPsec

Fabric IPsec Gateway supports digital certificates for IPsec authentication of Fabric Extend tunnels. To support different certificates for different IPsec tunnels, you can configure multiple certificate authority (CA) trustpoints and identity subject certificates.

If you are not familiar with digital certificates, see [Digital Certificate/PKI](#) on page 3380 for additional background information like digital certificate terminology.

Fabric IPsec Gateway supports multiple CA trustpoints and multiple identity subject certificates. You can use different certificates for different IPsec tunnels. Fabric IPsec Gateway acts like a hub to isolate IPsec domains. To use Public Key Infrastructure (PKI) with IPsec Fabric Extend technology, all devices must acquire the digital-signed certificates. The CA server can be accessed from the devices, a public network, or an internal network. Each device must configure a profile for the CA server. The switch uses Simple Certificate Enrollment Protocol (SCEP) to obtain the trusted, signed certificates.

To use IPsec with Digital Certificates:

- Configure the Fabric Extend tunnels.
- Configure the authentication method as RSA-signature. For more information, see [Configure Public Key Infrastructure for IPsec Tunnels](#) on page 1147.
- Configure certificate information in Fabric IPsec Gateway.

Fabric IPsec Gateway supports both online and offline certificate management simultaneously.

Online Certificate Provisioning

The switch uses IPsec Simple Certificate Enrollment Protocol (SCEP) to obtain the CA certificate, and then validates the CA certificate against the certificate chain.



Note

Extreme Networks validated the Fabric IPsec Gateway SCEP implementation with EJBCA CA Server only. Fabric IPsec Gateway SCEP cannot currently use Win CA like digital certificate support in the network operating system on the switch.

Use trustpoints to manage and track CAs and certificates. The switch can enroll with a trustpoint to obtain an identity certificate. You must configure the CA URL, the CA common name, and select the HTTP request type to configure the CA server trustpoint.

Configure the certificate subject parameters to provide the device distinguished name (DN) and key name for the generated key pair (the private key). If you do not configure a private key, the switch generates one. The switch validates the returned certificate against the trustpoint's CA certificate.

You can remove subject certificates from the CA trustpoint or clean the CA trustpoint only if the subject-label is not configured on an IPsec tunnel.

Offline Certificate Provisioning

Offline certificate management supports switches that cannot communicate with the CA to obtain the identity certificate online by certificate enrollment operation.

The switch generates the certificate signing request (CSR) using the subject DN and the private key that you configure in the CLI. If you do not configure a private key, the switch generates one.

Transfer the CSR to the offline CA to be signed. Retrieve the signed certificate to validate against the original CSR. You must manually transfer all certificates in the certificate chain to the switch. The signed certificate must include the subject-label to map it to a locally-generated CSR for validation.

You must manually download Certificate Revocation List (CRL) files. You can remove offline subject certificates only if the subject-label is not configured on an IPsec tunnel.

IPsec NAT-T

Table 109: IPsec NAT-T product support

Feature	Product	Release introduced
IPsec NAT-T	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW Supported using Fabric IPsec Gateway
	7520 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7720 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3 Supported on VSP4900-12MXU-12XE and VSP4900-24XE Supported using Fabric IPsec Gateway
	VSP 7400 Series	VOSS 8.2 Supported using Fabric IPsec Gateway

IP security (IPsec) Network Address Translation Traversal (NAT-T) allows IPsec tunnel traffic through Network Address Translation (NAT) or Port Address Translation (PAT) points in the network.

IPsec NAT-T Considerations

The following considerations apply to IPsec NAT-T:

- You must configure one side of the IPsec NAT-T tunnel as an IPsec responder. If IPsec is configured on the IPsec Initiator device and subsequently configured on the IPsec Responder device, IPsec must be restarted on the Initiator device. If IPsec is not restarted, it can take approximately 3 minutes for the adjacency to open.
- You must configure the aggregator device as the IPsec Responder device, and configure the branch device as the IPsec Initiator device.
- Among all the IPsec responders, the system uses the lowest configured maximum transmission unit (MTU) value of any responder IPsec tunnel as the MTU value for all IPsec responder-only tunnels. The system uses the lowest configured IPsec tunnel

MTU value regardless of manually configured MTU tunnel values. For non-responder IPsec tunnels or VXLAN tunnels, the configured and visible MTU value for the tunnel is used for fragmentation and reassembly.

- If both the IPsec Initiator device and the IPsec Responder device are behind NAT, you must configure the IPsec Initiator device with the public IP address of the NAT router connected to the IPsec Responder device.
- You must add route table entries on the IPsec Responder device with the public IP address and private IP address of the remote NAT for the IPsec Initiator device. A configured route table is required for IPsec NAT-T Fabric Extend (FE) connectivity.
- You must add a route table entry on the IPsec Initiator device with the public IPsec Remote NAT IP address for the IPsec Responder device. A configured route table is required for IPsec NAT-T Fabric Extend (FE) connectivity.

IPsec Compression

Table 110: Fabric Extend - IPsec Compression product support

Feature	Product	Release introduced
IPsec compression	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW Supported using Fabric IPsec Gateway
	7520 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7720 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.7 Supported on VSP4900-12MXU-12XE and VSP4900-24XE Supported using Fabric IPsec Gateway
	VSP 7400 Series	VOSS 8.7 Supported using Fabric IPsec Gateway

IPsec compression reduces the size of the IP datagram to improve the communication performance between hosts connected behind Backbone Edge Bridges (BEB).

**Tip**

As a best practice, use IPsec compression only for Fabric Extend tunnels where latency is greater than 70ms.

The following list identifies how you can implement IPsec compression:

- You can configure multiple IPsec Fabric Extend (FE) adjacencies with or without compression simultaneously.
- You must enable IPsec compression on both BEBs to use IPsec compression for an FE adjacency.
- You cannot configure IPsec compression if fragmentation before encryption is already enabled.
- You can change the IPsec compression configuration only if IPsec is disabled.

To configure IPsec compression, see [Configure IPsec Compression](#) on page 1146.

IPsec Fragmentation Before Encryption

Table 111: Fabric Extend - IPsec Fragmentation before Encryption product support

Feature	Product	Release introduced
IPsec fragmentation before encryption	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW Supported using Fabric IPsec Gateway
	7520 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7720 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3.1 Supported on VSP4900-12MXU-12XE and VSP4900-24XE Supported using Fabric IPsec Gateway
	VSP 7400 Series	VOSS 8.3.1 Supported using Fabric IPsec Gateway

The best practice is to enable fragmentation before encryption only for an IPsec adjacency over a WAN.

Configure IPsec fragmentation of the packets to occur before encryption and IPsec encapsulation. Packets are fragmented based on the tunnel maximum transmission unit (MTU) without the IPsec header so that the final packet does not exceed the tunnel MTU. The MTU value is a per tunnel configuration, which means packet fragmentation occurs per tunnel. For a tunnel with this functionality enabled, packets that egress the specific network-to-network interface (NNI) port are encapsulating security payload (ESP) packets only.



Note

You cannot configure IPsec compression if fragmentation before encryption is already enabled.

For more information, see [Enable Fragmentation Before Encryption on Fabric IPsec Gateway VM](#) on page 1147.

Virtual Services Configuration using CLI

Perform the procedures in this section to configure Extreme Integrated Application Hosting (IAH) virtual services on the switch using the command line interface (CLI).

Access a Virtual Service Console

The virtual services running on a Virtual Machine (VM) require a console for configuration and monitoring purposes.

About This Task

Perform this procedure to access the virtual service console port for the specific VM.

This command also works in Global Configuration mode but does not display if you use the question mark (?) help in that mode.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Enter the following command to access the virtual service console:
`virtual-service WORD<1-128> console`



Note

Type CTRL+Y to exit the console.

Example

```
Switch:1>enable
Switch:1#virtual-service tpvm console
```

Variable Definitions

The following table defines parameters for the **virtual-service** command.

Variable	Value
<code>WORD<1-128></code>	Specifies the virtual service name.
<code>console</code>	Accesses the console for the specific virtual service.

Install a Virtual Service

A virtual service provides the ability to support additional applications or services and simultaneously support the regular switching functionality. Each virtual service provides an Open Virtual Appliance (OVA) image, which is installed on Extreme Integrated Application Hosting (IAH) through ExtremeCloud IQ Site Engine.

Before You Begin

Use FTP, SFTP, or SCP to transfer the OVA image to the `/var/lib/insight/packages/` directory on the switch.



Note

The Fabric IPsec Gateway image includes no integrity check. Use SCP to copy the file to the switch and confirm the file size before installation.

About This Task

Perform this procedure to install a package file to a specific location indicated by a virtual service name. This procedure also verifies if the package is in OVA format, and if a certificate is provided in the package.

This command also works in Global Configuration mode but does not display if you use the question mark (?) help in that mode.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Install the virtual service package:

```
virtual-service WORD<1-128> install package WORD<1-512>
```

Variable Definitions

The following table defines parameters for the **virtual-service** command.

Variable	Value
<i>WORD<1-128></i>	Specifies the virtual service name.
<i>install</i>	Installs the virtual service package.
<i>package WORD<1-512></i>	Specifies the package name and path.

Configure a Virtual Service

Before You Begin

- You must enable trunking on the Extreme Integrated Application Hosting (IAH) port when you use SR-IOV and OVS connection types. For more information about enabling trunking, see [MultiLink Trunking and Split MultiLink Trunking](#) on page 2657.
- Ensure the switch has the Ethernet drivers installed as per the SR-IOV standard, to support the VT-d and the SR-IOV connection type for the configured virtual ports.

About This Task

Perform this procedure to configure a virtual service on the switch.



Note

- Following procedure lists the general sequence to configure a virtual service.
- The names of Ethernet ports appearing in a specific Virtual Machine (VM) are not correlated to the configured virtual port names. Each VM renames the Ethernet ports as per its requirements, after they are discovered during the VM initialization.
- By default, the system displays all virtual ports of OVS connection type first in the alphabetical order of their configured names, followed by the virtual ports of SR-IOV and VT-d connection types.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a VLAN:



Note

Virtual service configuration supports port-based VLANs only.

```
vlan create <2-4059> name WORD<0-64> type {port-mstprstp <0-63>}  
[color <0-32>]
```

3. Add the IAH and faceplate port to the VLAN:

```
vlan members add <1-4059> {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

4. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

5. Enable the IAH and faceplate ports:

```
no shutdown
```

6. Exit to Global Configuration mode:

```
exit
```

7. (Optional) Create a virtual service:

```
virtual-service WORD<1-128>
```


8. (Optional) Configure the number of CPU cores to be assigned to the virtual service created:

```
virtual-service WORD<1-128> num-cores <num_cores>
```

9. (Optional) Configure the memory size to be assigned to the virtual service created:

```
virtual-service WORD<1-128> mem-size <mem-size>
```

10. (Optional) Configure the disk to be assigned to the virtual service created:

```
virtual-service WORD<1-128> disk WORD<1-32> size <1-30>
```

11. Configure the virtual port connection type:



Note

Ensure the connection type you configure for the virtual port matches the connection type supported by the IAH port.

```
virtual-service WORD<1-128> vport WORD<1-32> connect-type {ovs | sriov  
| vtd}
```

12. Configure the IAH port to associate with the connection type:

```
virtual-service WORD<1-128> vport WORD<1-32> port WORD<1-32>
```



Important

You cannot configure two virtual services with conflicting connect types on the same IAH port. You cannot configure two virtual services with VT-d connect type on the same IAH port.

13. Configure the NIC type of the IAH port:

```
virtual-service WORD<1-128> vport WORD<1-32> port WORD<1-32> nic-type  
{virtio | e1000}
```

14. Add the virtual port to the VLAN created:

```
virtual-service WORD<1-128> vport WORD<1-32> vlan <1-4096>
```

15. Enable the virtual service:

```
virtual-service WORD<1-128> enable
```

Example

Configuring the TPVM virtual service using IAH port 1/s1 with an SR-IOV connection type:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 1/s1
Switch:1(config-if)#encapsulation dot1q
Switch:1(config-if)#exit
Switch:1(config)#vlan create 10 name tpvm-lan-vlan type port-mstprstp 0
Switch:1(config)#vlan members add 10 1/s1,1/6/2
Switch:1(config)#interface GigabitEthernet 1/s1,1/6/2
Switch:1(config-if)#no shutdown
Switch:1(config-if)#exit
Switch:1(config)#virtual-service tpvm vport eth0 connect-type sriov
Switch:1(config)#virtual-service tpvm vport eth0 vlan 10
Switch:1(config)#virtual-service tpvm enable
```

Configuring the TPVM virtual service on IAH port 1/s2 with a VT-d connection type:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#vlan create 10 type port-mstprstp 0
Switch:1(config)#vlan member add 10 1/1,1/s2
Switch:1(config)#interface GigabitEthernet 1/s2,1/1
Switch:1(config-if)#no shutdown
Switch:1(config-if)#exit
Switch:1(config-if)virtual-service tpvm vport eth0 port 1/s2
Switch:1(config)#virtual-service tpvm enable
```

Variable Definitions

The following table defines parameters for the **vlan create** command.

Variable	Value
<2-4059>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.
color<0-32>	Specifies the color of the VLAN.
nameWORD<0-64>	Specifies a name for the VLAN to be created.
type {port-mstprstp<0-63>}	Creates a VLAN by port, with the STP instance ID ranging from 0 to 63. Note: MSTI instance 62 is reserved for SPBM if SPBM is enabled on the switch.

The following table defines parameters for the **vlan members** command.

Variable	Value
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code>add</code>	Adds ports to a specified VLAN ID.

The following table defines parameters for the **virtual-service** command.

Variable	Value
<code>WORD<1-128></code>	Specifies a name for virtual service.
<code>connect-type {ovs sriov vtd}</code>	Specifies the connection type for the virtual port created. The default is VT-d. The switch supports the following maximums for virtual ports: <ul style="list-style-type: none"> • OVS - 16 • SR-IOV - 16 • VT-d - 2
<code>disk WORD<1-32></code>	Specifies the disk assigned to the virtual service.
<code>mem-size <1-5120></code>	Specifies the memory size in Megabytes assigned to the virtual service. The default value is 1024 Megabytes.
<code>nic-type [virtio e1000]</code>	Specifies the Virtual Port NIC type. The default is e1000. Note: Configure this value only when the connect-type parameter is ovs .

Variable	Value
<code>num-cores <num_cores></code>	<p><code><num_cores></code> specifies the number of cores assigned to the virtual service. Different platforms support different ranges.</p> <p>Note: 5720-24MXW and 5720-48MXW support <code><1-2></code> cores. 7520 Series and 7720 Series support <code><1-6></code> cores.</p> <p>The default value is 1.</p>
<code>port WORD<1-32></code>	<p>Specifies the name of the Extreme Integrated Application Hosting (IAH) port associated with the virtual port. Depending on the hardware, the switch can support the following IAH ports:</p> <ul style="list-style-type: none"> • 1/s1 • 1/s2
<code>size <1-30></code>	Specifies the size of the disk in Gigabytes.
<code>vlan <1-4096></code>	Specifies the VLAN ID used by the virtual port.
<code>vport WORD<1-32></code>	Specifies the name of the virtual port.

Shut Down a Virtual Service

About This Task

Perform this procedure to disable the virtual service.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Disable the virtual service:

```
no virtual-service WORD<1-128> enable
```

Example

Disable the virtual service.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no virtual-service tpmv enable
```

Delete Virtual Service Resources

About This Task

Perform this procedure to delete the virtual service resource allocation.



Note

If a corresponding virtual machine is running, it is stopped, and then the virtual service configuration is deleted.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Delete the virtual service resource allocation:
`no virtual-service WORD<1-128> [disk WORD<1-32>] [vport WORD<1-32>]`

Example

Delete all virtual service resource allocation.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#no virtual-service tpvm
```

Uninstall a Virtual Service

Before You Begin

You must disable the virtual service before you uninstall it.

About This Task

Perform this procedure to uninstall a configured virtual service.



Note

If a virtual machine is running, it is stopped, and then the service directory is uninstalled.

This command also works in Global Configuration mode but does not display if you use the question mark (?) help in that mode.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Uninstall a specific virtual service:
`virtual-service WORD<1-128> uninstall`

Example

```
Switch:1>enable
Switch:1#virtual-service tpvm uninstall
```

Variable Definitions

Use data in the following table to use the **virtual-service** command.

Variable	Value
<i>WORD<1-128></i>	Specifies the virtual service name.
<i>uninstall</i>	Uninstalls the specified virtual service name.

Display Virtual Service Configuration**About This Task**

Perform this procedure to display the virtual service configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the virtual-service configuration:

```
show virtual-service config [WORD<1-128>]
```

Example

Display the configuration of all virtual services:

```
Switch:1>show virtual-service config
=====
                        Installed Packages
=====
Package:                FIGW-SHAPE
Package App Name:       FabricIPSecGW_VM_5.2.0.0_20.04
Package Version:        5.2.0.0
Package Name:           FabricIPSecGW_VM_5.2.0.0_20.04.ova
=====

                        Virtual Services Config
=====

Virtual Service :FIGW-SHAPE
Memory Assigned(M) : 8196
Number of Cores      : 6
Additional Disk Assigned:

VPort Information:
Name                 Vlan          Connect Type  Insight Port  NIC Type
eth0                 1             vtd           1/s1

Management Status :           Enabled
=====
```

Display Virtual Service Installation Status

About This Task

Perform this procedure to display the installation status for the specific virtual service. This procedure indicates if the installation finished successfully or failed to complete.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display installation status for a specific virtual service:

```
show virtual-service install WORD<1-128>
```

Example

Display installation status for a specific virtual service:

```
Switch:1>show virtual-service install tpvm
      Stage:   Convert
      Status:  In Progress
```

Display Virtual Services Resources

About This Task

Perform the following procedure to display the number of remaining virtual services resources on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display statistics for all virtual services configured on the switch or a specific virtual service:

```
show virtual-service statistics [WORD<1-128>]
```

Example

```
Switch:1>show virtual-service statistics
=====
                        Virtual Services
=====
Virtual Service :      figw
Package App Name :      FabricIPSecGW_VM_master.0.22_20.04
Package Name :      FabricIPSecGW_VM_master.0.22_20.04.ova
Package Version :      0.22

Memory Utilization (Mega Bytes)
  Allocated      Used      Available
    12288         1209      11079

CPU Utilization
  Allocated(# cores)      CPU Utilization (Total %)
             6              12

Disk Utilization
  Primary Disk Size :      10G

VPort Information:
  Name                  Vlan          Connect Type      Insight Port      NIC Type
```

```
eth0                                vtd                                1/s1
  Guest Intf Name : eth0
  MAC Address    : 42:fd:46:00:00:01
  IPv4 Address   : 0.0.0.0
  IPv6 Address   : fe80:0:0:0:40fd:46ff:fe00:1

Management Status : Enabled
Operational Status : Running
Uptime   :        0 day(s), 08:49:30
-----
=====
                          Hypervisor Remaining Resources
=====
Number of Cores Remaining: 0
Total Memory Remaining(M): 147
Total Disk Remaining(GB): 79
```

Run a VM command from Network Operating System (NOS) CLI

About This Task

Perform this task to run a virtual machine (VM) command from the NOS CLI.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Run the **ls** command for the VM configs directory from the CLI:
`virtual-service WORD<1-128> exec-command WORD<1-256>`
3. Run a Fabric IPsec Gateway command from the CLI:
`virtual-service WORD<1-128> figw-cli WORD<1-256>`



Note

To use this command, the Fabric IPsec Gateway image must be at least 5.0.0.0.

Examples

From NOS, list the contents of the `home/rwa/configs` directory in the Fabric IPsec Gateway VM:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#virtual-service figw exec-command "ls /home/rwa/configs"
config.cfg
figw_cli.log
shadov.txt
```


From NOS, configure the source VLAN ID for the IPsec tunnel in the Fabric IPsec Gateway VM:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#virtual-service figw figw-cli "set global ipsec-tunnel-src-vlan 71"
```

Variable Definitions

The following table defines parameters for the **virtual-service** command.

Variable	Value
<i>WORD</i> <1-128>	Specifies the virtual service name.
<i>WORD</i> <1-256>	Specifies the VM command to run. To include spaces in the syntax, include the text string in quotation marks (").

Copy VM Files

About This Task

Perform this task to copy files between the Network Operating System (NOS) and a VM, or between VMs.

This command also works in Global Configuration mode but does not display if you use the question mark (?) help in that mode.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Copy a file:
virtual-service copy-file *WORD*<1-256> *WORD*<1-256>

Examples

Copy a file from the NOS to a VM:

```
Switch:1>enable
Switch:1#virtual-service copy-file /intflash/config_figw.cfg figw:/home/rwa/configs/config.cfg
```

Copy a file from a VM to the NOS:

```
Switch:1>enable
Switch:1#virtual-service copy-file figw:/home/rwa/configs/config.cfg /intflash/config_figw.cfg
```

Copy a file between VMs:

```
Switch:1>enable
Switch:1#virtual-service copy-file figw:/home/rwa/configs/config.cfg figw2:/home/rwa/configs/config.cfg
```

Variable Definitions

The following table defines parameters for the **virtual-service copy-file** command.

Variable	Value
<i>WORD</i> <1-256>	<p>Specifies the source and destination file to copy.</p> <p>To specify a VM location, use the format <VM_name>:<VM_file_path/file name>.</p> <p>To specify a NOS location, use the format </file_path/file name> where the valid path can be one of the following:</p> <ul style="list-style-type: none">• /intflash• /extflash• /usb• /var/lib/insight/packages

Change a VM User Password from Network Operating System (NOS) CLI

About This Task

Perform this task to change the password for a VM user. The password must be greater than, or equal to, 8 characters.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Change the password:

```
virtual-service WORD<1-128> change-user-pass WORD<1-20>
```
3. Enter the new password.
4. Enter the new password a second time.

Example

Change the password for the rwa user account in the Fabric IPsec Gateway VM:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#virtual-service figw change-user-pass rwa
Enter password : *****
Re-enter password : *****
```

Variable Definitions

The following table defines parameters for the **virtual-service WORD<1-128> change-user-pass** command.

Variable	Value
WORD<1-20>	Specifies the user name.
WORD<1-128>	Specifies the virtual service name.

Upgrade a Virtual Service

If Extreme Networks makes a new version of the virtual service available, uninstall the original virtual service and install the newer virtual service. The following generic procedure can apply to all virtual services. For a Fabric IPsec Gateway-specific procedure, see [Upgrade a Fabric IPsec Gateway VM](#) on page 1135.



Important

You can perform an upgrade of Linux inside the virtual service by standard Linux upgrade procedures. For example, TPVM is Ubuntu based, so you can use **sudo apt-get update** and **sudo apt-get upgrade**. If you complete such an upgrade, Extreme Networks is not responsible for the behavior of the VM; it has not been tested with every version of the network operating system (NOS).

Before You Begin

- Ensure the new virtual service image version is compatible with the NOS release that runs on the switch. Only use the virtual service image version that is posted with the NOS release image. For more information about image files in a release, see *Fabric Engine Release Notes*. If necessary, upgrade the NOS image before you upgrade the virtual service image.
- If you installed applications in the Third Party Virtual Machine (TPVM), you must migrate important data for those applications before you perform this procedure.
- If you created new users in the TPVM, follow standard Linux procedures to back up user names and passwords.
- For Fabric IPsec Gateway, back up the configuration files (*.cfg) and the `shadow.txt` file, which is an encrypted file that contains the authentication keys for the IPsec tunnels. You can use the **ls** command within the VM to see the file list. Use FTP within the VM to transfer the files for backup or see [Copy VM Files](#) on page 1123.
- Use FTP or SFTP to transfer the new OVA image to the `/var/lib/insight/packages/` directory on the switch.



Note

The Fabric IPsec Gateway image includes no integrity check. Use SCP to copy the file to the switch and confirm the file size before installation.

About This Task

When you uninstall the original virtual service, the system removes the complete virtual service configuration from the configuration file.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Disable the virtual service:
`no virtual-service WORD<1-128> enable`
3. Uninstall the virtual service:
`virtual-service WORD<1-128> uninstall`
4. Install the virtual service package using the new OVA image:
`virtual-service WORD<1-128> install package WORD<1-512>`
5. Reconfigure the virtual service; for more information, see [Configure a Virtual Service](#) on page 1113.
6. Remove the original OVA image from the `/var/lib/insight/packages/` directory on the switch:
`remove WORD<1-255>`

Virtual Services Configuration using EDM

Perform the procedures in this section to configure Extreme Integrated Application Hosting (IAH) virtual services on the switch using the Enterprise Device Manager (EDM).

View Virtual Services Resources

About This Task

Perform the following procedure to view the number of remaining virtual services resources on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **Globals** tab.

Globals Field Descriptions

Use data in the following table to use the **Globals** tab.

Name	Description
DiskRemain	Shows the remaining disk space available, in Gigabytes (GB).
NumCoresRemain	Shows the remaining number of CPU cores available.
MemSizeRemain	Shows the remaining amount of memory size available, in Megabytes (MB).
CopySourceFile	<p>Specifies the source file to copy. To specify a location, use the format: {VM_NAME_SRC}:{VM_FILE_PATH} or {NOS_FILE_PATH}. For example, figw:/home/rwa/configs/ipsec1.cfg identifies a file located in the VM. /intflash/ipsec1.cfg identifies a file located in the NOS. The valid path for a NOS location can be one of the following:</p> <ul style="list-style-type: none"> • /intflash • /extflash • /usb • /var/lib/insight/packages
CopyDestinationFile	<p>Specifies the destination file to copy. To specify a location, use the format: {VM_NAME_DST}:{VM_FILE_PATH} or {NOS_FILE_PATH}. For example, figw:/home/rwa/configs/ipsec1.cfg identifies a file located in the VM. /intflash/ipsec1.cfg identifies a file located in the NOS. The valid path for a NOS location can be one of the following:</p> <ul style="list-style-type: none"> • /intflash • /extflash • /usb • /var/lib/insight/packages
CopyAction	<p>Specifies an action to copy a file from source to destination. The switch supports the following action:</p> <ul style="list-style-type: none"> • Copy • None
ScalarsName	Specifies the virtual service name. You must specify the virtual service name if you use this tab to change the password or run a command.

Name	Description
ExecuteCommand	Specifies the Virtual Machine (VM) command to run. To include spaces in the syntax, include the text string in quotation marks (").
User	Specifies the virtual service user name. The range is 0-20 characters.
Password	Specifies the virtual service password.
FigwCli	Specifies the command to send to the Fabric IPsec Gateway VM. For more information about Fabric IPsec Gateway commands, see <i>Fabric Engine Command References</i> .

Configure a Virtual Service

Before You Begin

You must configure at least one virtual port to enable the virtual service. For more information, see [Configure Virtual Ports](#) on page 1130.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **Virtual Service** tab.
4. Select **Insert**.
5. In **Name**, type a unique name.
6. (Optional) In **NumCores**, type a value.
7. (Optional) In **MemSize**, type a value.
8. Select **Insert**.
9. In **Enable** for the newly inserted row, change the value to true.
10. Select **Apply**.

Virtual Service Field Descriptions

Use data in the following table to use the **Virtual Service** tab.

Name	Description
Name	Specifies the name of the virtual service. Every virtual service must have a unique name.
NumCores	Specifies the number of CPU cores assigned to the virtual service. The default is 1.
MemSize	Specifies the memory size (in Megabytes) assigned to the virtual service. The default value is 1024 Megabytes.

Name	Description
Enable	Enables the virtual service. Note: You must configure at least one virtual port to enable the virtual service.
PackageInfoName	Shows the package name used by the virtual service.
PackageAppName	Shows the application name used by the virtual service.
PackageAppVersion	Shows the application version used by the virtual service.
UtilCpuAllot	Specifies the number of CPUs allocated to the virtual service.
UtilCpuUtil	Specifies the average percentage of CPU utilization over the past 30 seconds.
UtilMemAllot	Specifies the memory (in Megabytes) allocated to the virtual service.
UtilMemUsed	Specifies the memory used (in Megabytes) by the virtual service.
UtilMemAvailable	Specifies the memory available (in Megabytes) for the virtual service.
State	Specifies the operational state of the virtual service.
UpTime	Specifies the operational time of the virtual service.

Configure Disks to be Used by the Virtual Service

About This Task

Perform the following procedure to configure the number of disks to be used by the virtual service configured on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **Disks** tab.
4. Select **Insert**.
5. In **ServName**, type the virtual service name.
6. In **Name**, type the disk name.
7. (Optional) In **Size**, type a value.
8. Select **Insert**.

Disks Field Descriptions

Use data in the following table to use the **Disks** tab.

Name	Description
ServName	Specifies the virtual service name. Note: The specified name must match the virtual service name configured on the switch.
Name	Specifies the name of the disk used by the virtual service.
Size	Specifies the disk size (in Gigabytes). The default is 10 Gigabytes.
SizeAllot	Shows the disk size (in Megabytes) allocated to the virtual service.
SizeAvailable	Shows the available disk storage space (in Megabytes).
SizeUsed	Shows the amount of disk storage space (in Megabytes) used by the virtual service.

Configure Virtual Ports

Before You Begin

- You must enable trunking on the Extreme Integrated Application Hosting (IAH) port when you use SR-IOV and OVS connection types.
- Ensure the switch has the Ethernet drivers installed as per the SR-IOV standard, to support the VT-d and the SR-IOV connection type for the configured virtual ports.

About This Task

Perform the following procedure to configure virtual ports to be used by the virtual service configured on the switch.

**Note**

The names of Ethernet ports appearing in a specific Virtual Machine (VM) are not correlated to the configured virtual port names. Each VM renames the Ethernet ports as per its requirements, after they are discovered during the VM initialization.

By default, the system displays all virtual ports of OVS connection type first in the alphabetical order of their configured names, followed by the virtual ports of SR-IOV and VT-d connection types.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.

3. Select the **VPorts** tab.
4. Select **Insert**.
5. In **Virtual Service Name**, type the virtual service name.
6. In **Interface Name**, type a name for the virtual port.
7. (Optional) In **VlanIdList**, type a VLAN ID.
8. (Optional) In **ConnectType**, select a connection type.

**Note**

Ensure the connection type you configure for the virtual port matches the connection type supported by the IAH port.

9. Select **Insert**.

VPorts *Field Descriptions*

Use data in the following table to use the **VPorts** tab.

Name	Description
Virtual Service Name	Specifies the virtual service name. Note: The specified name must match the virtual service name configured on the switch.
Interface Name	Specifies the virtual port.
VlanIdList	Specifies the VLAN ID to which the virtual port is assigned.
ConnectType	Specifies the virtual port connect type. The default is VT-d. The switch supports the following maximums for virtual ports: <ul style="list-style-type: none">• OVS - 16• SR-IOV - 16• VT-d - 2

Name	Description
Port	Specifies the name of the Extreme Integrated Application Hosting port associated with the virtual port. Depending on hardware, the switch can support the following Extreme Integrated Application Hosting ports: <ul style="list-style-type: none">• 1/s1• 1/s2
NicType	Specifies the Virtual Port NIC type. The default is e1000. <ul style="list-style-type: none">• virtio• e1000 Note: Configure this value only when the ConnectType is ovs.

Install a Virtual Service

Before You Begin

- Use FTP or SFTP to transfer the OVA image to the `/var/lib/insight/packages/` directory on the switch.

About This Task

Perform the following procedure to configure the package information to be used by the virtual service.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **Application** tab.
4. Select **Insert**.
5. In **Name**, type the virtual service name.
6. Next to **PackageName**, select the ellipsis, select the package to install, and then select **Ok**.
7. Select **Insert**.

Application Field Descriptions

Use data in the following table to use the **Application** tab.

Name	Description
Name	Specifies the name of the virtual service.
PackageName	Specifies the name and location of the package.
InstallResult	Shows the status of the virtual service installation.
InstallStage	Shows the stages of a package installation.
PackageAppName	Shows the application name used by the virtual service.
PackageAppVersion	Shows the application version used by the virtual service.

Run a VM command from EDM

About This Task

Perform this task to run a VM command from EDM. To include spaces in the syntax, include the text string in quotation marks (").

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select the **Virtual Service**.
3. Select the **Globals** tab.
4. In **ScalarsName**, type a virtual service name.
5. In **ExecuteCommand**, type the **ls** command for the VM `configs` directory. For example, "`ls /home/rwa/configs`".
6. In **FigwCli**, type a command to send it to the Fabric IPsec Gateway. For more information about Fabric IPsec Gateway commands, see *Fabric Engine Command References*.
7. Select **Apply**.

Copy VM Files

About This Task

Perform this task to copy files between the Network Operating System (NOS) and a VM or between VMs. The valid path for a NOS location can be one of the following:

- `/intflash`
- `/extflash`
- `/usb`
- `/var/lib/insight/packages`

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **Globals** tab.
4. In **CopySourceFile**, type a location of the source file using the following format: `{VM_NAME_SRC}:{VM_FILE_PATH}` or `{NOS_FILE_PATH}`. For example, `figw:/home/rwa/configs/ipsec1.cfg` identifies a file located in the VM. `/intflash/ipsec1.cfg` identifies a file located in the NOS.
5. In **CopyDestinationFile**, type a location of the destination file using the following format: `{VM_NAME_DST}:{VM_FILE_PATH}` or `{NOS_FILE_PATH}`. For example, `figw:/home/rwa/configs/ipsec1.cfg` identifies a file located in the VM. `/intflash/ipsec1.cfg` identifies a file located in the NOS.
6. In **CopyAction**, select **Copy**.
7. Select **Apply**.

Change a VM User Password from EDM

About This Task

Perform this task to change the password for a VM user.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **Globals** tab.
4. In **ScalarsName**, type a virtual service name.
5. In **User**, type the user name.
6. In **Password**, type a password.
7. Select **Apply**.

View Virtual Services Package File Information

About This Task

Perform the following procedure to view information about the package files available in the `/var/lib/insight/packages` directory, which you can use to install a new virtual service.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **Virtual Service**.
3. Select the **PackageFile** tab.

PackageFile Field Descriptions

Use data in the following table to use the **PackageFile** tab.

Name	Description
Name	Shows the name and absolute path information for package files available in the <code>/var/lib/insight/packages</code> directory.
Date	Shows the date and time when the package file was added to the directory.
Size	Shows the size (in bytes) of the package file.


Fabric IPsec Gateway Configuration using CLI

Perform the procedures in this section to configure services like IPsec, fragmentation and reassembly, and to manage the Fabric IPsec Gateway Virtual Machine using the command line interface (CLI).

Upgrade a Fabric IPsec Gateway VM

If Extreme Networks makes a new version of the Fabric IPsec Gateway available, uninstall the original virtual service, and then install the newer virtual service.

Before You Begin

- Ensure the new virtual service image version is compatible with the NOS release that runs on the switch. Only use the virtual service image version that is posted with the NOS release image. For more information about image files in a release, see *Fabric Engine Release Notes*. If necessary, upgrade the NOS image before you upgrade the virtual service image.
-  **Note**
The Fabric IPsec Gateway image includes no integrity check. Use SCP to copy the file to the switch and confirm the file size before installation.

About This Task

Steps in this procedure include examples or links to background procedures if you are unfamiliar with how to complete a particular step.

Procedure

1. Within the VM, save the configuration. For more information, see [Save Running Configuration to a File](#) on page 1159.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#virtual-service figwOld console
Connected to domain figw5.2 Escape character is ^Y
FIGW> save config
```

```
File already exists, do you want to overwrite [y/n]: y
FIGW>
```

2. Copy the configuration files (*.cfg), the shadov.txt file, which is an encrypted file that contains the authentication keys for the IPsec tunnels, and the default-config-file.txt file from the VM to intflash within the NOS. For more information, see [Run a VM command from Network Operating System \(NOS\) CLI](#) on page 1122 and [Copy VM Files](#) on page 1123.

```
Switch:1(config)#mkdir figw
Switch:1(config)#virtual-service figwOld exec-command "ls /home/rwa/configs/"
config.cfg
figw.cfg
figw_cli.log
new.cfg
shadov.txt
Switch:1(config)#virtual-service copy-file figwOld:/home/rwa/configs/config.cfg /
intflash/figw/config.cfg
Switch:1(config)#virtual-service copy-file figwOld:/home/rwa/configs/new.cfg /intflash/
figw/new.cfg
Switch:1(config)#virtual-service copy-file figwOld:/home/rwa/configs/figw.cfg /
intflash/figw/figw.cfg
Switch:1(config)#virtual-service copy-file figwOld:/home/rwa/default-config-file.txt /
intflash/figw/default-config-file.txt
Switch:1(config)#virtual-service copy-file figwOld:/home/rwa/configs/shadov.txt /
intflash/figw/shadov.txt
```

3. Verify the file copy:

```
Switch:1#ls figw/
Listing Directory /intflash/figw/:
drwxr-xr-x 2 0 0 4096 Jun 17 13:46 ./
drwxr-xr-x 31 0 0 4096 Jun 17 13:43 ../
-rw-r--r-- 1 0 0 851 Jun 17 13:44 config.cfg
-rw-r--r-- 1 0 0 8 Jun 17 13:46 default-config-file.txt
-rw-r--r-- 1 0 0 0 Jun 17 13:45 figw.cfg
-rw-r--r-- 1 0 0 851 Jun 17 13:45 new.cfg
-rw-r--r-- 1 0 0 32 Jun 17 13:45 shadov.txt
```

4. Disable the virtual service:

```
no virtual-service WORD<1-128> enable
```

5. Uninstall the virtual service:

```
virtual-service WORD<1-128> uninstall
```



Note

When you uninstall the original virtual service, the system removes the complete virtual service configuration from the configuration file.

6. Install the virtual service package using the new image:

```
virtual-service WORD<1-128> install package WORD<1-512>
```

7. Reconfigure the virtual service. For more information, see [Configure a Virtual Service](#) on page 1113.

8. Copy the files you saved from the old VM to the same folder path in the new VM:

```
Switch:1(config)#virtual-service copy-file /intflash/figw/config.cfg figwNew:/home/rwa/
configs/config.cfg
Switch:1(config)#virtual-service copy-file /intflash/figw/figw.cfg figwNew:/home/rwa/
configs/figw.cfg
Switch:1(config)#virtual-service copy-file /intflash/figw/new.cfg figwNew:/home/rwa/
configs/new.cfg
Switch:1(config)#virtual-service copy-file /intflash/figw/shadov.txt figwNew:/home/rwa/
configs/shadov.txt
Switch:1(config)#virtual-service copy-file /intflash/figw/default-config-file.txt
figwNew:/home/rwa/default-config-file.txt
```

9. Verify the file copy:

```
Switch:1(config)#virtual-service figwNew exec-command "ls /home/rwa/configs"
config.cfg
figw.cfg
figw_cli.log
new.cfg
shadov.txt
```

10. Reboot the Fabric IPsec Gateway VM. For more information, see [Reboot Fabric IPsec Gateway VM](#) on page 1165.



Tip

As an alternative, you can disable and reenab the Fabric IPsec Gateway virtual service.

11. Verify the running configuration of the new VM matches the configuration of the old VM:

```
Switch:1(config)#virtual-service figwNew figw-cli "show running-config"
set global ipsec-tunnel-src-vlan 30
set global ipsec-tunnel-src-ip 30.30.30.2/24
set global lan-intf-vlan 100
set global lan-intf-ip 100.100.100.2/24
set global lan-intf-gw-ip 100.100.100.102
set global fe-tunnel-src-ip 102.102.102.102
set global wan-intf-gw-ip 30.30.30.102
set global mtu 1950
set global services sshd enable
set ipsec 104 auth-key *****
set ipsec 104 responder-only true
set ipsec 104 fe-tunnel-dest-ip 104.104.104.104
set ipsec 104 fragment-before-encrypt enable
set ipsec 104 admin-state enable
set ipsec 105 auth-key *****
set ipsec 105 responder-only true
set ipsec 105 fe-tunnel-dest-ip 105.105.105.105
set ipsec 105 fragment-before-encrypt enable
set ipsec 105 admin-state enable
set ipsec 107 auth-key *****
set ipsec 107 responder-only true
set ipsec 107 fe-tunnel-dest-ip 192.168.22.107
set ipsec 107 admin-state enable
```

12. Remove the original image from the /var/lib/insight/packages/ directory on the switch:

```
remove WORD<1-255>
```

Configure FTP Connection to an IP Address

Fabric IPsec Gateway Virtual Machine (VM) provides a File Transfer Protocol (FTP) CLI to copy the configuration files to the VM.

About This Task

Perform this procedure to configure an FTP connection to a specific IP Address.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure FTP connection:

```
ftp {A.B.C.D}
```

Example

Configuring FTP connection to 192.0.2.50:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> ftp 192.0.2.50
```

Variable Definitions

The following table defines the variable for **ftp** command.

Variable	Value
{A.B.C.D}	Specifies the IP Address to establish the FTP connection with.

Display the Default Directory on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display content in the default directory on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display the configured directory:

```
ls
```

Example

Displaying the configured directory on the VM.

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> ls
coupled.cfg
```

Load Configuration File to Fabric IPsec Gateway VM

About This Task

Perform this procedure to load a configuration file to the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Load a specific configuration file to the VM :

```
load WORD <1-255>
```

Example

Loading a configuration file to the VM.

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> load coupled.cfg
```

Variable Definitions

The following table defines the variable for **load** command.

Variable	Value
<i>WORD</i> <1–255>	Specifies the configuration file name.

Ping an IP Address on Fabric IPsec Gateway VM

About This Task

Perform this procedure to ping an IP Address on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Ping an IP Address:

```
ping {A.B.C.D}
```

Example

Pinging an IP Address on the VM.

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> ping 192.0.2.35
```

Variable Definitions

The following table defines parameters for the **ping** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the IP address.

Configure Global Parameters on Fabric IPsec Gateway VM

About This Task

Perform this procedure to configure IPsec source IP address, Local Area Network (LAN) interface IP and gateway IP address, maximum transmission unit (MTU) value, and so on globally, on the Fabric IPsec Gateway Virtual Machine (VM).



Note

You must perform this procedure only after the VM boots up.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure IPsec source IP address for a Fabric Extend (FE) tunnel for IPsec in decoupled mode:

```
set global ipsec-tunnel-src-ip {A.B.C.D/X}
```

3. Assign VLAN ID to the configured IPsec source IP address:

```
set global ipsec-tunnel-src-vlan <2-4059>
```

4. Configure the LAN interface IP address on the first Ethernet interface (eth0) of Fabric IPsec Gateway VM:

```
set global lan-intf-ip {A.B.C.D/X}
```

5. Assign VLAN ID to the configured LAN interface IP address:

```
set global lan-intf-vlan <2-4059>
```

6. Configure the LAN interface gateway IP address on the VOSS switch:

```
set global lan-intf-gw-ip {A.B.C.D}
```

7. Configure the logical interface gateway IP address, to add routes for FE tunnels that need Fragmentation:

```
set global fe-tunnel-gw-ip {A.B.C.D}
```

8. Configure the logical interface source IP address for the FE tunnel:

```
set global fe-tunnel-src-ip {A.B.C.D}
```



Note

The logical interface source IP address must be same as the source IP address configured on the VOSS switch.

9. Configure the global MTU value:

```
set global mtu <mtu-value>
```



Note

- The switch applies the global MTU value, if you do not configure MTU during the IPsec tunnel configuration.
- If an IPsec tunnel is not using the fragmentation and reassembly capabilities, the default MTU value is 1950.

10. Configure the Wide Area Network (WAN) interface gateway IP address, which is the next hop for IPsec tunnels.

```
set global wan-intf-gw-ip {A.B.C.D}
```

11. Configure the virtual reassembly interface IP address:

```
set global virtual-reassembly-intf-ip {A.B.C.D/X}
```



Note

You must configure the virtual reassembly interface IP address to use the fragmentation and reassembly service.

12. Assign VLAN ID to the configured virtual reassembly interface IP address:

```
set global virtual-reassembly-intf-vlan <2-4059>
```

13. Disable IPsec on all configured tunnels:

```
set global ipsec-disable
```

14. Set IPsec log level:

```
set global ipsec-log-level <-1-5>
```

Example

Configuring global parameters on Fabric IPsec Gateway VM to configure an IPsec tunnel between two switches:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> set global ipsec-tunnel-src-ip 192.0.2.10/24
FIGW> set global ipsec-tunnel-src-vlan 101
FIGW> set global lan-intf-ip 192.0.2.20/24
FIGW> set global lan-intf-vlan 30
FIGW> set global lan-intf-gw-ip 192.0.2.30
FIGW> set global fe-tunnel-src-ip 192.0.2.40
FIGW> set global wan-intf-gw-ip 192.0.2.50
FIGW> set global mtu 1950
```

Variable Definitions

The following table defines parameters for the **set global** command.

Variable	Value
<i>ipsec-tunnel-src-ip</i> {A.B.C.D/X}	Specifies the source IP address and subnet mask for IPsec tunnel.
<i>ipsec-tunnel-src-vlan</i> <2-4059>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.
<i>lan-intf-ip</i> {A.B.C.D/X}	Specifies the IP address and subnet mask for Local Area Network (LAN) interface.
<i>lan-intf-vlan</i> <2-4059>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.
<i>lan-intf-gw-ip</i> {A.B.C.D}	Specifies the gateway IP address for LAN interface.
<i>fe-tunnel-gw-ip</i> {A.B.C.D}	Specifies the gateway IP address for Fabric Extend (FE) tunnel.
<i>fe-tunnel-src-ip</i> {A.B.C.D}	Specifies the source IP address for FE tunnel.
<i>mtu</i> <750-9000>	Specifies the Maximum Transmission Unit (MTU) value. Note: If an IPsec tunnel is not using the fragmentation and reassembly capabilities, the default MTU value is 1950.
<i>wan-intf-gw-ip</i> {A.B.C.D}	Specifies the Wide Area Network (WAN) interface gateway IP address.
<i>virtual-reassembly-intf-ip</i> {A.B.C.D/X}	Specifies the virtual-reassembly interface IP address and subnet mask on the Fabric IPsec Gateway (VM). Note: You must configure the virtual reassembly interface IP address to use the fragmentation and reassembly service.
<i>virtual-reassembly-intf-vlan</i> <2-4059>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.

Variable	Value
<i>ipsec-disable</i>	Disables IPsec operationally on all tunnels in the Fabric IPsec Gateway VM.
<i>ipsec-log-level</i> <-1-5>	Specifies the IPsec log levels on Fabric IPsec Gateway VM. Following are the three levels: <ul style="list-style-type: none"> • -1: Absolutely Silent • 0-4: Log levels • 5: Clear Logs

Configure IPsec Tunnels on Fabric IPsec Gateway VM

About This Task

Perform this procedure to configure IPsec tunnels on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure the Maximum Transmission Unit (MTU) value for the specific IPsec tunnel:

```
set ipsec <1-255> mtu <1300 - 9000>
```



Note

The MTU range <1300-9000> is applicable for FE tunnels with IPsec and fragmentation and reassembly capabilities.

3. Configure the ESP cipher suite for the IPsec tunnel:

```
set ipsec <1-255> esp <aes128gcm16-sha256 | aes256-sha256 |  
aes256gcm16-sha256>
```

4. Configure the authentication key for specific IPsec tunnel:

```
set ipsec <1-255> auth-key WORD <1-32>
```



Note

Do not use special characters ?, \, &, <, >, #.

5. Configure VXLAN destination IP address for IPsec tunnel:

```
set ipsec <1-255> fe-tunnel-dest-ip {A.B.C.D}
```



Note

The VXLAN destination IP address for IPsec tunnel must be the same as the VXLAN destination IP address for FE tunnel.

6. Configure the IPsec destination IP address for the specific tunnel deployed in decoupled mode:

```
set ipsec <1-255> ipsec-dest-ip {A.B.C.D}
```

7. Configure a name for the IPsec tunnel:

```
set ipsec <1-255> tunnel-name WORD <1-64>
```

8. Identify if the specific tunnel is a responder or initiator in Network Address Translation (NAT) cases:

```
set ipsec <1-255> responder-only <true | False>
```

9. Enable the IPsec on a specific tunnel:

```
set ipsec <1-255> admin-state enable
```

Example

Configure parameters for IPsec tunnel on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> set ipsec 1 ipsec-dest-ip 192.0.2.5
FIGW> set ipsec 1 mtu 1950
FIGW> set ipsec 1 auth-key abcd
FIGW> set ipsec 1 tunnel-name Tunnel-to-BEB2
FIGW> set ipsec 1 fe-tunnel-dest-ip 192.0.2.15
FIGW> set ipsec 1 esp aes256gcm16-sha256
FIGW> set ipsec 1 admin-state enable
```

Variable Definitions

The following table defines parameters for the **set ipsec** command.

Variable	Value
<1-255>	Specifies the unique ID for the IPsec tunnel.
<i>admin-state</i> <enable disable>	Enables or disables IPsec on the specific IPsec tunnel.
<i>auth-key</i> WORD <1-32>	Specifies the pre-shared authentication key. Note: Do not use special characters ?, \, &, <, >, #.

Variable	Value
<code>encryption-key-length <128 256></code>	Specifies the encryption key length for the IPsec tunnel. The default encryption key length is 128. As a best practice, use the newer <code>esp</code> parameter instead; the <code>encryption-key-length</code> parameter remains for backward compatibility.
<code>esp <aes128gcm16-sha256 aes256-sha256 aes256gcm16-sha256></code>	Specifies the ESP cipher suites for the IPsec tunnel. The default is aes128gcm16-sha256. aes256-sha256 is not supported in the current release.
<code>fe-tunnel-dest-ip {A.B.C.D}</code>	Specifies the destination IP address for Fabric Extend (FE) tunnel.
<code>ipsec-dest-ip {A.B.C.D}</code>	Specifies the destination IP address for IPsec tunnel.
<code>mtu <1300-9000></code>	Specifies the Maximum Transmission Unit (MTU) value for the FE tunnel with both IPsec and fragmentation and assembly capabilities.
<code>responder-only <true false></code>	Specifies if the IPsec session in the FE tunnel is in responder only mode or initiator mode. When in responder mode the FE tunnel only responds to the incoming request and does not initiate the IPsec connection. By default, both sides of IPsec connection are initiators in the FE tunnel. Configure the IPsec tunnel to be in responder only mode when there is Network Address Translation (NAT) between the IPsec connection. For more information about NAT, see IPsec NAT-T on page 1108.
<code>tunnel-name WORD <1-64></code>	Specifies a name for the IPsec tunnel.
<code>egress-shaping-rate <1-1000></code>	Specifies the egress shaper rate for the IPsec tunnel.

Configure IPsec Compression

Before You Begin

Ensure IPsec fragmentation before encryption is disabled.

About This Task

Perform this procedure to enable IPsec compression on Fabric IPsec Gateway Virtual Machine (VM).



Note

By default, IPsec compression is disabled. You must enable IPsec compression on both ends of the adjacency.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Enable IPsec compression:

```
set ipsec compression
```

Enable Fragmentation Before Encryption on Fabric IPsec Gateway VM

Perform this procedure to fragment packets larger than the IPsec tunnel maximum transmission unit (MTU) before the packets are sent for encryption.

Before You Begin

- Ensure IPsec is disabled on the tunnel. The administrative state must be disabled before you can enable or disable fragmentation before encryption.
- Configure the IPsec destination IP address or enable responder mode.

About This Task

By default, fragmentation before encryption is disabled.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Enable fragmentation before IPsec encryption:

```
set ipsec <1-255> fragment-before-encrypt enable
```

Configure Public Key Infrastructure for IPsec Tunnels

Before You Begin

- Configure the Fabric Extend tunnels between the branch and hub switches.
- Configure digital certificates on the switch using Fabric IPsec Gateway virtual machine.

About This Task

Support for IPsec authentication and encryption of Fabric Extend tunnels is provided using Fabric IPsec Gateway. You can use a digital certificate to authenticate IPsec for Fabric Extend

The default IPsec authentication method for Fabric Extend tunnels is a pre-shared key. If you configure the authentication method to RSA signature, the tunnels use the installed digital certificate.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure the authentication type as RSA signature:

```
set ipsec <1-255> auth-method rsasig
```

Variable Definitions

The following table defines parameters for the **set ipsec** command.

Variable	Value
<1-255>	Specifies the tunnel ID.
<subject-label>	Specifies the subject identity.
cert-subject-nameWORD<1-45>	Specifies the digital certificate subject name to be used as the identity certificate. If a subject name is not specified, the default certificate subject name is Global.

Configure Egress Shaping Rate for IPsec Tunnels on Fabric IPsec Gateway VM

Before You Begin

Before you can configure the egress shaping rate for the IPsec tunnel on the VM, you must first disable the IPsec tunnel.

About This Task

Perform this procedure to configure the egress shaping rate for IPsec tunnels on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Disable the IPsec tunnel:

```
delete ipsec <1-255> admin-state enable
```

3. Configure the egress shaping rate for the IPsec tunnel:

```
set ipsec <1-255> egress-shaping-rate <1-1000>
```

4. Enable the IPsec tunnel:

```
set ipsec <1-255> admin-state-enable
```

Example

Configuring egress-shaping-rate for the IPsec tunnel on the Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> delete ipsec 1 admin-state enable
FIGW> set ipsec 1 egress-shaping-rate 200
FIGW> set ipsec 1 admin-state enable
```

Variable Definitions

The following table defines parameters for the **set ipsec** command.

Variable	Value
<1-255>	Specifies the unique ID for the IPsec tunnel.
<i>admin-state</i> <enable disable>	Enables or disables IPsec on the specific IPsec tunnel.
<i>egress-shaping-rate</i> <1-1000>	Specifies the egress shaping rate for the IPsec tunnel.

Configure the Subject Identity on Fabric IPsec Gateway VM

About This Task

Use this procedure to configure the subject parameters to identify the device.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure the distinguished name of the device:

```
set certificate subject <subject-label> DN <name>
```

3. (Optional) Configure the name of the generated key-pair:

```
set certificate subject <subject-label> key-label <key-label>
```

Example

```
Switch:1>enable
Switch:1#virtual-service FIGW console
FIGW>set certificate subject ExtremeLab DN "CN=subca5, OU=Test, O=Extreme, L=Town,
ST=State, E=email@extremenetworks.com"
FIGW>set certificate subject ExtremeLab key-label key1
```

Variable Definitions

The following table defines parameters for the **set certificate subject** command.

Variable	Value
<i>DN <name></i>	Specifies the distinguished name. You can create a comma-separated list.
<i>key-label <key-label></i>	Specifies the key name of the generated key pair. This parameter is optional. If you do not configure one, the switch generates one the same as the subject-label.
<i><subject-label></i>	Specifies the subject identity. You cannot use the following special characters: <ul style="list-style-type: none"> • question mark (?) • backslash (\) • ampersand (&) • less than (<) • greater than (>) • pound (#)

Generate the Key Pair on Fabric IPsec Gateway VM

Before You Begin

- Configure an EJBCA CA server.
- Configure a route from Fabric IPsec Gateway to the EJBCA CA server.

About This Task

Use the following procedure to generate the private and public key pair. By default, VOSS generates a 2,048 RSA key when the system starts. You can use this procedure to generate a new key.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Generate the key:

```
certificate generate key <type> <size> <key-label>
```

Example

```
Switch:1>enable
Switch:1#virtual-service FIGW console
FIGW>certificate generate key rsa 2048 key_rsa
fingerprint: 09ac0c64b9bf3ad04dc67f20942c674e
```

Variable Definitions

The following table defines parameters for the **certificate generate key** command.

Variable	Value
<i>key-label</i>	Specifies the key name of the generated key pair.
<i>size</i>	Specifies the size of key-pair to be generated. The switch supports 2048.
<i>type</i>	Specifies the type of cryptography algorithm used to generate the key-pair. The switch uses only rsa as the cryptography algorithm type.

Configure a Trustpoint CA on Fabric IPsec Gateway VM

About This Task

Use this procedure to configure the certificate authority (CA) to use Simple Certificate Enrollment Protocol (SCEP) with a CA server for online certificate provisioning.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure the trusted CA URL:

```
set certificate ca-trustpoint <ca-label> ca-url <ca-url>
```

3. Configure the common name of the CA:

```
set certificate ca-trustpoint <ca-label> caname <caname>
```

4. Configure the HTTP request type to support the type of CA:

```
set certificate ca-trustpoint <ca-label> get-method <post | get>
```

5. Configure the appropriate action:

- Configure the trustpoint, authenticate the trustpoint CA by getting the certificate of the CA, and store the CA certificate locally:

```
certificate ca <ca-trustpoint> caAuth
```

- Generate the certificate enrollment request, get the digital certificate, and store it locally, associating with the trustpoint CA:

```
certificate ca <ca-trustpoint> enroll <subject-label>
```

- Get the Certificate Revocation List (CRL) from the CDP and store into a file:

```
certificate get crl-from <A.B.C.D> <user> <file-path> <cacert-  
filename>
```

- Get the CA certificate obtained from the trustpoint CA:

```
certificate get cacert-from <A.B.C.D> <user> <file-path>
```

- Get the subject certificate obtained from the trustpoint CA:

```
certificate get signedcert-from <A.B.C.D> <user> <file-path>  
<subject-label>
```

- Release the locally stored certificate associated with the trustpoint CA after revocation:

```
certificate ca <ca-trustpoint> remove <subject-label>
```

- Remove all certificates from the CA trustpoint:



Note

You can clean the CA trustpoint only if the subject-label is not configured on an IPsec tunnel.

```
certificate ca <ca-trustpoint> clean
```

Example

```
Switch:1>enable
Switch:1#virtual-service FIGW console
FIGW>set certificate ca-trustpoint caExtremeEJBCA ca-url http://192.0.2.9:8080/ejbcas/publicweb/apply/scep/test/pkiclient.exe
FIGW>set certificate ca-trustpoint caExtremeEJBCA caname subca5
FIGW>set certificate ca-trustpoint caExtremeEJBCA get-method post
```

Variable Definitions

The following table defines parameters for the **set certificate ca-trustpoint** command.

Variable	Value
<code><ca-label></code>	Specifies the name of the certificate authority (CA). The name can use alphanumeric characters and is case-sensitive. The maximum length is 45 characters.
<code>ca-url <ca-url></code>	Specifies the trusted CA URL.
<code>caname <caname></code>	Specifies the name of the owner of the device or user.
<code>get-method <post get></code>	Specifies the HTTP request style. You can use post for EJBCA or get for Win2012 CA. The default value is post.

The following table defines parameters for the **certificate ca** command.

Variable	Value
<code><ca-trustpoint></code>	Specifies the name of the certificate authority. The name can be alphanumeric and is case-sensitive. The maximum length is 45 characters.
<code><subject-label></code>	Specifies the subject identity.

The following table defines parameters for the **certificate get** command.

Variable	Value
<code>cacert-from <A.B.C.D> <user> <file-path></code>	Specifies where to obtain the CA certificate. Specify the IP address, user name, and remote file path.
<code>crl-from <A.B.C.D> <user> <file-path> <cacert-filename></code>	Specifies where to obtain the Certificate Revocation List. Specify the IP address, user name, remote file path, and the CA certificate file to verify the CRL.
<code>signedcert-from <A.B.C.D> <user> <file-path> <subject-label></code>	Specifies where to obtain the subject certificate. Specify the IP address, user name, remote file path, and subject label.

Generate the Certificate Signing Request on Fabric IPsec Gateway VM

About This Task

Use this procedure to generate a certificate signing request (CSR) and store it into a file. This CSR is required to obtain the offline subject certificate.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Generate the CSR:

```
certificate generate csr <subject-label>
```

3. Configure where to send the CSR for signing:

```
certificate send-csr-to <A.B.C.D> <user> <remote-path> <subject-label>
```

Variable Definitions

The following table defines parameters for the **certificate generate csr** command.

Variable	Value
<i><subject-label></i>	Specifies the subject identity.

The following table defines parameters for the **certificate send-csr-to** command.

Variable	Value
<i><A.B.C.D></i>	Specifies the IP address for the certificate authority.
<i><remote-path></i>	Specifies the file path on the certificate authority.
<i><subject-label></i>	Specifies the subject identity.
<i><user></i>	Specifies the user name for the certificate authority.

Remove Keys and Certificates on Fabric IPsec Gateway VM

Before You Begin

You can remove subject certificates from the certificate authority (CA) trustpoint only if the subject-label is not configured on an IPsec tunnel.

About This Task

Use this procedure to remove keys or certificates from the certificate store.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Remove a key:

```
certificate remove key <key-label>
```

3. Remove a specific certificate from the store:

```
certificate remove offline-cacert <filename>
```

4. Remove a Certificate Revocation List (CRL) certificate from the store:

```
certificate remove offline-crl <filename>
```

5. Remove signed certificates for a specific subject label:

```
certificate remove offline-subject-certs <subject-label>
```

6. Remove a specific identity certificate from the CA trustpoint:

```
certificate ca <ca-trustpoint> remove <subject-label>
```

7. Remove all certificates from the CA trustPoint:

```
certificate ca <ca-trustpoint> clean
```

Variable Definitions

The following table defines parameters for the **certificate remove** command.

Variable	Value
<i>key</i> <key-label>	Specifies the key name to remove.
<i>offline-cacert</i> <filename>	Specifies the certificate filename to remove.
<i>offline-crl</i> <filename>	Specifies the Certificate Revocation List (CRL) certificate filename to remove.
<i>offline-subject-certs</i> <subject-label>	Specifies the subject label for which to remove signed certificates.

The following table defines parameters for the **certificate ca** command.

Variable	Value
<ca-trustpoint>	Specifies the name of the certificate authority. The name can be alphanumeric and is case-sensitive. The maximum length is 45 characters.
<subject-label>	Specifies the subject identity.

View the Certificate Details on Fabric IPsec Gateway VM

About This Task

Use this procedure for the following tasks:

- Display the digital certificate for a certificate type or list all the certificate details from the local store.
- Display the certificate authority (CA) details for a trustpoint CA name or list all the CA details from the local store if the CA name is not specified.
- Display the configured key details for a key name.
- Display the configured subject details.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display all digital certificates:

```
show certificates all
```

3. Display the CA details:

```
show certificates cacert [<ca-label>]
```

4. Display Certificate Revocation List (CRL) certificate details:

```
show certificates crl [<ca-label>]
```

5. Display the certificate signing request (CSR) details:

```
show certificates csr [<ca-label>]
```

6. Display the name and public key of all the key-pairs:

```
show certificates keys
```

7. Display the details of signed certificates:

```
show certificates signed [<ca-label>]
```

Example

```
Switch:1>enable
Switch:1#virtual-service FIGW console
FIGW>show certificates keys key_rsa
Key Label:      key_rsa
private key with:
pubkey:         RSA 2048 bits
keyid:          ef:4c:1d:a7:cc:84:6f:87:da:e4:de:99:07:3d:96:fc:9a:d1:c9:f4
subjkey:        cb:d1:67:a0:da:9c:05:ce:c0:0d:a3:5c:1b:ba:ce:3f:ff:af:8f:77
```

Variable Definitions

The following table defines parameters for the **show certificates** command.

Variable	Value
ca <ca-label>	Specifies the name of the certificate authority (CA). If you do not specify the name, the command displays the details of all configured CAs.

View the Certificate Configuration on Fabric IPsec Gateway VM

About This Task

Use this procedure to view the certificate configuration for the VM.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display all configured entries:

```
show running-config
```

3. Display the CA trustpoint configuration:

```
show certificate-config ca-trustpoint [<ca-label>]
```

4. Display the subject-related configuration:

```
show certificate-config subject [<subject-label>]
```

Examples

```
Switch:1>enable
Switch:1#virtual-service FIGW console
FIGW>show certificate-config ca-trustpoint
certificate {
  ca-trustpoint {
    ca-label a;
    caname subCaVpn;
    ca-url http://10.2.38.35:8080/ejbca/publicweb/apply/scep/test/pkiclient.exe;
    get-method post;
  }
}

Switch:1>enable
Switch:1#virtual-service FIGW console
FIGW>show certificate-config subject
certificate {
  subject {
    subject-label fig;
    DN CN=FIGW;
    key-label gigi;
  }
  subject {
    subject-label figv;
    DN CN=figvpn;
  }
}
```

Configure Logical Interface Tunnel on Fabric IPsec Gateway VM

About This Task

Perform this procedure to configure a Fabric Extend (FE) tunnel with only fragmentation and reassembly capabilities, on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Configure the logical interface destination IP address for the specific tunnel:

```
set logical-intf-tunnel <1-255> fe-tunnel-dest-ip {A.B.C.D}
```

3. Configure the Maximum Transmission Unit (MTU) value for the specific tunnel:

```
set logical-intf-tunnel <1-255> mtu <750-9000>
```



Note

The MTU range <750-9000> is applicable for FE tunnels with only fragmentation and reassembly capabilities.

4. Configure tunnel name:

```
set logical-intf-tunnel <1-255> tunnel-name WORD <1-64>
```

5. Configure the egress shaping rate for a specific tunnel:

```
set logical-intf-tunnel <1-255> egress-shaping-rate <1-1000>
```

Example

Configuring logical interface tunnel on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> set logical-intf-tunnel 2 fe-tunnel-dest-ip 192.0.2.50
FIGW> set logical-intf-tunnel 2 mtu 1300
FIGW> set logical-intf-tunnel 2 egress-shaping-rate 5
FIGW> set logical-intf-tunnel 2 tunnel-name Tunnel-to-BEB2
```

Variable Definitions

The following table defines parameters for the **set logical-intf-tunnel** command.

Variable	Value
<1-255>	Specifies the unique ID for the logical interface tunnel.
<i>fe-tunnel-dest-ip</i> {A.B.C.D}	Specifies the FE tunnel destination IP address for the logical interface.
<i>mtu</i> <750-9000>	Specifies the Maximum Transmission Unit (MTU) value for the FE tunnel with only fragmentation and assembly capabilities.
<i>tunnel-name</i> WORD <1-64>	Specifies a name for the the logical interface tunnel.
<i>egress-shaping-rate</i> <1-1000>	Specifies the egress shaping rate for the logical interface tunnel.

Save Running Configuration to a File

About This Task

Perform this procedure to save the current configuration on Fabric IPsec Gateway Virtual Machine (VM) to a specific file.

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Save configuration to the default configuration file:

```
save config [-y]
```

3. Save configuration to a specific file in the Fabric IPsec Gateway VM.

```
save config file WORD <1-255> [-y]
```

Example

Save the Fabric IPsec Gateway configuration to file "test":

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW>save config file test.txt
File already exists, do you want to overwrite [y/n]: y
```

Save the Fabric IPsec Gateway configuration to file "test", forcing the switch to overwrite the file without confirmation:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> save config file test.txt -y
```

Variable Definitions

The following table defines parameters for the **save config** command.

Variable	Value
<i>file</i> <i>WORD</i> <1-255>	Specifies the name of file to save the configuration of the Fabric IPsec Gateway VM.
<i>-y</i>	Forces the switch to overwrite the configuration file without confirmation.

Remove Configuration File from Fabric IPsec Gateway VM

About This Task

Perform this procedure to remove a specific configuration file from Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Remove the configuration file:

```
remove WORD <1-255>
```

Example

Remove configuration file "test" from Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> remove test
```

Variable Definitions

The following table defines parameters for the **remove** command.

Variable	Value
<i>WORD</i> <1-255>	Specifies the configuration file name that the system removes from Fabric IPsec Gateway VM.

Delete Global Configuration on Fabric IPsec Gateway VM

About This Task

Perform this procedure to delete the global parameters that you configure on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Delete configuration of specific global parameters:

```
delete global <fe-tunnel-gw-ip | fe-tunnel-src-ip | ipsec-disable |  
ipsec-tunnel-src-ip | ipsec-tunnel-src-vlan | lan-intf-gw-ip | lan-  
intf-ip | lan-intf-vlan | mtu | virtual-reassembly-intf-ip | virtual-  
reassembly-intf-vlan | wan-intf-gw-ip>
```

Example

Deleting the global Maximum Transmission Unit (MTU) configuration on Fabric IPsec Gateway VM:

```
Switch:1> enable  
Switch:1# virtual-service figw console  
Connected to domain figw  
Escape character is ^Y  
  
<cr>  
FIGW> delete global mtu
```

Variable Definitions

The following table defines parameters for the **delete global** command.

Variable	Value
<i>fe-tunnel-gw-ip</i>	Deletes the global gateway IP address for Fabric Extend (FE) tunnel.
<i>fe-tunnel-src-ip</i>	Deletes the global source IP address for FE tunnel.

Variable	Value
<i>ipsec-disable</i>	Deletes the global IPsec configuration.
<i>ipsec-tunnel-src-ip</i>	Deletes the global source IP address and subnet mask for IPsec tunnel.
<i>ipsec-tunnel-src-vlan</i>	Deletes the global source VLAN configuration for IPsec tunnel.
<i>lan-intf-gw-ip</i>	Deletes the global gateway IP address on the Local Area Network (LAN) interface.
<i>lan-intf-ip</i>	Deletes the global IP address and subnet mask on LAN interface.
<i>lan-intf-vlan</i>	Deletes the global VLAN configuration on LAN interface.
<i>mtu</i>	Resets the Maximum Transmission Unit (MTU) value to its default, that is 1950 bytes.
<i>virtual-reassembly-intf-ip</i>	Deletes the global virtual-reassembly interface IP address and subnet mask.
<i>virtual-reassembly-intf-vlan</i>	Deletes the global virtual-reassembly interface VLAN configuration.
<i>wan-intf-gw-ip</i>	Deletes the global gateway IP address on the Wide Area Network (WAN) interface.

Delete IPsec Tunnel Configuration on Fabric IPsec Gateway VM

Before You Begin

You must disable the IPsec administrative state on the tunnel before you can remove IPsec configuration.

About This Task

Perform this procedure to delete the configuration of a specific IPsec tunnel on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Delete the configuration of a specific tunnel:

```
delete ipsec <1-255> <admin-state enable | auth-key | encryption-key-length | fe-tunnel-dest-ip | fragment-before-encrypt enable | ipsec-dest-ip | mtu | responder-only | tunnel-name | egress-shaping-rate>
```


Example

Delete configuration on IPsec tunnel ID 2:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW>delete ipsec 2 admin-state enable
FIGW>delete ipsec 2 auth-key
FIGW>delete ipsec 2 tunnel-name
FIGW>delete ipsec 2 fragment-before-encrypt enable
```

Variable Definitions

The following table defines parameters for the **delete ipsec** command.

Variable	Value
<1-255>	Specifies the unique ID of the configured IPsec tunnel.
<i>admin-state enable</i>	Disables the IPsec status on the specific IPsec tunnel.
<i>auth-key</i>	Deletes the authentication key that you configure on the specific IPsec tunnel.
<i>encryption-key-length</i>	Resets the encryption key length for the specific IPsec tunnel to its default value, that is 128 bit.
<i>fe-tunnel-dest-ip</i>	Deletes the destination IP address that you configure on the Fabric Extend (FE) tunnel.
<i>fragment-before-encrypt enable</i>	Disables the fragmentation of packets before IPsec encryption on the tunnel. By default, fragmentation before encryption is disabled.
<i>ipsec-dest-ip</i>	Deletes the destination IP address that you configure on the IPsec tunnel.
<i>mtu</i>	Resets the Maximum Transmission Unit (MTU) value for the specific IPsec tunnel to the MTU value configured globally.
<i>responder-only</i>	Deletes the mode that you configure for the IPsec session in FE tunnel.
<i>tunnel-name</i>	Deletes the name that you configure for the IPsec tunnel.
<i>egress-shaping-rate</i>	Deletes the egress shaping rate for the IPsec tunnel.

Delete Logical Interface Tunnel Configuration on Fabric IPsec Gateway VM

About This Task

Perform this procedure to delete configuration of a specific logical interface tunnel on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Delete configuration of specific logical interface tunnel:

```
delete logical-intf-tunnel <1-255> < fe-tunnel-dest-ip | mtu | egress-  
shaping-rate>
```

Example

Deleting the destination IP address for Fabric Extend (FE) tunnel configured on the logical interface tunnel with ID 3.

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> delete logical-intf-tunnel 3 fe-tunnel-dest-ip
```

Variable Definitions

The following table defines parameters for the **delete logical-intf-tunnel** command.

Variable	Value
<i><1-255></i>	Specifies the unique ID of the logical interface tunnel.
<i>fe-tunnel-dest-ip</i>	Deletes the destination IP address that you configure on the logical interface tunnel.
<i>mtu</i>	Resets the Maximum Transmission Unit (MTU) value for the specific logical interface tunnel to the MTU value configured globally.
<i>egress-shaping-rate</i>	Deletes the egress shaping rate on the logical interface tunnel.

Display Data in a File on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the data in a specific file on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display data in a file:

```
more WORD <1-255>
```

Example

Display the data from coupled.cfg file:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> more coupled.cfg
set global ipsec-tunnel-src-vlan 125
set global ipsec-tunnel-src-ip 192.0.2.10/24
set global lan-intf-vlan 30
set global lan-intf-ip 192.0.2.20/24
set global lan-intf-gw-ip 192.0.2.25
set global fe-tunnel-src-ip 192.0.2.45
set global wan-intf-gw-ip 192.0.2.11
set global mtu 1950
set ipsec 1 auth-key *****
set ipsec 1 fe-tunnel-dest-ip 192.0.2.50
set ipsec 1 encryption-key-length 128
set ipsec 1 admin-state enable
```

Reboot Fabric IPsec Gateway VM

About This Task

Perform this procedure to reboot the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Reboot the VM:

```
reboot
```

Example

Rebooting Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> reboot
```

Reset Current Configuration on Fabric IPsec Gateway VM

About This Task

Perform this procedure to reset the current configuration on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Reset current configuration:

```
reset-config
```



Note

Reboot the Fabric IPsec Gateway VM after you reset the configuration.

Example

Resetting current configuration on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> reset-config
```

Traceroute to an IP address on Fabric IPsec Gateway VM

About This Task

Perform this procedure to traceroute to an IP address on Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Traceroute to an IP address:

```
traceroute {A.B.C.D}
```

Example

Traceroute to IP address.

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> traceroute 192.0.2.100
```

Variable Definitions

The following table defines parameters for the **traceroute** command.

Variable	Value
{A.B.C.D}	Specifies the IP address to initiate traceroute to.

Display the Default Configuration File on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the default configuration file on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display default configuration file:

```
show default-config-file
```

Example

Displaying default configuration file on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> show default-config-file
coupled.cfg
```

Display IPsec Logs on Fabric IPsec Gateway

About This Task

Perform this procedure to display IPsec session logs on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display IPsec session logs:

```
show ipsec-logs
```

Example

Displaying IPsec session logs on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> show ipsec-logs
<<Month dd>> <<hh:mm:ss>> 15[IKE] <ipsec0-192.0.2.10|29> sending DPD request
<<Month dd>> <<hh:mm:ss>> 15[ENC] <ipsec0-192.0.2.10|29> generating INFORMATIONAL request
11832 [ ]
<<Month dd>> <<hh:mm:ss>> 15[NET] <ipsec0-192.0.2.10|29> sending packet: from
192.0.2.30[500] to 192.0.2.10[500] (76 bytes)
<<Month dd>> <<hh:mm:ss>> 13[NET] <ipsec0-192.0.2.10|29> received packet: from
192.0.2.10[500] to 192.0.2.30[500] (76 bytes)
<<Month dd>> <<hh:mm:ss>> 13[ENC] <ipsec0-192.0.2.10|29> parsed INFORMATIONAL response
11832 [ ]
<<Month dd>> <<hh:mm:ss>> 11[NET] <ipsec0-192.0.2.10|29> received packet: from
192.0.2.10[500] to 192.0.2.30[500] (76 bytes)
<<Month dd>> <<hh:mm:ss>> 11[ENC] <ipsec0-192.0.2.10|29> parsed INFORMATIONAL request
12924 [ ]
<<Month dd>> <<hh:mm:ss>> 11[ENC] <ipsec0-192.0.2.10|29> generating INFORMATIONAL
response 12924 [ ]
<<Month dd>> <<hh:mm:ss>> 11[NET] <ipsec0-192.0.2.10|29> sending packet: from
192.0.2.30[500] to 192.0.2.10[500] (76 bytes)
<<Month dd>> <<hh:mm:ss>> 06[IKE] <ipsec0-192.0.2.10|29> sending DPD request
```

```
<<Month dd>> <<hh:mm:ss>> 06[ENC] <ipsec0-192.0.2.10|29> generating INFORMATIONAL request
11833 [ ]
<<Month dd>> <<hh:mm:ss>> 06[NET] <ipsec0-192.0.2.10|29> sending packet: from
192.0.2.30[500] to 192.0.2.10[500] (76 bytes)
--More-- (q = quit)
```

Display IPsec Routes on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the IPsec routes configured on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display IPsec routes installed:

```
show ipsec-routes
```

Example

Displaying the IPsec routes configured on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> show ipsec-routes
192.0.2.30 via 192.0.2.20 dev eth0.125 mtu lock 1950
192.0.2.1/24 dev eth0.30 proto kernel scope link src 192.0.2.2
192.0.2.10 via 192.0.2.45 dev eth0.30
192.0.2.100/24 dev eth0.125 proto kernel scope link src 192.0.2.60
192.0.2.11/16 dev docker0 proto kernel scope link src 192.0.2.12 linkdown
```

Display IPsec Encryption Statistics on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the IPsec encryption statistics on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:
enable

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display IPsec encryption statistics:
show ipsec-stats

Example

Displaying IPsec encryption statistics on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> show ipsec-stats
src 192.0.2.30 dst 192.0.2.40
  proto esp spi 0xc0c2d9cd(3233995213) reqid 1(0x00000001) mode tunnel
  replay-window 32 seq 0x00000000 flag af-unspec (0x00100000)
  aead rfc4106(gcm(aes)) 0xa9c1923a4b4c5618ea2f3596de821261218bdea2 (160 bits) 128
  anti-replay context: seq 0x0, oseq 0x138, bitmap 0x00000000
  lifetime config:
    limit: soft (INF)(bytes), hard (INF)(bytes)
    limit: soft (INF)(packets), hard (INF)(packets)
    expire add: soft 3268(sec), hard 3600(sec)
    expire use: soft 0(sec), hard 0(sec)
  lifetime current:
    475650(bytes), 312(packets)
    add <<yyyy-mm-dd>> <<hh:mm:ss>> use <<yyyy-mm-dd>> <<hh:mm:ss>>
  stats:
    replay-window 0 replay 0 failed 0
src 192.0.2.40 dst 192.0.2.30
  proto esp spi 0xc92b08e5(3375040741) reqid 1(0x00000001) mode tunnel
  replay-window 32 seq 0x00000000 flag af-unspec (0x00100000)
  aead rfc4106(gcm(aes)) 0x9ca3568095298cefaaa709b9b932eb5141bd252c (160 bits) 128
  anti-replay context: seq 0x135, oseq 0x0, bitmap 0xffffffff
  lifetime config:
    limit: soft (INF)(bytes), hard (INF)(bytes)
    limit: soft (INF)(packets), hard (INF)(packets)
    expire add: soft 3341(sec), hard 3600(sec)
    expire use: soft 0(sec), hard 0(sec)
  lifetime current:
    470953(bytes), 309(packets)
    add <<yyyy-mm-dd>> <<hh:mm:ss>> use <<yyyy-mm-dd>> <<hh:mm:ss>>
  stats:
    replay-window 0 replay 0 failed 0
```

Display the Status of IPsec Tunnels on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the status of configured IPsec tunnel on the Fabric IPsec Gateway Virtual Machine (VM):

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display the status of IPsec tunnels configured on the VM:

```
show ipsec-status
```

Example

Displaying the status of configured IPsec tunnel on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> show ipsec-status
Status of IKE charon daemon (strongSwan 5.3.5, Linux 4.4.0-128-generic, x86_64):
  uptime: 13 days, since <<month, day hh:mm:ss year>>
  malloc: sbrk 2433024, mmap 0, used 369408, free 2063616
  worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 3
  loaded plugins: charon test-vectors aes rc2 sha1 sha2 md4 md5 random nonce x509
  revocation constraints
    pubkey pkcs1 pkcs7 pkcs8 pkcs12 pgp dnskey sshkey pem openssl fips-prf gmp agent xcbc
    hmac gcm attr
    kernel-netlink resolve socket-default connmark stroke updown
Listening IP addresses:
  192.0.2.40
  192.0.2.20
Connections:
ipsec0-192.0.2.5: 192.0.2.40...192.0.2.5 IKEv2, dpddelay=3s
ipsec0-192.0.2.5: local: [192.0.2.60] uses pre-shared key authentication
ipsec0-192.0.2.5: remote: [192.0.2.5] uses pre-shared key authentication
ipsec0-192.0.2.5: child: 192.0.2.60/32 === 192.0.2.5/32 TUNNEL, dpdaction=restart
Security Associations (1 up, 0 connecting):
ipsec0-192.0.2.5[29]: ESTABLISHED 21 hours ago,
192.0.2.40[192.0.2.60]...192.0.2.5[192.0.2.5]
ipsec0-192.0.2.5[29]: IKEv2 SPIs: dcf0a2d545d40679_i 55006e07252b9934_r*, pre-shared key
reauthentication in 2 hours
ipsec0-192.0.2.5[29]: IKE proposal: AES_CBC_128/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_2048
ipsec0-192.0.2.5[377]: INSTALLED, TUNNEL, reqid 1, ESP SPIs: c92b08e5_i c0c2d9cd_o
ipsec0-192.0.2.5[377]: AES_GCM_16_128, 291247 bytes_i (190 pkts, 6s ago), 297523 bytes_o
(194 pkts, 1s ago), rekeying in 30 minutes
ipsec0-192.0.2.5[377]: 192.0.2.60/32 === 192.0.2.5/32
```

Displays the IPsec Configuration on the Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the IPsec configuration on the Fabric IPsec Gateway Virtual Machine (VM):

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display the IPsec configuration on the configured on the VM:

```
show ipsec-config <1-255>
```

Example

Displaying the IPsec configuration on the configured on the VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

FIGW> show ipsec-config 2
ipsec {
    tunnel_id 2;
    encryption-key-length 128;
    fe-tunnel-dest-ip 10.10.10.10;
    ipsec-dest-ip 70.70.70.73;
    mtu 1950;
    responder-only false;
    tunnel-name ----;
    auth-method psk;
    cert-subject;
    auth-key *****;
    egress-shaping-rate 110;
    fragment-before-encrypt enable;
    admin-state enable
```

Display the Logical Interface IPsec Configurations on the Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the logical interface IPsec configurations on the Fabric IPsec Gateway Virtual Machine (VM):

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display the logical interface IPsec configurations on the VM:

```
show logical-intf-config <1-255>
```

Example

Displaying the logical interface IPsec configurations on the on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW>
FIGW> show logical-intf-config 1
logical-intf-tunnel {
    tunnel_id 1;
    tunnel-name ----;
    fe-tunnel-dest-ip 20.20.20.20;
    mtu 1950;
    egress-shaping-rate 110;
```

Display Current Configuration on Fabric IPsec Gateway VM

About This Task

Perform this procedure to display the parameters configured currently on the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display the parameters currently configured on the VM:

```
show running-config
```

Example

Displaying the parameters configured on Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> show running-config
set global ipsec-tunnel-src-vlan 125
set global ipsec-tunnel-src-ip 192.0.2.1/24
set global lan-intf-vlan 30
set global lan-intf-ip 192.0.2.10/24
set global lan-intf-gw-ip 192.0.2.25
```

```
set global fe-tunnel-src-ip 192.0.2.55
set global wan-intf-gw-ip 192.0.2.11
set global mtu 1950
set ipsec 1 auth-key *****
set ipsec 1 fe-tunnel-dest-ip 192.0.2.70
set ipsec 1 encryption-key-length 128
set ipsec 1 admin-state enable
```

Display Current Version of Fabric IPsec Gateway VM

About This Task

Display current version of the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Display current version of the VM:

```
show version
```

Example

Displaying current version of the Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW>show version
FabricIPSecGW_VM_4.0.0.0
```

Log Out of Fabric IPsec Gateway VM

About This Task

Perform this procedure to log out of the Fabric IPsec Gateway Virtual Machine (VM).

Procedure

1. Enter Fabric IPsec Gateway Configuration mode:

```
enable
```

```
virtual-service WORD<1-128> console
```



Note

Type **CTRL+Y** to exit the console.

2. Log out of VM:


```
exit
```

Example

Logging out of the Fabric IPsec Gateway VM:

```
Switch:1> enable
Switch:1# virtual-service figw console
Connected to domain figw
Escape character is ^Y

<cr>
FIGW> exit
```



Fabric Configuration Workflow

This section describes the generic workflow to configure SPBM and IS-IS infrastructure and services on your network.



Note

This section is a high-level overview. For further details on the SPBM and IS-IS infrastructure and configuration, see the linked sections in the Procedure sections column of the table.

Table 112: Fabric Connect configuration workflow

Steps	Procedure sections	Description
1	Fabric Basics and Layer 2 Services on page 1178	As a first step, you must configure your basic infrastructure for Shortest Path Bridging MAC (SPBM). This section also contains information about how to configure Fabric Extend, which enables the enterprise to extend Fabric Connect technology over Layer 2 or Layer 3 core networks. To use Fabric Extend with ExtremeCloud SD-WAN, see SD-WAN on page 3304.
Purpose: Services configuration		
2	Layer 2 VSN Configuration on page 1514	After you complete the infrastructure configuration, configure the appropriate services for the network to run on top of the base architecture. This includes: <ul style="list-style-type: none">• Layer 2 and Layer 3 VSNs• IP Shortcuts• Inter-VSN routing
	Fabric Layer 3 Services on page 1580	
	Inter-VSN Routing Configuration on page 1559	
Purpose: Fabric interoperation		

Table 112: Fabric Connect configuration workflow (continued)

Steps	Procedure sections	Description
3	Distributed Virtual Routing on page 915	<p>You can also configure Fabric functionality such as the following features:</p> <ul style="list-style-type: none">• Distributed Virtual Routing (DvR) <p>DvR is a technology for router redundancy in a Fabric deployment where IP subnets are stretched across multiple switches.</p> <ul style="list-style-type: none">• IP Multicast over Fabric Connect• SPB-PIM Gateway <p>SPB-PIM GW provides multicast inter-domain communication between an SPB network and a PIM network. The SPB-PIM GW can also connect two independent SPB domains.</p> <ul style="list-style-type: none">• VXLAN Gateway <p>VXLAN Gateway terminates virtual extensible LAN (VXLAN) tunnels that stretch emulated Layer 2 segments over an IP network.</p>
	IP Multicast over Fabric Connect on page 1946	
	SPB-PIM Gateway configuration on page 3557	
	VXLAN Gateway on page 4170	
Purpose: Operations and Management		
4	Connectivity Fault Management on page 3799	<p>To debug connectivity issues and isolate network faults in the SPBM network, use Connectivity Fault Management (CFM).</p>



Fabric Basics and Layer 2 Services

[SPBM and IS-IS Infrastructure Fundamentals](#) on page 1178

[Fabric Extend](#) on page 1299

[Fabric Attach](#) on page 1334

[Endpoint Tracking](#) on page 1408

[Multi-area SPB](#) on page 1422

[Layer 2 VSN Configuration](#) on page 1514

[Inter-VSN Routing Configuration](#) on page 1559

[SBPM Reference Architectures](#) on page 1569

The topics in this section provide conceptual and procedural information about basic Fabric infrastructure and the following Layer 2 Services:

- Layer 2 VSN
- Inter-VSN Routing

SPBM and IS-IS Infrastructure Fundamentals

Shortest Path Bridging MAC (SPBM) is a next generation virtualization technology that revolutionizes the design, deployment, and operations of carriers and service providers, along with enterprise campus core networks and the enterprise data center. SPBM provides massive scalability while at the same time reducing the complexity of the network.

SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet based link-state protocol that provides all virtualization services in an integrated model. In addition, by relying on endpoint service provisioning only, the idea of building your network once and not touching it again becomes a true reality. This technology provides all the features and benefits required by carrier-grade, enterprise and service provider deployments without the complexity of alternative technologies, for example, Multiprotocol Label Switching (MPLS).

SPBM simplifies deployments by eliminating the need to configure multiple points throughout the network. When you add new connectivity services to an SPBM network you do not need intrusive core provisioning. The simple endpoint provisioning is done where the application meets the network, with all points in between automatically provisioned through the robust link-state protocol, Intermediate-System-to-Intermediate-System (IS-IS).

Most Ethernet based networks use 802.1Q tagged interfaces between the routing switches. SPBM uses two Backbone VLANs (B-VLANs) that are used as the transport instance. A B-VLAN is not a traditional VLAN in the sense that it does not flood unknown, broadcast or multicast traffic, but only forwards based on IS-IS provisioned backbone MAC (BMAC) tables. After you configure the B-VLANs and the IS-IS protocol is operational, you can map the services to service instances.

SPBM uses IS-IS to discover and advertise the network topology, which enables it to compute the shortest path to all nodes in the SPBM network. SPBM uses IS-IS shortest path trees to populate forwarding tables for the individual BMAC addresses of each participating node.

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (CMAC) addresses in a backbone MAC (BMAC) address pair. MAC-in-MAC encapsulation defines a BMAC destination address (BMAC-DA) and a BMAC source address (BMAC-SA). Encapsulating customer MAC addresses in BMAC addresses improves network scalability (no end-user CMAC learning is required in the core) and also significantly improves network robustness (loops have no effect on the backbone infrastructure.)

The SPBM BMAC header includes a Service Instance Identifier (I-SID) with a length of 32 bits with a 24-bit ID. I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. You can use I-SIDs in a Virtual Services Network (VSN) for VLANs or VRFs across the MAC-in-MAC backbone:

- **Unicast**

- For a Layer 2 VSN, the device associates the I-SID with a customer VLAN, which the device then virtualizes across the backbone. Layer 2 VSNs associate one VLAN per I-SID.
- With Layer 3 VSN, the device associates the I-SID with a customer VRF, which the device virtualizes across the backbone. Layer 3 VSNs associate one VRF per I-SID.
- With Inter-VSN routing, Layer 3 devices, routers, or hosts connect to the SPBM cloud using the SPBM Layer 2 VSN service. The Backbone Core Bridge can transmit traffic between different VLANs with different I-SIDs.
- With IP shortcuts, no I-SID is required, forwarding for the Global Routing Table (GRT) is done using IS-IS based shortest path BMAC reachability.

For more information on Fabric Layer 3 services, see [Fabric Layer 3 Services](#) on page 1580.

- **Multicast**

- With Layer 2 VSN with IP multicast over Fabric Connect, the BEB associates a data I-SID with the multicast stream and the scope I-SID is based on the Layer 2 VSN I-SID.
- With Layer 3 VSN with IP multicast over Fabric Connect, the BEB associates a data I-SID with the multicast stream and the scope I-SID is based on the Layer 3 VSN I-SID.

- With IP Shortcuts with IP multicast over Fabric Connect, the BEB associates a data I-SID with the multicast stream, but there is no I-SID for the scope, which is the Global Routing Table (GRT).

For more information on IP multicast over Fabric Connect, see [IP Multicast over Fabric Connect](#) on page 1946.



Note

Inter-VSN routing for IP multicast over Fabric Connect is not supported.

The switch supports the IEEE 802.1aq standard of SPBM, which allows for larger Layer 2 topologies and permits faster convergence.

Multiple tenants using different SPBM services

The following figure shows multiple tenants using different services within an SPBM metro network. In this network, you can use some or all of the SPBM implementation options to meet the needs of the community while maintaining the security of information within VLAN members.




Figure 59: Multi-tenant SPBM metro network

To illustrate the versatility and robustness of SPBM even further, the following figure shows a logical view of multiple tenants in a ring topology. In this architecture, each tenant has its own domain where some users have VLAN requirements and are using Layer 2 VSNs and others have VRF requirements and are using Layer 3 VSNs. In all three domains, they can share data center resources across the SPBM network.




Figure 60: SPBM ring topology with shared data centers

spbm-config-mode boot flag

Shortest Path Bridging (SPB) and Protocol Independent Multicast (PIM) cannot interoperate with each other on the switch at the same time. The boot flag called **spbm-config-mode** ensures that SPB and PIM stay mutually exclusive.

- The **spbm-config-mode** boot flag is enabled by default. This enables you to configure SPB and IS-IS, but you cannot configure PIM either globally or on an interface.
- If you disable the boot flag, save the configuration and reboot with the saved configuration. After you enable the flag, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.



Important

- Any change to the **spbm-config-mode** boot flag requires a reboot for the change to take effect.
- If you disable the boot flag, save the configuration and reboot with the saved configuration. After you disable the flag, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.

For more information, see [IP Multicast](#) on page 1696.

vxlan-gw-full-interworking-mode Boot Configuration Flag

The VXLAN Gateway implementation is available in the following modes:

- **Base Interworking Mode** – This is the default mode. In this mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.
- **Full Interworking Mode** – This mode supports the Base mode communication between VXLAN and traditional VLAN environments as well as VXLAN-to-VXLAN communication and all SPB functionality including vIST and SMLT. To enter this mode, you must enable the **vxlan-gw-full-interworking-mode** boot configuration flag.



Note

Changing the mode requires a reboot for the change to take effect.

MAC-in-MAC Encapsulation

To forward customer traffic across the core network backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (CMAC) addresses in a backbone MAC (BMAC) address pair. MAC-in-MAC encapsulation defines a BMAC source address (BMAC-SA) and a BMAC destination address (BMAC-DA) to identify the backbone source and destination addresses.

The originating node creates a MAC header that is used for delivery from end to end. As the MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end.

Encapsulating customer MAC addresses in BMAC addresses improves network scalability (no end-user CMAC learning is required in the core) and also significantly improves network robustness (loops in access networks do not impact forwarding results in the backbone infrastructure.)

I-SID

SPBM introduces a service instance identifier called I-SID. SPBM uses I-SIDs to separate services from the infrastructure. After you create an SPBM infrastructure, you can add additional services (such as VLAN extensions or VRF extensions) by provisioning the endpoints only. The SPBM endpoints are Backbone Edge Bridges (BEBs), which mark the boundary between the core MAC-in-MAC SPBM domain and the edge customer 802.1Q domain. I-SIDs are provisioned on the BEBs to be associated with a particular service instance. In the SPBM core, the bridges are Backbone Core Bridges (BCBs). BCBs forward encapsulated traffic based on the BMAC-DA.

The SPBM BMAC header includes a Service Instance Identifier (I-SID) with a length of 32 bits with a 24-bit ID. I-SIDs identify a service instance for virtualized traffic in an

encapsulated SPBM frame. You can use I-SIDs in a Virtual Services Network (VSN) for VLANs or VRFs across the MAC-in-MAC backbone:

- For a Layer 2 VSN, the I-SID is associated with a customer VLAN, which is then virtualized across the backbone. Layer 2 VSNs offer an any-any LAN service type. Layer 2 VSNs associate one VLAN per I-SID.
- For a Layer 2 VSN with IP multicast over Fabric Connect, the BEB associates a data I-SID with the multicast stream and a scope I-SID that defines the scope as Layer 2 VSN. A multicast stream with a scope of Layer 2 VSN can only transmit a multicast stream for the same Layer 2 VSN.
- For a *Transparent Port UNI*, the I-SID is associated with a port or MLT, which is then virtualized across the backbone. *Transparent Port UNI* associates multiple ports or MLT to an I-SID.
- For a Layer 3 VSN, the I-SID is associated with a customer VRF, which is also virtualized across the backbone. Layer 3 VSNs are always full-mesh topologies. Layer 3 VSNs associate one VRF per I-SID.
- For a Layer 3 VSN with IP multicast over Fabric Connect, the BEB associates a data I-SID with the multicast stream and a scope I-SID that defines the scope as Layer 3 VSN. A multicast stream with a scope of Layer 3 VSN can only transmit a multicast stream for the same Layer 3 VSN.
- For IP Shortcuts with IP multicast over Fabric Connect, the BEB associates a data I-SID with the multicast stream and defines the scope as Layer 3 GRT. A multicast stream with a scope of Layer 3 GRT can only transmit a multicast stream for a Layer 3 GRT.
- When you add or remove an I-SID on a VLAN IP interface for 5520 Series, 5720 Series, the switch disables and enables routing on the VLAN interface.

**Note**

I-SID configuration is required only for virtual services such as Layer 2 VSN and Layer 3 VSN. With IP Shortcuts with unicast, no I-SID is required, forwarding for the Global Routing table is done using IS-IS based shortest path BMAC reachability.

I-SID to VLAN binding is used to automatically determine the path between client and server in order to attach network devices to FA Zero touch services.

BCBs and BEBs

Table 113: Fabric Mode product support

Feature	Product	Release introduced
Fabric BCB mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Fabric BEB mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The boundary between the core MAC-in-MAC SPBM domain and the edge customer 802.1Q domain is handled by Backbone Edge Bridges (BEBs). I-SIDs are provisioned on the BEBs to be associated with a particular service instance.

In the SPBM core, the bridges are referred to as Backbone Core Bridges (BCBs). BCBs forward encapsulated traffic based on the BMAC-DA.

**Important**

SPBM separates the payload from the transport over the SPBM infrastructure. Configure all virtualization services on the BEBs at the edge of the network. There is no provisioning required on the core SPBM switches. This provides a robust carrier grade architecture where configuration on the core switches never needs to be touched when adding new services.

A BEB performs the same functionality as a BCB, but it also terminates one or more Virtual Service Networks (VSN). A BCB does not terminate any VSNs and is unaware of

the VSN traffic it transports. A BCB simply knows how to reach any other BEB in the SPBM backbone.

VLANs without member ports

If a VLAN is attached to an I-SID there must be another instance of that same I-SID in the SPBM network.

- If another instance of that I-SID exists, the device designates that VLAN as operationally up regardless of whether it has a member port or not.

When the VLAN is operationally up, the IP address of the VLAN will be in the routing table.

- If no matching instance of the I-SID exists in the SPBM network, then that VLAN has no reachable members and does not act as a network-to-network interface (NNI).

The VLAN does not act as a UNI interface because it does not have a member port.

Therefore, the device does not designate the VLAN as operationally up because the VLAN does not act as a UNI or an NNI.

If the device acts as a BCB with two VLANs configured and two I-SIDs, there must be a UNI side with the corresponding I-SID existing in the network.

If the device acts as both BEB and BCB, then there must be a member port in that VLAN to push out the UNI traffic.

Basic SPBM Network Topology

The following figure shows a basic SPBM network topology, specifically a Layer 2 VSN. Switches A and D are the Backbone Edge Bridges (BEB) that provide the boundary between the customer VLANs (C-VLAN) and the Backbone. Switches B and C are the Backbone Core Bridges (BCB) that form the core of the SPBM network.




Figure 61: SPBM L2 VSN

SPBM uses IS-IS in the core so that all BEBs and BCBs learn the IS-IS System-ID (BMAC) of every other switch in the network. For example, BEB-A uses IS-IS to build an SPBM unicast forwarding table containing the BMAC of switches BCB-B, BCB-C, and BEB-D.

The BEBs provide the boundary between the SPBM domain and the virtualized services domain. For a Layer 2 VSN service, the BEBs map a C-VLAN to an I-SID based on local service provisioning. Any BEB in the network that has the same I-SID configured can participate in the same Layer 2 VSN.

In this example, BEB A and BEB D are provisioned to associate C-VLAN 20 with I-SID 100. When BEB A receives traffic from C-VLAN 20 that must be forwarded to the far-end location, it performs a lookup and determines that C-VLAN 20 is associated with I-SID 100 and that BEB D is the destination for I-SID 100. BEB A then encapsulates the data and C-MAC header into a new BMAC header, using its own nodal BMAC: A as the source address and BMAC: D as the destination address. BEB A then forwards the encapsulated traffic to BCB B.

To forward traffic in the core toward the destination node D, BCB B and BCB C perform Ethernet switching using the BMAC information only.

At BEB D, the node strips off the BMAC encapsulation, and performs a lookup to determine the destination for traffic with I-SID 100. BEB D identifies the destination on the C-VLAN header as C-VLAN 20 and forwards the packet to the appropriate destination VLAN and port.

E-Tree and Private VLAN Topology

Table 114: E-Tree and Private VLANs product support

Feature	Product	Release introduced
E-Tree and Private VLANs	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Routing on Private VLANs	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-4X-XT
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

Ethernet Private Tree (E-Tree) extends Shortest Path Bridging MAC (SPBM) to Private VLANs (PVLAN).

Transport within the SPBM network is achieved by associating the private VLAN with an I-SID. Flooded traffic from both promiscuous and isolated devices is transported over the same I-SID multicast tree and suppression for spoke-to-spoke traffic is done on the egress SPB Backbone Edge Bridge (BEB). This means the Private VLAN IDs are globally significant and must be the same on all BEBs.

The following list provides details for E-Tree and Private VLAN topology:

- E-Tree associates a Private VLAN with an I-SID.



Note

The same I-SID could be attached to a regular VLAN. In that case, all ports on the regular VLAN function like Promiscuous ports on the PVLAN.

- Other SPB BEBs can associate a regular CVLAN to the same I-SID that E-Tree uses.



Note

The CVLAN ID must match the primary PVLAN ID.

- CVLAN devices assigned to the same I-SID that E-Tree uses have Promiscuous connectivity within the segment.

The following figure shows a basic E-Tree network topology consisting of groups of private VLANs connected by the SPBM core network.




Figure 62: Sample E-Tree configuration

Private VLAN Port Types

The private VLAN port type is isolated, promiscuous, or trunk. If the port is a member of an MLT, then the port inherits the private VLAN type of the MLT.

In terms of network topology, the isolated port is considered a spoke. The isolated port, or spoke, does not communicate with any other isolated port in the network. The isolated port only communicates with the promiscuous ports, or hubs.

E-Tree and Private VLAN Limitations

The following limitations apply to E-Tree and Private VLAN topology:

- A port that is of Private VLAN type trunk must be tagged. Isolated and Promiscuous Private VLAN ports can be either tagged or untagged.
- When a port or MLT that has a Private VLAN type set to Isolated or Promiscuous is added to a private VLAN, if that port is used by other non private VLANs, then those non private VLANs are removed.

- Private VLAN type Isolated or Promiscuous and Flex-UNI are mutually exclusive on MLT ports.
- A port which is Private VLAN type Isolated and is tagged can belong to only one Private VLAN.

IS-IS

Table 115: IS-IS product support

Feature	Product	Release introduced
IS-IS authentication with SHA-256	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Suspend duplicate system ID detection	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 115: IS-IS product support (continued)

Feature	Product	Release introduced
Multiple IS-IS parallel adjacencies	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
IS-IS tags for external routes (BGP and OSPF)	4220 Series	Not Supported
	5320 Series	Fabric Engine 9.1
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1

SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based, link-state protocol (IS-IS). IS-IS provides virtualization services using a pure Ethernet technology base. SPBM also uses IS-IS to discover and advertise the network topology, which enables it to compute the shortest path to all nodes in the SPBM network.

IS-IS is a link-state, interior gateway protocol that was developed for the International Organization for Standardization (ISO). ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System-to-Intermediate System (IS-IS).

To provide a loop-free network and to learn and distribute network information, SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol. IS-IS is designed to find the shortest path from any one destination to any other in a dynamic fashion. IS-IS creates any-to-any connectivity in a network in an optimized, loop-free manner, without the long convergence delay experienced with the Spanning Tree Protocol. IS-IS does not block ports from use, but rather employs a specific path. As such, all links are available for use.

IS-IS dynamically learns the topology of a network and constructs unicast and multicast mesh connectivity. IS-IS parallel adjacency support allows you to configure multiple IS-IS links between the two nodes. Each node in the network calculates a shortest-path tree to every other network node based on System-IDs (BMAC addresses). Only one adjacency with the shortest path is selected as an active adjacency.

**Note**

Only an active interface with an active adjacency is added into local SPF calculations. This mechanism ensures the local node selects the shortest path and has the same view as the rest of the SPB network.

In the SPBM environment for Layer 2 VSNs, IS-IS carries only pure Layer 2 information with no requirement for an underlying IP control plane or forwarding path. IS-IS runs directly over Layer 2.

**Note**

SPBM carries Layer 3 information for Layer 3 VSNs.

In SPBM networks, IS-IS performs the following functions:

- Discovers the network topology
- Builds shortest path trees between the network nodes:
 - Forwards unicast traffic
 - Determines the forwarding table for multicast traffic
- Communicates network information in the control plane:
 - Service Instance Identifier (I-SID) information

SPBM can distribute I-SID service information to all SPBM nodes, as the I-SIDs are created. SPBM includes I-SID information in the IS-IS Link State protocol data units (PDUs). When a new service instance is provisioned on a node, its membership is flooded throughout the topology using an IS-IS advertisement.

Standard TLVs

IS-IS uses Type-Length-Value (TLV) encoding. SPBM employs IS-IS as the interior gateway protocol and implements additional TLVs to support additional functionality. The switch also supports Sub-TLVs. TLVs exist inside IS-IS packets and Sub-TLVs exist as additional information in TLVs.

The switch supports and is in full compliance with standard 802.1aq TLVs. The IEEE ratified the 802.1aq standard that defines SPBM and the Type-Length-Value (TLV)

encoding that IS-IS uses to support SPBM services. The following table lists all the TLVs that the switch supports.

Table 116: Standard TLVs

TLV	Description	Usage
1	Area addresses — The Area Addresses TLV contains the area addresses to which the IS-IS is connected.	IS-IS area
22	Extended IS reachability — The Extended IS Reachability TLV contains information about adjacent neighbors.	IS-IS adjacencies Sub-TLV 29: SPBM link metric is carried within this TLV.
129	Protocols supported — The Protocol supported TLV carries the Network Layer Protocol Identifiers (NLPID) for the Network Layer protocols where the IS-IS can be used.	SPBM in addition to existing NLPID (IPV4 0xCC, IPV6 0xE..), IEEE 802.1aq defined SPBM NLPID as 0xC1.
135	Extended IP reachability — The Extended IP Reachability TLV 135 is used to distribution IP reachability between IS-IS peers.	SPBM uses this existing IS-IS TLV to carry IP Shortcut routes in the Global Routing Table (GRT).

Table 116: Standard TLVs (continued)

TLV	Description	Usage
143	<p>Multi-topology port aware capability (MT-Port-Capability) TLV</p> <p>This TLV carries the SPB instance ID in a multiple SPB instances environment. This TLV is carried within IS-IS Hello Packets (IIH), only when parallel links exist.</p>	<p>This TLV carries the following SPBM Sub TLV:</p> <ul style="list-style-type: none"> Sub-TLV 6: SPB B-VID Sub TLV indicates the mapping between a VLAN and its equal cost tree (ECT) algorithm. To form an adjacency, both nodes must have a matching primary (B-VLAN, ECT) pair, and secondary (B-VLAN, ECT) pair, the number of B-VLANs must be equal, B-VLAN values must match, ECT values for the B-VLANs must match. Used in IS-IS Hellos only. MCID Sub TLV: The MCID is a digest of the VLANs and MSTI. Neighboring SPBM nodes must agree on the MCID to form an adjacency. The MCID is set to all zeros (0). <p>After the switch receives a non-zero MCID Sub TLV, it reflects content back to the neighbor.</p> <ul style="list-style-type: none"> Link L1 Metric Sub-TLV 7: Contains L1 metric of the link
144	<p>Multi-topology Capability (MT-Capability) TLV.</p> <p>This TLV carries the SPB instance ID in a multiple SPB instance environment. This TLV is carried within LSPs.</p> <p>In multicast over Fabric Connect, TLV 144 on the BEB bridge, where the sender is located, has the transmit (Tx) bit set. On the BEB bridge, where the receiver is located the receive (Rx) bit is set.</p>	<p>TLV 144 is the service identifier TLV. TLV 144 advertizes B-MAC and I-SID information.</p> <p>This TLV carries the following Sub TLVs:</p> <p>Sub-TLV 1: SPB instance Sub TLV contains a unique SPSourceID (nickname) to identify the SPBM node within this SPB topology.</p> <p>Sub-TLV 3: SPB Service ID (I-SID) is stored in TLV 144 sub-TLV 3. Sub-TLV 3 carries service group membership (I-SIDs) for a particular SPBM B-VLAN.</p>

Table 116: Standard TLVs (continued)

TLV	Description	Usage
184	SPBM IP VPN reachability — IS-IS TLV 184 is used to advertise SPBM L3 VSN route information across the SPBM cloud.	IP reachability for Layer 3 VSNs
185	IPVPN multicast TLV with IPMC sub TLV — The IPVPN multicast TLV contains information about the scope I-SID.	<p>TLV 185 on the BEB bridge, where the source is located, displays the multicast source and group addresses and has the transmit (Tx) bit set. Each multicast group has its own data I-SID that maps to the source and group addresses.</p> <p>As part of the IPVPN TLV, sub-TLVs define IPv4 unicast, IPv6 unicast and IPv4 multicast information.</p> <p>Layer 2 VSN IP multicast over Fabric Connect and Layer 3 VSN IP multicast over Fabric Connect (using VRF) use TLV 185.</p>
186	IP multicast TLV (GRT) — TLV 186 on the BEB bridge, where the source is located, displays the multicast source and group addresses and has the transmit (Tx) bit set. Each multicast group has its own data I-SID that maps to the source and group addresses.	<p>IP Shortcuts with IP multicast over Fabric Connect use TLV 186.</p> <p>All multicast streams are constrained within the level in which they originate, which is called the scope level.</p>
236	IPv6 Reachability — The IPv6 reachability TLV 236 is used to distribute IPv6 network reachability between IS-IS peers.	SPBM uses the existing IS-IS TLV to carry IPv6 shortcut routes through the SPBM core.

For more information on IP multicast over Fabric Connect, see [IP Multicast over Fabric Connect](#) on page 1946.

IS-IS Hierarchies

Intermediate System-to-Intermediate System (IS-IS) is a dynamic routing protocol that operates within an autonomous system (or domain). IS-IS provides support for hierarchical routing, which enables you to partition large routing domains into smaller areas. When used separately from Shortest Path Bridging MAC (SPBM), IS-IS uses a

two-level hierarchy, dividing the domain into multiple Level 1 areas and one Level 2 area. When used separately from SPBM, the Level 2 area serves as backbone of the domain, connecting to all the Level 1 areas. SPBM currently uses only Level 1 areas.

**Important**

The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled.

You can configure the Level 1 metric as 1 to 16,777,215.

Auto Level 1 Metric

You can also configure the Level 1 metric as **auto**. When you use this setting, the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed. This ensures that the fastest port speed determines the network route when you configure IS-IS and SPBM.

The Auto parameter is supported on Auto-sense enabled ports also.

Any speed change within the interface causes the port to recalculate the L1 metric. In the case of MultiLink Trunking (MLT), this update is triggered by every individual port that forms the trunk.

**Note**

You cannot configure the value to **auto** on IS-IS logical interfaces because these interfaces only support a static metric configuration.

Note the following operational considerations for the Level 1 metric auto feature:

- When updating from an earlier release, all L1 metrics previously configured on static interfaces are maintained. However, the auto-sense interfaces use the auto metric by default.
- The L1 auto metric cannot be used with the statically configured L1 metric. When you enable the auto metric, the statically configured metric is deleted. When you enable a statically configured L1 metric, the auto metric is disabled.
- To avoid conflicts, L1 auto metric should be set on both ends of the link. Otherwise, the metric is selected based on the IS-IS protocol - the highest value between the two peers.
- The IS-IS L1 metric is calculated in megabits per second (Mbps), based on the operational link speed. For example: a L1 metric = 2 000 000 / link speed on a port with a speed of 1000 Mbps, then the L1 metric = 2000000 / 1000 = 2000.
- The auto-sense IS-IS L1 metric settings are configured globally and are applied to all auto-sense ports where IS-IS is in use. All auto-sense ports will have the same L1 metric settings applied to them.
- By default, the L1 auto metric is disabled on static IS-IS interfaces and enabled on auto-sense ports.

IS-IS PDUs

Intermediate System to Intermediate System Hello (IIH) packets discover IS-IS neighbors and establish and maintain IS-IS adjacencies. An IIH is sent in every Hello-interval to maintain the established adjacency. If a node has not heard IIHs from its neighbor for (hello-interval x hello-multiple) seconds, the node tears down the adjacency. IIH carries TLV 143 and SPB B-VLAN Sub-TLV (among other sub-TLVs). For two nodes to form an adjacency the B-VLAN pairs for primary B-VLAN and secondary B-VLAN must match.

Link State Packets (LSP) advertise link state information. The system uses the link state information to compute the shortest path. LSP also advertises MT-capability TLV 144 and SPB instance Sub-TLV, and SPB I-SIDs Sub-TLV.

Complete Sequence Number Packets (CSNP) contain the most recent sequence numbers of all LSPs in the database. CSNP notifies neighbors about the local LSDB. After a neighbor receives a CSNP, it compares the LSPs in the CSNP with the LSP in the local LSDB. If the neighbor is missing LSPs, it sends a Partial Sequence Number Packets (PSNP) to request the missing LSPs. This process synchronizes the LSDBs among neighbors. A synchronized LSDB among all nodes in the network is crucial to producing a loop-free shortest path.

IS-IS Configuration Parameters

IS-IS System Identifiers

The IS-IS system identifiers consist of three parts:

- **System ID** — The system ID is any 6 bytes that are unique in a given area or level. The system ID defaults to the base Mac Address of the chassis but you can configure a non-default value. The system ID must use a unicast MAC address; do not use a multicast MAC address. A MAC address that has the low order bit 1 set in the highest byte is a multicast MAC address. For example, the following are multicast MAC addresses: x1xx.xxxx.xxxx, x3xx.xxxx.xxxx, x5xx.xxxx.xxxx, x7xx.xxxx.xxxx, x9xx.xxxx.xxxx, xBxx.xxxx.xxxx, xDxx.xxxx.xxxx, and xFxx.xxxx.xxxx.
- **Manual area** — The manual area or area ID is up to 13 bytes long. The first byte of the area number (for example, 49) is the Authority and Format Indicator (AFI). The next bytes are the assigned domain (area) identifier, which is up to 12 bytes (for example, 49.0102.0304.0506.0708.0910.1112). IS-IS supports a maximum of three manual areas, but the switch software only supports one manual area.
- **NSEL** — The last byte (00) is the n-selector. In this implementation, this part is automatically attached. There is no user input accepted.

The Network Entity Title (NET) is the combination of all three global parameters.

All routers have at least one manual area. Typically, a Level 1 router does not participate in more than one area.

The following are the requirements for system IDs:

- All IS-IS enabled routers must have one manual area and a unique system ID.
- All routers in the same area must have the same area ID.

- All routers must have system IDs of the same length (6 bytes).
- All IS-IS enabled routers must have a unique nickname.

PSNP Interval

You can change the PSNP interval rate. A longer interval reduces overhead, while a shorter interval speeds up convergence.

CSNP Periodic and Interval Rate

You can configure the CSNP periodic and interval rate. A longer interval reduces overhead, while a shorter interval speeds up convergence.

Parameters for the Link State Packet

Link state packets (LSPs) contain vital information about the state of adjacencies, which must be exchanged with neighboring IS-IS systems. Routers periodically flood LSPs throughout an area to maintain synchronization. You can configure the LSP to reduce overhead or speed up convergence.

The following list describes IS-IS parameters related to LSPs:

- The **max-lsp-gen-interval** is the time interval at which the generated LSP is refreshed. The default is 900 seconds with a range of 30 to 900.
- The **retransmit-lsp-interval** is the minimum amount of time between retransmission of an LSP. When transmitting or flooding an LSP an acknowledgement (ACK) is expected. If the ack is not received within **retransmit-lsp-interval**, the LSP is re-transmitted. The default is 5 seconds with a range of 1 to 300.

Point-to-point Mode

All SPBM links are point-to-point links. The switch does not support broadcast links.

IS-IS Interface Authentication

Configure IS-IS interface authentication to improve security and to guarantee that only trusted routers are included in the IS-IS network. Interface level authentication only checks the IIH PDUs. If the authentication type or key in a received IIH does not match the locally-configured type and key, the IIH is rejected. By default, authentication is disabled.

You can use either one of the following authentication methods:

- Simple password authentication — Uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication — Creates a Message Digest (MD5) key.
- SHA-256 — Adds a Hash-based Message Authentication Code (HMAC) digest to each IS-IS Hello packet.



Important

If the `.isis_md5key.txt` and `.isis_simplekey.txt` are missing, IS-IS adjacencies cannot be established.

Password Considerations

To reset the authentication password type, you must set the type to none.

The switch software supports only interface level authentication. The switch software does not support area level or domain level authentication.

SHA-256 Considerations

IS-IS Hello packets are sent periodically to discover IS-IS neighbors, and to establish and maintain IS-IS adjacencies. If you enable SHA-256 authentication, the switch adds an HMAC-SHA256 digest to each Hello packet.



Note

The interfaces used to make the adjacencies must have SPBM configured.

The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. If the digests match, the packet is accepted. If the digests do not match, the receiving switch discards the packet.

Directly connected switches must share the same key (secret), which can have a maximum length of 16 characters.

Hellos

To update the identities of neighboring routers, you can configure the:

- IS-IS Interface Level 1 Hello interval
- IS-IS Interface Level 1 Hello multiplier

IS-IS Interface Level 1 Hello Interval

IS-IS uses level 1 Hello packets to initialize and maintain adjacencies between neighboring routers.

You can configure the IS-IS interface level 1 Hello interval to change how often Hello packets are sent out from an interface level.

IS-IS Interface Level 1 Hello Multiplier

You can configure the IS-IS interface level 1 Hello multiplier to specify how many Hellos the switch must miss before it considers the adjacency with a neighboring switch down. By default, the hold (wait) time is the Hello interval multiplied by the Hello multiplier. By default, if the Hello interval is 9 and the Hello multiplier is 3, the hold time is 27. If the Hello multiplier is increased to 10, the hold time is increased to 90.

IS-IS Interface Level 1 Link Metric

You can configure the IS-IS interface level 1 link metric to overwrite the default metric value. By configuring the metric, you can specify a preferred path. Low cost reflects high-speed media, and high cost reflects slower media. For the wide metric, the value ranges from 1 to 16,777,215.

You can also configure the metric as *auto*. When you use the auto parameter, the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed. The IS-IS L1 metric is automatically calculated based on the operational link speed, in Mbps. The following table provides a translation between link speed and L1 metric.

Table 117: Auto L1 metric calculation

Link speed	Interface Type	L1 Metric
1 Gbps	Native Ethernet	2000
2 Gbps	MLT bundle	1000
10 Gbps	Native Ethernet	200
20 Gbps	MLT bundle	100
25 Gbps	Native Ethernet	80
40 Gbps	Native Ethernet	50
50 Gbps	MLT bundle	40
80 Gbps	MLT bundle	25
100 Gbps	Native Ethernet	20
120 Gbps	MLT bundle	17
160 Gbps	MLT bundle	13
200 Gbps	MLT bundle	10
400 Gbps	MLT bundle	5
1 Tbps	Future	2

- The switch only supports the wide metric.
- The total cost of a path equals the sum of the cost of each link.
- The default value for wide metrics is 10.



Note

When multiple paths exist to reach a node, the path with the lowest sum of metrics of the individual links is chosen. If the sum of the paths are the same, the one with the lowest number of hops is chosen. If the number of hops is the same as well, then the tie-breaking is done by the system ID.

For the primary B-VLAN, the path that has a node with the lowest system ID is chosen. Whereas, for the secondary B-VLAN, the path that has a node with the highest system ID is chosen.

Disabling IS-IS

You can disable IS-IS globally or at the interface level. If IS-IS is globally disabled, then all IS-IS functions stop. If IS-IS is enabled at the global level and disabled at one of the interface levels, then IS-IS continues on all other interfaces.

Overload Bit

A node sends the overload bit in LSP updates to inform other devices whether to use that node to pass transit traffic. For example, when a device receives an LSP with an overload bit, the device ignores that LSP in its Shortest Path First (SPF) calculation to avoid sending transit traffic through the overloaded node; however, the overloaded node can still receive traffic destined to itself.

The system activates the overload bit on bootup and clears it after 20 seconds. You can use the **overload-on-startup** parameter to control the time before the overload bit is cleared after bootup.

You can permanently configure the overload bit using the **overload** parameter. If you use this parameter, the system does not clear the overload bit after bootup and sends it in all LSP updates. If the overload bit is configured, other devices do not include this node for use as a transit node in IS-IS computations. By default, the **overload** parameter is set to false.

The **overload** and **overload-on-startup** parameters are configured under the `router isis` configuration mode in the CLI.

When IS-IS is enabled on a switch, the switch delays a reset by two seconds so that LSPs with the overload bit can be sent to all Backbone Edge Bridges (BEB) and Backbone Core Bridges (BCB) in the SPB domain.

IS-IS Route Tagging

You can apply a tag for IPv4 routes that BGP or OSPF redistribute into IS-IS. This configuration ensures that if one BEB redistributes a route into IS-IS from either protocol, other BEBs do not redistribute tagged IS-IS routes back into BGP or OSPF.

After you configure a route tag and then configure the protocol to automatically filter tagged routes, when BGP or OSPF receive a route with the matching tag, the switch blocks redistribution for that route. You can also create route-map policies to match the IS-IS route tag. For information about route policies, see [Route Filtering and IP Policies](#) on page 3249.

The switch sends an administrative tag using TLV 135 for the GRT and TLV 184 for VRFs. The switch also transmits the tags into Multi-area SPB. If an area includes platforms that cannot calculate the TLV length in TLV 135, you can disable tag transmission.



Caution

To use IS-IS Route Tagging on GRT IS-IS routes, you must also configure the metric-type as external. If you want to use IS-IS tags on GRT as internal routes, all Fabric nodes must be above a minimum software version. Any switch in the SPB Fabric that runs earlier software versions triggers an exception if you use metric type internal. To ensure this does not occur, if you attempt to configure a tag and the metric-type is not external, the switch reminds you to upgrade the software on all devices. You must ensure all devices in the network run the minimum required software.

For more information about the metric type, see [Configure the Metric Type for IS-IS Routes](#) on page 1260.

Table 118: Minimum software required

NOS	Minimum software versions
Fabric Engine	8.10.6.1 and later 9.0.5.1 and later 9.1 and later
VOSS	8.10.6.1 and later 9.0.5.1 and later 9.1 and later
VSP 8600 Series	8.1.7 and later

For CLI configuration, see the following tasks:

- [Configure SPBM IPv4 Shortcuts](#) on page 1605
- [Configure SPBM IPv4 Layer 3 VSN](#) on page 1657
- [Configure Route Redistribution to BGP](#) on page 590
- [Configure Route Redistribution to OSPF](#) on page 2809
- [Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution](#) on page 1447
- [Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution on a VRF Instance](#) on page 1451

For EDM configuration, see the following tasks:

- [Configure Route Redistribution to BGP](#) on page 648
- [Configure Route Redistribution to OSPF](#) on page 2893
- [Configure Multi-area SPB Layer 3 Redistribution in Home Area](#) on page 1476
- [Configure Multi-area SPB Layer 3 Redistribution in Remote Area](#) on page 1477
- [Configure IPv4 IS-IS Redistribution](#) on page 1634

SPBM B-VLAN

Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network.



Note

SPB internally uses spanning tree group (STG) 63 or Multiple Spanning Tree Instance (MSTI) 62. STG 63 or MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use STG 63 or MSTI 62 in the configuration, you must delete STG 63 or MSTI 62 before you can configure SPBM.

This VLAN is used for both control plane traffic and dataplane traffic.

**Note**

Always configure two B-VLANs in the core to allow load distribution over both B-VLANs.

SPBM alters the behavior of the VLAN. When a B-VLAN is associated with an SPBM network the following VLAN attributes and behaviors are modified for the B-VLAN:

- Flooding is disabled
- Broadcasting is disabled
- Source address learning is disabled
- Unknown MAC discard is disabled

You cannot add ports to a B-VLAN manually, IS-IS enabled ports are automatically added to the B-VLAN.

Essentially the BMAC addresses are programmed into the B-VLAN Forwarding Information Bases (FIBs) by IS-IS instead of the traditional VLANs flooding and learning approach.

Modification of the VLAN behavior is necessary to ensure proper control over the SPBM traffic.

Pre-populated FIB

An Ethernet network usually learns MAC addresses as frames are sent through the switch. This process is called reverse learning and is accomplished through broadcast.

SPBM does not permit broadcast flooding of traffic on the B-VLAN to prevent looping accomplished through flooding packets with unknown destinations (although multicast traffic is supported). As such, MAC addresses must be distributed within SPBM. This is accomplished by carrying the necessary BMAC addresses inside the IS-IS link state database. To that end, SPBM supports an IS-IS TLV that advertises the I-SID and BMAC information across the network. This functionality enables the powerful end-point-provisioning of SPBM.

These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB) to maximize efficiency and to enable Reverse Path Forwarding Check (RPFC) to operate properly.

RPFC

A loop prevention mechanism is required at Layer 2 to stop wayward traffic from crippling the network. Reverse Path Forwarding Check (RPFC) is the chosen method of suppressing loop traffic with SPBM. RPFC was originally designed for IP traffic at Layer 3 where it checks the source address of the packet against the routing entry in the routing table. The source address must match the route for the port it came in on otherwise the packet is illegitimate and therefore dropped.

With SPBM, the node matches the source BMAC address against the ingress port to establish validity. If the frame is not supposed to come in that port, it is immediately suppressed imposing a guaranteed loop control. If there is no VLAN FDB entry to the source MAC address with the outgoing port as the ingress port, the frame will be dropped.

IP Unicast FIB, Unicast FIB, and Unicast Tree

IP Unicast FIB

In SPBM, Backbone MAC (BMAC) addresses are carried within the IS-IS link-state database. To do this, SPBM supports an IS-IS Type-Length-Value (TLV) that advertises the Service Instance Identifier (I-SID) and BMAC information across the network. Each node has a System ID, which also serves as BMAC of the switch. These BMAC addresses are populated into the SPBM Forwarding Information Base (FIB).

When the network topology is discovered and stored in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If you only enable IP Shortcuts on the Backbone Edge Bridges, I-SIDs are never exchanged in the network as IP Shortcuts allow Global Routing Table (GRT) IP networks to be transported across IS-IS.

The **show isis spbm ip-unicast-fib** or **show isis spbm ipv6-unicast-fib** command displays all of the IS-IS routes in the IS-IS LSDB. The IP ROUTE PREFERENCE column in the show output displays the IP route preference.

Routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. IS-IS accept policies enable you to change the route preference for incoming routes. If the same route is learned from multiple sources with different route preferences, then the routes are not considered equal cost multipath (ECMP) routes. The route with the lowest route preference is the preferred route. In Layer 2, in the event of a tie-break between routes from multiple sources, the tie-breaking is based on cost and hop count.

Unicast FIB

The unicast computation runs a single Dijkstra (unlike all pair Dijkstras for multicast). SPBM produces only one Shortest Path First (SPF) tree and the tree is rooted on the computing node.

The unicast computation generates an entry for each node in the network. The Destination Address (DA) for that entry is the system-id of the node. In addition, if a node advertises MAC addresses other than the system-id, each MAC address has an entry in the unicast FIB table, and the shortest path to that MAC should be exactly the same as the path to the node.

Unicast FIB entries are installed to the vlan-fdb table.

SPBM Restrictions

RSTP and MSTP

The following list identifies restrictions and limitations associated with RSTP and MSTP:

- RSTP mode does not support SPBM.
- A C-VLAN-level loop across SPBM network-to-network interface (NNI) ports cannot be detected and needs to be resolved at the provisional level.
- SPBM NNI ports are not part of the Layer 2 VSN C-VLAN, and BPDUs are not transmitted over the SPBM tunnel. SPBM can only guarantee loop-free topologies consisting of the NNI ports. You should always use Simple Loop Prevention Protocol (SLPP) in an SMLT environment.



Note

Deploy SLPP on C-VLANs to detect loops created by customers in their access networks. However, SLPP is not required on B-VLANs, and it is not supported. The B-VLAN active topology is controlled by IS-IS that has loop mitigation and prevention capabilities built into the protocol.

- SPB internally uses spanning tree group (STG) 63 or Multiple Spanning Tree Instance (MSTI) 62. STG 63 or MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use STG 63 or MSTI 62 in the configuration, you must delete STG 63 or MSTI 62 before you can configure SPBM.
- You must configure SPBM B-VLANs on all devices in the same MSTP region. MSTP requires this configuration to generate the correct digest.
- Configure the SPBM B-VLANs to use matching VLAN IDs.

Best Practices for SPB Regarding MSTP

Use NNI ports exclusively to transport traffic for SPB-based services and not be configured as members of any VLANs other than SPB B-VLANs. In releases that do not support `nni-mstp`, when an SPBM IS-IS interface is created on an NNI port or an MLT, MSTP is automatically disabled for MSTI-62 on the port/MLT. However, MSTP is not automatically disabled on NNI ports for the CIST (default MSTI). In releases that support the **`boot config flags nni-mstp`** command, the default behavior of the MSTP NNI ports is that CIST is disabled automatically on the NNI and the NNI ports cannot be members of any VLANs other than B-VLANs. The default boot config flags `nni-mstp` must be set to false (which is the default). The following example shows the command to disable the MSTP on the NNI ports.

```
Switch:1(config)#interface gigabitEthernet 1/8
Switch:1(config-if)#no spanning-tree mstp
```

Coexistence of MSTP and SPB-Based Services on NNI Ports

In releases that do not support `nni-mstp` boot configuration, you can support the coexistence of non-SPB based services on the NNI ports, by adding NNI ports as members of VLANs, except for B-VLANs. These other VLANs rely on the use of MSTP for Loop prevention. The network operator must carefully consider the implications of

keeping MSTP enabled on the NNI ports because any MSTP topology changes detected on the NNI ports impacts all services and causes most dynamically learned information on the UNI side to be flushed and relearned. This includes, but is not limited to, all customer MAC and ARP records. This can also cause all the UNI ports on a BEB to be temporarily put into a spanning-tree blocking state before transitioning to a forwarding state again. The net result is that MSTP topology changes on the NNI ports adversely impact traffic for SPB-based services. Therefore, the NNI ports be used exclusively for SPB traffic.

SPBM IS-IS

The following list identifies restrictions and limitations associated with SPBM IS-IS:

- The switch does not support IP over IS-IS as defined by RFC 1195. IS-IS protocol is only to facilitate SPBM.
- The switch uses level 1 IS-IS. The switch does not support level 2 IS-IS. The CLI command **show isis int-12-cont1-pkts** is not supported because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.
- The IS-IS standard defines wide (32bit) metrics and narrow (8 bits) metrics. The switch supports the wide metric.
- To run IS-IS on an MLT, add the ports to the MLT, and then enable IS-IS on the MLT.

SPBM NNI SMLT

The switch does not support NNI on SMLT links.

VLACP

VLACP is generally used when a repeater or switch exists between connected switches to detect when a connection is down even when the link LED is lit. You can enable VLACP on Ethernet ports that are NNI, as well as Ethernet ports that are part of a NNI MLT.

SNMP Traps

On each SPBM peer, if you configure the SPBM B-VLANs to use different VLAN IDs, for example, VLAN 10 and 20 on one switch, and VLAN 30 and 40 on the second, the system does not generate a trap message to alert of the mismatch because the two switches cannot receive control packets from one another. Configure the SPBM B-VLANs to use matching VLAN IDs.

System MTU

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

IP Multicast over Fabric Connect

IP multicast over Fabric Connect cannot connect to existing Protocol Independent Multicast (PIM) networks that connect to SPB originated streams or that add PIM network streams into the SPB network. SPB-PIM Gateway (SPB-PIM GW), however, provides multicast interdomain communication between an SPB network and a PIM

network. SPB-PIM GW accomplishes this interdomain communication across a special Gateway VLAN. The Gateway VLAN communicates with the PIM network through the PIM protocol messaging and translates the PIM network requirements into SPB language, and vice versa. For more information about SPB-PIM GW, see [SPB-PIM Gateway configuration](#) on page 3557.

Other

The following list identifies other restrictions or considerations:

- You cannot use 3.33.33 as the SPB nickname because of a conflict with reserved IPv6 Ethernet multicast address 33:33:xx:xx:xx:xx.
- The software does not support I-SID filters.
- You cannot enable C-VLAN and B-VLAN on the same port.
- To ensure proper cleanup of MAC tables after you run the **no spbm** command, save the configuration, and then reboot the switch.

Network Load Balancing (NLB)

Table 119: Network Load Balancing product support

Feature	Product	Release introduced
Network Load Balancing (NLB) - multicast operation	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Network Load Balancing (NLB) - unicast operation	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

SPBM supports Network Load Balancing (NLB) Unicast and Multicast modes.

NLB is a clustering technology available with Microsoft Windows 2000, Microsoft Windows 2003, Microsoft Windows 2008, and Microsoft Windows 2012 server family of operating systems. You can use NLB to share the workload among multiple clustering servers. NLB uses a distributed algorithm to load balance TCP/IP network traffic across a number of hosts, enhancing the scalability and availability of mission critical, IP based services, such as web, VPN, streaming media, and firewalls. NLB also provides high availability by detecting host failures and automatically redistributing traffic to remaining operational hosts.

SPBM Script

Table 120: run spbm installation script product support

Feature	Product	Release introduced
run spbm installation script	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

You can use a CLI script to quickly configure the SPB and IS-IS infrastructure to enable Fabric Connect on a switch. You can use the SPB script, rather than manually configure the minimum SPBM and IS-IS parameters.

You can use the command **run spbm** to quickly configure the following:

- Configure the SPB Ethertype.
- Create an SPB instance.
- Create an SPBM backbone VLAN and associate it to the SPB instance.
- Create an SPBM secondary backbone VLAN and associate it to the SPB instance.
- Add an SPB nickname.
- Create a manual area.
- Enable IS-IS on one of the switch interfaces.
- Enable IS-IS globally.
- Configure the IS-IS system name.
- Configure the IS-IS system ID.

- Configure SPBM port and MLT interfaces.
- Clean up any SPBM configuration.

The following table displays the default values applied if you use the **run spbm** command. The SPB script creates some of the default values based on the MAC address of the switch, including the nickname and System ID value.

Parameter	Default values
Ethertype	0x8100
Primary B-VLAN	4051
Secondary B-VLAN	4052
Manual area	49.0000
Nickname	Derived from the chassis MAC
System name	IS-IS system name. By default, the system uses the global system name as the IS-IS system name, until you manually configure it. Also, the system does not save the global system name in the running configuration on the switch.
System ID value	Derived from the chassis MAC, using a different algorithm from that for the Nickname



Note

The SPB script only creates the SPBM instance, VLAN, and other parameters if they do not exist. For example, if the SPBM instance and VLAN already exist, the SPB script does not create them. If the SPB script is unable to create a parameter because it already exists, then the script stops and the system displays an error message.

IS-IS external metric

The software supports the IS-IS external metric to differentiate between internal and external routes with Accept Policies.

With this feature you can use IS-IS to:

- change the external metric-type of a route when redistributing it from another protocol to IS-IS through route redistribution using a route-map.
- change the external metric-type of a route when accepting a remote IS-IS route with the help of IS-IS accept policies using a route-map.
- match the external metric-type when redistributing IS-IS routes into other protocols using the match option in the route-map.
- match the external metric-type when accepting a remote IS-IS route with the help of IS-IS accept policies by using a route-map
- process the external metric-type in the route selection process.

The IS-IS metric type can also be set using the base redistribute command without using the route-map.

SPB Ethertype

The switch aligns the SPB ethertype to BCB's locally configured SPB ethertype. The BCBs mark the BTAG Ethertype of a transit MAC-in-MAC packet to match its locally configured value when it exits on a different network-to-network interface (NNI) port, even if the BTAG Ethertype on the incoming packet (CFM or SPB) does not match its configured value.



Note

IS-IS Hello packets are always marked with 0x8100 ethertype, and do not change according to the BCB's locally configured values.

Equal Cost Multipath Pathlist with Fabric Connect

Table 121: ECMP Pathlist with Fabric Connect product support

Feature	Product	Release introduced
ECMP Pathlist with Fabric Connect (IS-IS routes)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1.5
	VSP 7400 Series	VOSS 8.1.5
ECMP support for Fabric Extend	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.0

If you use Equal Cost Multipath (ECMP) in a Shortest Path Bridging (SPB) scenario, the Intermediate System-to-Intermediate System (IS-IS) protocol sends multiple routes with the same destination to the routing manager. IS-IS can add up to eight equal cost routes with the same destination to the routing table and the router uses one route for

traffic forwarding based on load management. Use the ECMP Pathlist feature to control how many equal-cost paths to add to the routing manager for the same destination.

For information about how to configure ECMP Pathlist, see [Configure ECMP](#) on page 2126 and [Configure ECMP](#) on page 2151.

FAN Transit

Fabric Area Network (FAN) transit refers to the ability of a switch to forward traffic between SPB nodes participating in a FAN. The switch is neither a part of the FAN nor does it originate or synchronize FAN traffic. The switch only forwards the traffic between the FAN end-points.

For information on how to verify the functioning of a transit switch within a FAN, see [View Detailed FAN Transit Information](#) on page 1255.

Auto-nickname for SPBM

Table 122: Auto-nickname for SPBM product support

Feature	Product	Release introduced
Auto-nickname for SPBM	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.10
	5420 Series	Fabric Engine 8.10
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.10
	VSP 7400 Series	VOSS 8.10

Each device that joins the Fabric Area Network (FAN) can automatically self-assign a random nickname from a fixed range. Automatic nicknames generate 10 seconds after the device joins the FAN. After the first IS-IS adjacency establishes, the device searches the network for a nickname server and if one is found, the device replaces the automatic nickname with the dynamic nickname assigned by the server.

The following considerations apply to automatic nickname assignment:

- Automatic nicknames generate from the range F.F0.00-F.FF.FF.
- Maximum of 4096 devices can acquire auto-nicknames.
- Configure a dynamic nickname server outside of the auto-nickname range to prevent conflicts.
- Automatic nickname behavior is always enabled unless the device has a configured static nickname.

For information about dynamic and static nicknames, see [Dynamic Nickname Assignment](#) on page 1211.


Dynamic Nickname Assignment

Table 123: Dynamic Nickname Assignment product support

Feature	Product	Release introduced
Dynamic Nickname Assignment Client mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0.5
Dynamic Nickname Assignment Server mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0.5
Extends the assignment behavior with a <i>prefix</i> parameter	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

Dynamic Nickname Assignment

Dynamic Nickname Assignment is a service that provides unique nicknames to compatible switches across a Fabric Area Network (FAN).



You can configure a node in a FAN as a nickname server. The nickname server cannot be started until you configure it with a static nickname. As a best practice, configure at least two nickname servers in a FAN to provide redundancy.



Note

Configure the nickname server with a static nickname that is outside any configured dynamic nickname server range in the network, or you can configure the nickname server with a static nickname from the first 10 values of the configured dynamic nickname server range in the network.

Dynamic Nickname Assignment Considerations

Consider the following information when implementing this feature:

- You must configure a nickname server to assign unique nicknames to clients based on established policies.
- You can configure multiple nickname servers in a FAN to provide resiliency.
- Dynamic Nickname Assignment is not supported in a FAN that contains ERS 4900 or ERS 5900 products.

Static and Dynamic Nickname Servers

You can use static nickname assignment and Dynamic Nickname Assignment in the same FAN.

You can configure Dynamic Nickname Assignment using a range prefix that can use a range from 0.00.00 to F.FF.FF. This method provides 256 groups that cover the range of 0.00.00 to F.FF.FF.

**Note**

You can configure the nickname server with a static nickname from the first 10 values of the configured dynamic nickname server range in the network.

Do not use nicknames from the dynamic nickname range when you assign nicknames statically to non-server nodes. However, if there are existing nodes in the network with static nicknames in the dynamic nickname range, it is not a requirement to change their nickname assignment. If a node is assigned a dynamic nickname that is being used in the network, duplicate nickname protection is initiated. If the node that has the dynamic nickname loses the nickname election, it requests a different nickname from the nickname server. If a node with a static nickname loses the nickname election, IS-IS is disabled on that node and you must manually re-assign the nickname and re-enable IS-IS.

You can configure nicknames from a dynamic range if the nickname server is not started.

**Note**

You must disable Dynamic Nickname Assignment before you can change the nickname prefix.

Debugging

A node must be a member of a FAN to host Dynamic Nickname Assignment applications. FAN connectivity enables the exchange of information between nickname clients and servers, such as nickname requests or nickname assignments. You can use Connectivity Fault Management (CFM) to debug connectivity issues or isolate faults. For more information about CFM, see [Connectivity Fault Management](#) on page 3799.

The following scenarios provide examples to help debug possible failure cases in single and multi-area configurations.

Example: Same SPB static nicknames in the same area

Configuration:

- Existing Switch1 with static nickname of 7.44.65 in area 49.0001
- New Switch2 with static nickname 7.44.65 in area 49.0001

Result:

- Because Switch1 has been in the network longer, it wins the election and remains operational when you add the new switch.

- IS-IS on Switch2 does not activate.
- Nodes in the Fabric identify a nickname conflict in log messages on the console.

Log messages:

```
CP1 [08/29/22 10:09:20.543:EDT] 0x00004727 00000000 GlobalRouter SNMP INFO SPBM detected
adj UP on Port1/5 , neighbor c8be.35db.6084
CP1 [08/29/22 10:09:20.596:EDT] 0x0000471b 00000000 GlobalRouter SNMP WARNING SPBM ISIS
detected duplicate nickname 7.44.65 from node c8be.35db.6084 (UNKNOWN), 2 nodes use this
nickname
CP1 [08/29/22 10:09:46.852:EDT] 0x00004726 00000000 GlobalRouter SNMP INFO SPBM detected
adj DOWN on Port1/5 , neighbor c8be.35db.6084 (5420M-24W-4YE-VOSS)
CP1 [08/29/22 10:09:46.862:EDT] 0x001dc6f6 00000000 GlobalRouter ISIS WARNING FIB-
used node bcad.ab77.4465 has the duplicate nickname 7.44.65 with Non-FIB-used node
c8be.35db.6084

1 2022-08-29T10:09:19.947-04:00 ybr-sw-idf2-vsp7400-1 CP1 - 0x0000471b - 00000000
GlobalRouter SNMP WARNING SPBM ISIS detected duplicate nickname 7.44.65 from node
c8be.35db.6084 (UNKNOWN), 2 nodes use this nickname
1 2022-08-29T10:09:23.291-04:00 ybr-sw-idf2-vsp7400-1 CP1 - 0x0000c5ec - 00300001.196
DYNAMIC CLEAR GlobalRouter HW INFO Link Up(1/5)
1 2022-08-29T10:09:46.178-04:00 ybr-sw-idf2-vsp7400-1 CP1 - 0x001dc6f6 - 00000000
GlobalRouter ISIS WARNING FIB-used node bcad.ab77.4465 has the duplicate nickname 7.44.65
with Non-FIB-used node c8be.35db.6084
```

Example: Same SPB nicknames (Switch1 static and Switch2 dynamic) in the same area

Configuration:

- Existing Switch1 with static assignment a.05.fb in area 49.0001
- New Switch2 with dynamic assignment a.05.fb in area 49.0001

Result:

- Switch2 recognizes the conflicting nickname immediately and requests a new nickname from the nickname server.
- Switch1 is unaffected.

Log messages:

```
1 2022-08-29T14:50:33.624Z Ybr-Sw-IDF1-2-NA CP1 - 0x00000008 - 00000000 GlobalRouter SW
WARNING Dynamic nickname client: detected duplicate nickname 0xa05fb. Requesting a new
nickname from the server
1 2022-08-29T14:51:43.744Z Ybr-Sw-IDF1-2-NA CP1 - 0x00004782 - 00000000 GlobalRouter SNMP
INFO SPBM Nickname a.05.fc dynamically assigned by the nickname server 209e.f792.d884
(ybr-sw-mdf-vsp7400-3)
```

Example: Same SPB nicknames (Switch1 dynamic and Switch2 static) in the same area

Configuration:

- Existing Switch1 with dynamic nickname a.05.fc in area 49.0001
- New Switch2 with static nickname a.05.fc in area 49.0001

Result:

- Because Switch1 has been in the network longer, it wins the election and remains operational when you add the new switch.
- IS-IS on Switch2 does not activate.

- Nodes in the Fabric identify a nickname conflict in log messages on the console.

Log messages:

```
1 2022-08-29T17:19:48.910Z Ybr-Sw-IDF1-2-NA CP1 - 0x001dc6de - 00000000
GlobalRouter ISIS WARNING Duplicate Nickname ID a.05.fc detected with our Node
HostName:5420M-24W-4YE-VOSS Chassis AC:c8:be:35:db:60:00 and HostName:ybr-sw-idf1-1
Chassis MAC:bc:ad:ab:77:44:01.Election won locally.
1 2022-08-29T17:19:48.910Z Ybr-Sw-IDF1-2-NA CP1 - 0x0000471b - 00000000 GlobalRouter SNMP
WARNING SPBM ISIS detected duplicate nickname a.05.fc from node bcad.ab77.4465 (UNKNOWN),
2 nodes use this nickname
1 2022-08-29T17:19:48.924Z Ybr-Sw-IDF1-2-NA CP1 - 0x001dc6de - 00000000
GlobalRouter ISIS WARNING Duplicate Nickname ID a.05.fc detected with our Node
HostName:5420M-24W-4YE-VOSS Chassis MAC:c8:be:35:db:60:00 and HstName:ybr-sw-idf1-1
Chassis MAC:bc:ad:ab:77:44:01.Election won locally.
1 2022-08-29T17:19:48.925Z Ybr-Sw-IDF1-2-NA CP1 - 0x001dc6de - 00000000
GlobalRouter ISIS WARNING Duplicate Nickname ID a.05.fc detected with our Node
HostName:5420M-24W-4YE-VOSS Chassis MAC:c8:be:35:db:60:00 and HstName:ybr-sw-idf1-1
Chassis MAC:bc:ad:ab:77:44:01.Election won locally.
1 2022-08-29T17:19:48.925Z Ybr-Sw-IDF1-2-NA CP1 - 0x001dc6de - 00000000
GlobalRouter ISIS WARNING Duplicate Nickname ID a.05.fc detected with our Node
HostName:5420M-24W-4YE-VOSS Chassis MAC:c8:be:35:db:60:00 and HstName:ybr-sw-idf1-1
Chassis MAC:bc:ad:ab:77:44:01.Election won locally.
1 2022-08-29T17:21:09.027Z Ybr-Sw-IDF1-2-NA CP1 - 0x001dc6f6 - 00000000 GlobalRouter ISIS
WARNING FIB-used node c8be.35db.6084 has the duplicate nickname a.05.fc with Non-FIB-used
node bcad.ab77.4465
```

Example: Same SPB nickname in different areas

Configuration:

- Switch1 with static nickname of 7.44.65 in area 49.0001
- Switch2 dynamically gets nickname 7.44.65 in area 49.0007

Result:

- IS-IS remains operational on both switches.
- Boundary nodes have issues forwarding traffic destined for one of the switches.

Example: Dynamic client in one area with the nickname server in another area

Configuration:

- Nickname server in area 49.0007
- Switch1 nickname client (no static assignment) in area 49.0001

Result:

- Switch1 never receives a nickname from the server across the remote links.

Auto nickname conflicts with existing static or dynamic nickname in the same area

Configuration:

- Existing Switch1 with static or dynamic nickname of f.f2.c0 in area 49.0001
- New Switch2 self-assigns nickname of f.f2.c0 in area 49.0001

Result:

- Switch2 recognizes the conflicting nickname immediately and self-assigns a new nickname from the range FF000-FFFFF until it finds one not in the LSDB.
- Switch1 is unaffected.

Log message:

```
0x003dc00c - 00000000 GlobalRouter SNMP INFO SPBM Dynamic nickname client:
detected duplicate self-generated nickname 0xff2c0. Discarding and trying a new one.
```

MSTP-Fabric Connect Multi Homing

Table 124: MSTP-Fabric Connect Multi Homing product support

Feature	Product	Release introduced
MSTP-Fabric Connect Multi Homing	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The MSTP-Fabric Connect Multi Homing feature allows MSTP or RSTP network to be multi-homed into a Fabric Connect network, providing a loop-free topology. MSTP-Fabric Connect Multi Homing enables an MSTP network to be multihomed into the SPB Fabric network through single node-to-multiple nodes or multiple nodes-to-multiple nodes.



Important

You must enable MSTP-Fabric Connect Multi Homing before you establish multihoming with an MSTP network.

MSTP-Fabric Connect Multi Homing uses I-SID 16777003. The switch creates this I-SID automatically and it cannot be modified.

MSTP-Fabric Connect Multi Homing is supported on SPBM mode only.

SPBM and IS-IS infrastructure configuration using CLI

This section provides procedures to configure SPBM and IS-IS using Command Line Interface (CLI).



Important

The **spbm-config-mode** boot flag must be enabled (default) before you can configure SPBM or IS-IS. To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

Run the SPBM Script

Use the following procedure to run the SPBM script to automate the minimum required SPBM and IS-IS parameters to enable Fabric Connect to operate on the switch.

Before You Begin

- Enable SPBM before running the SPBM script.
- Delete existing IS-IS interfaces before running this script. See [Remove Specific IS-IS and MLT Interfaces](#) on page 1220 for information on removing IS-IS interfaces.

About This Task

You can use this procedure to quickly configure the minimum SPBM and IS-IS parameters. However, a manual procedure is available instead of using this script. The default values are given in square brackets. You can input your values at the prompt or if you want to accept the default values, press **Enter**. This command first accepts all values and then removes existing SPBM configurations before configuring the entered values.



Note

This process causes the SPBM traffic to flap temporarily.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Run the SPBM script:

```
run spbm
```



Note

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

Run the SPBM script:

```
Switch:1(config)#run spbm
*****
*** This script will guide you through configuring the      ***
*** switch for optimal operation SPB.                      ***
*** -----                                              ***
*** The values in [] are the default values, you can      ***
*** input alternative values at any of the prompts.       ***
*** If you wish to terminate or exit this script         ***
*** enter ^C <control-C> at any prompt.                  ***
*** NOTE: THE COMMAND WILL TEMPORARILY FLAP IS-IS,SPBM    ***
*****
SPB Ethertype <0x8100,0x88a8> [0x8100]:
SPB primary BVLAN 2-4059 [4051]:300
SPB secondary BVLAN 2-4059 [4052]:400
ISIS system id <xxxx.xxxx.xxxx> [a051.c6eb.7c65]:0200.0000.0100
SPB nickname <x.xx.xx> [b.7c.65]:0.02.02
SPB Manual Area <xx.xxxx.xxxx...xxx> [49.0000]:50
ISIS System Name [Switch]:BEB1
Enable SPBM multicast (y/n) [n]:y
Enable IP shortcuts (y/n) [n]:y
Loopback interface ID <1-256> [1]:1
Loopback interface IP and subnet <a.b.c.d/x>:20.1.1.1/24
Configure SPBM SMLT? (y/n) [n]:y
Peer system id <xxxx.xxxx.xxxx>:0200.0000.0200
SMLT virtual BMAC <0x00:0x00:0x00:0x00:0x00:0x00>:02:00:00:10:00:10
ISIS MLT interface <MLT ID LIST>[:1]
Enable CFM SPBM (y/n) [n]:y
Enter CFM SPBM MEPID <1-8191> [1]:2
Enter CFM SPBM level <0-7> [4]:4

****CONFIGURATION IN PROGRESS****
*SPBM enabled globally*
*SPBM instance 1 configured*
*SPBM BVLANS configured*
*SPBM SMLT configured*
*SPBM multicast enabled globally*
*IP shortcuts configured*
*SPBM SMLT configured*
*IS-IS enabled*
*IS-IS on port 1/5 configured*
*IS-IS on port 1/6 configured*
*IS-IS on MLT 1 configured*
*CFM SPBM configured*
****SCRIPT EXECUTION COMPLETE****
```

Remove Existing SPBM Configuration

Use the following procedure to remove existing SPBM configurations, disable CFM, and return the CFM MEP-ID and level configurations to default values.

Before You Begin

- Enable SPBM before running this script.
- If the switch uses Zero Touch Fabric Configuration, you must run the following commands before you perform this procedure:
 - **no auto-sense onboarding i-sid**
 - **no vlan i-sid <1-4059>**

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Run the script:

```
run spbm clean
```



Note

If the script causes a configuration conflict or cannot run a command, an error message displays and the script stops.

3. To ensure proper cleanup of MAC tables, save the configuration, and then reboot the switch.

Example

Run the script:

```
Switch:1(config)#no auto-sense onboarding i-sid
Switch:1(config)#no vlan i-sid
Switch:1(config)#run spbm clean
The following will delete all SPBM and interfaces and default the CFM configurations. Do
you want to continue? <y/n>[n]:y

Switch:1(config)#no router isis enable
Switch:1(config)#interface gigabitethernet 1/10
Switch:1(config-if)#no isis
Switch:1(config-if)#interface gigabitethernet 1/11
Switch:1(config-if)#no isis
Switch:1(config-if)#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no vlan 4051
Switch:1(config)#no vlan 4052
Switch:1(config)#router isis
Switch:1(config-isis)#no spbm 1
Switch:1(config-isis)#router isis
Switch:1(config-isis)#no ip-source-address
Switch:1(config-isis)#no system-id
Switch:1(config-isis)#no manual-area 49.0000
Switch:1(config-isis)#no cfm spbm enable
Switch:1(config)#cfm spbm level 4
Switch:1(config)#cfm spbm mepid 1
Switch:1(config)#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#

**SPBM configurations have been removed**
```

Configure the IS-IS Port Interfaces Using SPBM Script

Use the following procedure to run the SPBM script to configure the IS-IS port interfaces. As this command does not flap IS-IS or SPBM, it is particularly effective to use this command when SPBM is already configured and you require to configure additional ports or MLTs. Running the run spbm interface command does not alter existing IS-IS or SPBM configurations.

About This Task

You can use this procedure to quickly configure the minimum SPBM and IS-IS parameters. However, a manual procedure is available instead of using this script.



Note

You must enable SPBM before running the SPBM script.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Run the SPBM script:

```
run spbm interface
```



Note

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

Run the SPBM script:

```
Switch:1(config)# run spbm interface
```

```
*****
*** This script will guide you through configuring the          ***
*** switch for optimal operation SPB.                          ***
*** -----                                                  ***
*** The values in [] are the default values, you can          ***
*** input alternative values at any of the prompts.           ***
*** If you wish to terminate or exit this script              ***
*** enter ^C <control-C> at any prompt.                        ***
*****
ISIS port interfaces <a/b,c/d> []:1/2,1/4,1/8
ISIS MLT interface <MLT ID LIST> []:1
*IS-IS on port 1/2 configured*
*IS-IS on port 1/4 configured*
*IS-IS on port 1/8 configured*
*IS-IS on MLT-1 configured*
```

Remove Specific IS-IS and MLT Interfaces

Use the following procedure to remove specific IS-IS ports and MLT interfaces when you get the error IS-IS SPBM interfaces have been configured. Please delete these interfaces.

About This Task

This procedure removes existing IS-IS ports and MLT interfaces. You can choose which port and MLT interfaces need to be removed. This command does not alter the other SPBM or IS-IS configurations.



Note

You must enable SPBM before running the SPBM script.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Run the script:

```
run spbm interface clean
```



Note

If the script causes a configuration conflict or cannot execute a command, an error message displays and the script stops.

Example

Run the spbm interface clean script:

```
Switch:1(config)# run spbm interface clean
```

```
*****
*** This script will guide you through deleting the ***
*** IS-IS SPBM interfaces. ***
*** ----- ***
*** The values in [] are the default values. ***
*** If you wish to terminate or exit this script ***
*** enter ^C <control-C> at any prompt. ***
*****
ISIS port interfaces to be deleted <a/b,c/d>[:1/2,1/4,1/8
ISIS MLT interface <MLT ID LIST> [:1
IS-IS port 1/2 deleted
IS-IS port 1/4 deleted
IS-IS port 1/8 deleted
** 3 IS-IS port interfaces deleted **
MLT 1 deleted
** 1 IS-IS MLTs deleted **
```

Configure Minimum SPBM and IS-IS Parameters

Use the following procedure to configure the minimum required SPBM and IS-IS parameters to enable SPBM operation on the switch.

About This Task

If the IS-IS interface uses a port, you have a choice of how you can create the IS-IS interface in this task:

- Create the IS-IS interface using three separate **isis** CLI commands.
- Use the **auto-nni** command for a quick and simple way to configure the IS-IS interface.



Note

The **auto-nni** command is not supported on MLT interfaces and Fabric Extend Logical Interfaces.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable SPBM globally:

```
spbm
```

3. Enter IS-IS Router Configuration mode:

```
router isis
```

4. Create the SPBM instance (only one SPBM instance is supported):

```
spbm <1-100>
```

5. Add the SPBM B-VLAN to the SPBM instance:

```
spbm <1-100> b-vid {<vlan-id [-vlan-id] [,...]} [primary <1-4059>]
```

6. Configure the system nickname (2.5 bytes in the format <x.xx.xx>):

```
spbm <1-100> nick-name <x.xx.xx>
```



Note

Although it is not strictly required for SPBM operation, you should change the IS-IS system ID from the default BMAC value to a recognizable address to easily identify a switch (Log on to IS-IS Router configuration mode and use the `system-id <xxxx.xxxx.xxxx>` command). This helps to recognize source and destination addresses for troubleshooting purposes.

7. Configure an IS-IS manual area (1-13 bytes in the format <xx.xxxx.xxxx...xxxx>, only one manual area is supported.):

```
manual-area <xx.xxxx.xxxx...xxxx>
```

8. Exit IS-IS Router Configuration mode to Global Configuration mode:

```
exit
```

9. Create the SPBM backbone VLAN (B-VLAN):

```
vlan create <2-4059> type spbm-bvlan
```

10. Enter Interface Configuration mode, by specifying the ports or MLTs that are going to link to the SPBM network:

```
interface {GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}| mlt <1-512> }
```

11. Depending on the type of interface, use the appropriate command method to configure an IS-IS interface:

Command	Description
auto-nni	Configure an IS-IS interface on the selected ports. This command applies only to port interfaces and negates the need to use the following three commands.
isis	Create an IS-IS interface on the selected ports or MLTs.
isis spbm <1-100>	Enable the SPBM instance on the IS-IS interfaces on the selected ports or MLTs.
isis enable	Enable the IS-IS interface on the selected ports or MLTs.

12. Enable the port or MLT interface.

13. Exit Interface Configuration mode:

```
exit
```

14. Enable IS-IS globally:

```
router isis enable
```

15. Display the SPBM configurations:

```
show isis spbm
```

16. Display the global IS-IS configuration:

```
show isis
```

17. Display the interface IS-IS configuration:

```
show isis interface [verbose]
```

Examples

The following example includes both the longer command method and the shorter **auto-nni** command. Each option is commented (###) with a heading.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#spbm
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1
Switch:1(config-isis)#spbm 1 b-vid 4051,4052 primary 10
Switch:1(config-isis)#spbm 1 nick-name 1.11.16
Switch:1(config-isis)#manual-area c0.2000.000.00
Switch:1(config-isis)#exit
Switch:1(config)#interface GigabitEthernet 1/21
### Using auto-nni command###
Switch:1(config-if)#auto-nni
### Using 3 separate IS-IS commands###
Switch:1(config-if)#isis
Switch:1(config-if)#isis spbm 1
Switch:1(config-if)#isis enable
```

```

Switch:1(config-if)#exit
Switch:1(config)#vlan create 4051 type spbm-vlan
Switch:1(config)#vlan create 4052 type spbm-vlan
Switch:1(config)#router isis enable

Switch:1(config)#show isis spbm

```

```

=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY      NICK      LSDB      IP      IPV6      MULTICAST  SPB-PIM-GW      STP-MULTI  ORIGIN
INSTANCE  VLAN      NAME      TRAP
-----
1          4051-4052  4051          disable  disable  disable  disable  disable        disable    config

```

```

=====
                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
INSTANCE
-----
1          primary          00:00:00:00:00:00          N/A          N/A

```

```

-----
Total Num of SPBM instances: 1
-----

```

```

Switch:1>show isis

```

```

=====
                        ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID : 0014.c7e1.33df
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name : Switch1
ip source-address : 41.41.41.100
ipv6 source-address : 41:0:0:0:0:0:100
ip tunnel source-address : 11.11.12.11
Tunnel vrf : spboip
ip tunnel mtu : 1950
Num of Interfaces : 2
Num of Area Addresses : 1
inband-mgmt-ip :
backbone : disabled
Dynamically Learned Area : 00.0000.0000
FAN Member : No
Hello Padding : enabled
Multi-Area OperState : disabled
Multi-Area Flags : home-always-up
Designated BN Hostname : BN4435
Designated BN Home System ID : 0010.7400.4435
Designated BN Remote System ID : 0020.7400.4435

```

```

Switch:1# show isis interface verbose

```

```

=====
                        ISIS Interfaces
=====

```

INTERFACE	TYPE	LEVEL	OPER STATE	ADM STATE	ADJ	UP ADJ	ADM-SPBM L1-METRIC	OPER-SPBM L1-METRIC	ORIGIN	AREA	AREA-NAME	PORT MEMBERS
Mlt2	pt-pt	Level 1	UP	UP	1	1	10	(A) 10	CONFIG	HOME	area-9.00.02	1/16, 1/21
Port1/21	pt-pt	Level 1	UP	UP	1	1	10	(A) 10	CONFIG	HOME	area-9.00.02	
Home: 2 out of 2 Total Num of ISIS Interfaces												
Legend:												
(A): l1 metric is automatically updated based on detected interface speed.												

Variable Definitions

The following table defines parameters for the **isis** command.

Variable	Value
<i>enable</i>	Enables or disables the IS-IS circuit/interface on the specified port or MLT. The default is disabled. Use the no option to disable IS-IS on the specified interface.
<i>spbm</i> <1-100>	Enable the SPBM instance on the IS-IS interfaces.

The following table defines parameters for the **manual-area** command.

Variable	Value
<xx.xxx.xxx...xxx>	Specifies the IS-IS manual-area (1-13 bytes in the format <xx.xxx.xxx...xxx>). Only one manual area is supported. For IS-IS to operate, you must configure at least one area. Use the no option to delete the manual area.

The following table defines parameters for the **spbm** command.

Variable	Value
<1-100>	Creates the SPBM instance. Only one SPBM instance is supported.
<i>b-vid</i> {<vlan-id [-vlan-id] [,...]}>	Sets the IS-IS SPBM instance data VLANs. Use the no option to remove the specified B-VLAN from the SPBM instance.
<i>nick-name</i> <x.xx.xx>	Specifies a nickname for the SPBM instance globally. The value is 2.5 bytes in the format <x.xx.xx>. Use the no or default options to delete the configured nickname.
<i>primary</i> <1-4059>	Sets the IS-IS instance primary data B-VLAN.

The following table defines parameters for the **vlan create** command.

Variable	Value
<2-4059>	Specifies the VLAN ID. Creates an SPBM Backbone VLAN (B-VLAN). You can optionally specify a name for the SPBM B-VLAN.
<code>type {port-mstprstp protocol- mstprstp spbm- bvlan}</code>	Specifies the type of VLAN created. <ul style="list-style-type: none"> port-mstprstp — Create a VLAN by port. protocol-mstprstp — Create a VLAN by protocol. spbm-bvlan — Create an SPBM B-VLAN.

Job aid



Important

After you have configured the SPBM nickname and enabled IS-IS. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.
4. Enable IS-IS.
5. Wait up to 20 minutes for the LSPs with the original system ID to age out.



Note

To check the age out time, use the **show isis lsdb sysid <original-sys-id>** command on any of the other SPB nodes in the network. When there is no output from this command, proceed to the next step. The time left (in seconds) for the LSPs to age out is shown under the column **LIFETIME**.

6. Disable IS-IS.
7. Change the nickname to the original nickname.
8. Enable IS-IS.

Configure I-SIDs for Private VLANs

Before You Begin

- A private VLAN must be created. For more information about creating private VLANs, see [Create a Private VLAN](#) on page 4010.

About This Task

Assign one I-SID for each private VLAN.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Assign the I-SID to the primary and secondary VLAN.

```
vlan i-sid <1-4059> <0-16777215> [force]
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#vlan i-sid 5 75
Switch:1(config)#show vlan private-vlan
```

```
=====
                PRIVATE VLAN
=====
```

Primary VLAN	Primary ISID	Secondary VLAN	Secondary ISID
5	75	6	75

Variable Definitions

The following table defines parameters for the **vlan i-sid** command.

Variable	Value
<1-4059>	Specifies the VLAN ID. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<0-16777215>	Specifies the service instance identifier (I-SID). You cannot use I-SID 0x00ffffff. The system reserves this I-SID to advertise the virtual BMAC in an SMLT dual-homing environment. This value is the same for the primary and secondary VLANs.
<i>force</i>	Specifies the software must replace the existing VLAN-to-I-SID mapping, if one exists.

Display Global SPBM Parameters

Use the following procedure to verify the global SPBM configuration.

Procedure

1. Display the SPBM configuration:

```
show isis spbm
```
2. You can also use the following command to identify SPBM VLANs. For spbm-bvlan, the attribute TYPE displays spbm-bvlan instead of byport. For private VLANs, the attribute TYPE displays private instead of byport.

```
show vlan basic
```

Example

```
Switch:1>show isis spbm
```

ISIS SPBM Info										
SPBM INSTANCE	B-VID	PRIMARY VLAN	NICK NAME	LSDB TRAP	IP	IPV6	MULTICAST	SPB-PIM-GW	STP-MULTI HOMING	ORIGIN
1	4051-4052	4051	4.00.0a	disable	enable	enable	enable	N/A	disable	config

```
=====
```

ISIS SPBM SMLT Info					
SPBM INSTANCE	SMLT-SPLIT-BEB	SMLT-VIRTUAL-BMAC	SMLT-PEER-SYSTEM-ID	SMLT-PEER-NICK-NAME	SMLT-PEER-HOST-NAME
1	primary	00:00:00:00:00:00		N/A	N/A

```
=====
```

Total Num of SPBM instances: 1

```
=====
```

```
Switch# show vlan basic
```

Vlan Basic								
VLAN ID	NAME	TYPE	INST ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRfid	ORIGIN
1	Default	byPort	0	none	N/A	N/A	0	CONFIG
10	VLAN-10	spbm-bvlan	62	none	N/A	N/A	0	CONFIG
20	VLAN-20	spbm-bvlan	62	none	N/A	N/A	0	CONFIG
100	VLAN-100	byPort	0	none	N/A	N/A	0	CONFIG

```
All 5 out of 5 Total Num of Vlans displayed
```

Display Global IS-IS Parameters

Use the following procedure to display the global IS-IS parameters.

Procedure

1. Display IS-IS configuration information:

```
show isis
```
2. Display the IS-IS system-id:

```
show isis system-id
```
3. Display IS-IS net info:

```
show isis net
```

Example

```
Switch:1>show isis
=====
                        ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID : 0014.c7e1.33df
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name : Switch1
ip source-address : 41.41.41.100
ipv6 source-address : 41:0:0:0:0:0:0:100
ip tunnel source-address : 11.11.12.11
Tunnel vrf : spboip
ip tunnel mtu : 1950
Num of Interfaces : 2
Num of Area Addresses : 1
inband-mgmt-ip :
backbone : disabled
Dynamically Learned Area : 00.0000.0000
FAN Member : No
Hello Padding : enabled
Multi-Area OperState : disabled
Multi-Area Flags : home-always-up
Designated BN Hostname : BN4435
Designated BN Home System ID : 0010.7400.4435
Designated BN Remote System ID : 0020.7400.4435

Switch# show isis system-id
=====
                        ISIS System-Id
=====
SYSTEM-ID      AREA      AREA-NAME
-----
0014.c7e1.33df    HOME
Switch# show isis net
=====
                        ISIS Net Info
=====
NET
-----
c0.2000.0000.0000.14c7.e133.df00
```

Display IS-IS Areas

Use the following procedure to display IS-IS areas.

Procedure

Use the following procedure to display IS-IS areas.

```
show isis manual-area
```

Example

```
Switch# show isis manual-area
=====
```

ISIS Manual Area Address		
AREA ADDRESS	AREA	AREA-NAME
c0.2000.0000.00	HOME	

Configure SMLT Parameters for SPBM

Use the following procedure to configure the required Split MultiLink Trunking (SMLT) parameters to allow SPBM to interoperate with SMLT on the switch.



Note

- The assignment of primary and secondary roles to the vIST peers is automatic. The switch with the lower system ID (between the two vIST peers) is primary, and the switch with the higher system ID is secondary when default system-id values are being used.
- SMLT peer system ID is part of the required configuration. You must configure the SMLT peer system ID as the nodal MAC of the peer device. In the IS-IS network, the nodal MAC of devices should be eight apart from each other.
- When using the default hardware assigned system-id value, the SMLT Virtual BMAC is automatically derived by comparing the system-id values of the two vIST peers. A value of 0x01 plus the lower of the two system-id values is used as the SMLT Virtual BMAC.

When using a manually configured system-id value, the SMLT Virtual BMAC must also be manually configured.

- An I-SID must be assigned to every VLAN that is a member of a Layer 2 VSN. Also, if a Layer 2 VSN is created on one vIST Peer, it must also be created on the other vIST peer.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Disable IS-IS on the switch:

```
no router isis enable
```
3. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```
4. Specify the system ID of the vIST peer, so that if it goes down, the local peer can take over forwarding for the failed peer:

```
spbm <1-100> smlt-peer-system-id <xxxx.xxxx.xxxx>
```

5. (Optional) Configure the virtual BMAC, which is shared and advertised by both peers:

```
spbm <1-100> smlt-virtual-bmac <0x00:0x00:0x00:0x00:0x00:0x00>
```

6. Exit to Global Configuration mode:

```
exit
```

7. Enable IS-IS on the switch:

```
router isis enable
```

8. Display the SPBM SMLT configuration:

```
show isis spbm
```

Example

```
Switch:1>enable
Switch:1#configure terminal
```

Disable IS-IS on the switch:

```
Switch:1(config)#no router isis enable
```

Enter the IS-IS Router Configuration mode:

```
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 smlt-peer-system-id 0018.b0bb.b3df
Switch:1(config-isis)#router isis enable
```

```
Switch:1(config-isis)#show isis spbm
```

```
=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY  NICK      LSDB      IP      IPV6      MULTICAST  SPB-PIM-GW      STP-MULTI
INSTANCE          VLAN      NAME      TRAP                                     HOMING
-----
1          4051-4052  4051          disable  disable  disable  enable    disable    enable
=====

                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
INSTANCE
-----
1          primary          00:00:00:00:00:00      0018.b0bb.b3df      2.10.0e      Switch-2
=====

Total Num of SPBM instances: 1
=====
```

Variable Definitions

The following table defines parameters for the **spbm** command.

Variable	Value
<code>smlt-peer-system-id</code> <xxxx.xxxx.xxxx>	Specifies the IS-IS SPBM peer system ID. SMLT peer system ID is part of the required configuration. You must configure the SMLT peer system ID as the nodal MAC of the peer device. In the IS-IS network, the nodal MAC of devices should be eight apart from each other.
<code>smlt-virtual-bmac</code> <0x00:0x00:0x00:0x00:0x00:0x00>	Specifies a virtual MAC address that can be used by both peers. SMLT virtual BMAC is an optional configuration. Note: <ul style="list-style-type: none"> If SMLT virtual BMAC is not configured, the system derives SMLT virtual BMAC from the configured SMLT peer system ID and the nodal MAC of the device (IS-IS system ID). The system compares the nodal MAC of the device with the SMLT peer system ID configured and takes the small one, plus 0x01, as the SMLT virtual BMAC. The system also derives SMLT split BEB from the SMLT peer system ID and nodal MAC of the device. The device with the lower system ID is primary, the device with the higher system ID is secondary.

Configure Optional SPBM Parameters

Use the following procedure to configure optional SPBM parameters.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Configure the SPBM ethertype:

```
spbm ethertype {0x8100 | 0x88a8}
```
- Configure the optional link-state database (LSDB) trap global parameter. To configure this parameter, you must globally disable IS-IS on the switch:
 - Disable IS-IS on the switch:

```
no router isis enable
```
 - Enter IS-IS Router Configuration mode:

```
router isis
```
 - Enable a trap when the SPBM LSDB changes:

```
spbm <1-100> lsdb-trap enable
```
 - Enable IS-IS on the switch:

```
router isis enable
```
 - Exit IS-IS Router Configuration mode:

```
exit
```

4. Configure the optional SPBM interface parameters. To configure these parameters, you must disable IS-IS on the interface:

- a. Specify an SPBM interface to configure:

```
interface {GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][,...]} | mlt <mltid> }
```

- b. Disable IS-IS on the interface:

```
no isis enable
```

- c. Configure SPBM instance interface-type on IS-IS interface. SPBM supports only pt-pt:

```
isis spbm <1-100> interface-type {broadcast|pt-pt}
```

- d. Configure the IS-IS Interface level 1 metric:

```
isis spbm <1-100> ll-metric {<1-16777215> | auto}
```

- e. Enable IS-IS on the switch:

```
isis enable
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#spbm ethertype 0x8100
Switch:1(config-isis)#no router isis enable
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 lsdb-trap enable
Switch:1(config-isis)#router isis enable
Switch:1(config-isis)#exit
Switch:1(config)#interface gigabitethernet 1/7
Switch:1(config-if)#no isis enable
Switch:1(config-if)#isis spbm 1 interface-type pt-pt
Switch:1(config-if)#isis spbm 1 ll-metric 500
Switch:1(config-if)#isis enable
```

Variable Definitions

The following table defines parameters for the **spbm** command.

Variable	Value
<i>ethertype {0x8100 0x88a8}</i>	Configures the SPBM ethertype. The default value is 0x8100.
<i><1-100> lsdb-trap enable</i>	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disabled. Use the no or default options to disable LSDB traps.

The following table defines parameters for the **isis spbm** command.

Variable	Value
<code><1-100> interface-type {broadcast pt-pt}</code>	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT. SPBM only supports the point-to-point (pt-pt) interface type. The default is pt-pt. Use the no or default options to set this parameter to the default value of pt-pt.
<code><1-100> ll-metric {<1-16777215> auto}</code>	Configures the IS-IS interface level 1 metric on the specified port or MLT. The default value is 10. Use the no or default options to set this parameter to the default. Use auto and the network route is determined by summing the lowest value metrics. The metrics are inversely proportional to port speed.

Configure Optional IS-IS Global Parameters

Use the following procedure to configure optional IS-IS global parameters.

Procedure

1. Enter IS-IS Router Configuration mode:


```
enable
configure terminal
router isis
```
2. Configure optional IS-IS global parameters:
 - a. Specify the Complete Sequence Number Packet (CSNP) interval in seconds:


```
csnp-interval <1-600>
```
 - b. Configure the router type globally:


```
is-type {l1|l12}
```
 - c. Configure the maximum level, in seconds, between generated LSPs by this Intermediate System:


```
max-lsp-gen-interval <30-900>
```
 - d. Configure the IS-IS metric type:


```
metric {narrow|wide}
```
 - e. Set or clear the overload condition:


```
overload
```
 - f. Configure the overload-on-startup value in seconds:


```
overload-on-startup <15-3600>
```
 - g. Configure the Partial Sequence Number Packet (PSNP) in seconds:


```
psnp-interval <1-120>
```
 - h. Configure the minimum time between retransmission of an LSP:


```
retransmit-lsp-interval <1-300>
```


- i. Configure the SPF delay in milliseconds:
`spf-delay <0-5000>`
- j. Configure the name for the system:
`sys-name WORD<0-255>`
- k. Configure the IS-IS system ID for the switch:
`system-id <xxxx.xxxx.xxxx>`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#csnp-interval 10
Switch:1(config-isis)#is-type 11
Switch:1(config-isis)#max-lsp-gen-interval 800
Switch:1(config-isis)#metric wide
Switch:1(config-isis)#overload
Switch:1(config-isis)#overload-on-startup 30
Switch:1(config-isis)#psnp-interval 10
Switch:1(config-isis)#retransmit-lsp-interval 10
Switch:1(config-isis)#default sys-name
Switch:1(config-isis)#spf-delay 200
Switch:1(config-isis)#default system-id
```

Variable definitions

The following table defines parameters for the **csnp-interval** command.

Variable	Value
<1-600>	Specifies the CSNP interval in seconds. This is a system level parameter that applies for level 1 CSNP generation on all interfaces. A longer interval reduces overhead, while a shorter interval speeds up convergence. The default value is 10. Use the no or default options to set this parameter to the default value of 10.

The following table defines parameters for the **is-type** command.

Variable	Value
{11 112}	Sets the router type globally: <ul style="list-style-type: none"> 11: Level-1 router type 112: Not valid. <p>The default value is 11. Use the no or default options to set this parameter to the default value of 11.</p>

The following table defines parameters for the **max-lsp-gen-interval** command.

Variable	Value
<30-900>	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate System. The default value is 900 seconds. Use the no or default options to set this parameter to the default value of 900.

The following table defines parameters for the **metric** command.

Variable	Value
{ <i>narrow</i> <i>wide</i> }	Specifies the IS-IS metric type. Only wide is supported. The default value is wide. Use the no or default options to set this parameter to the default value of wide.

The following table defines parameters for the **overload** command.

Variable	Value
<i>overload</i>	Sets or clears the overload condition. The default value is disabled. Use the no or default options to set this parameter to the default value of disabled.

The following table defines parameters for the **overload-on-startup** command.

Variable	Value
<15-3600>	Specifies the IS-IS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after IS-IS startup. The default value is 20. Use the no or default options to set this parameter to the default value of 20.

The following table defines parameters for the **psnp-interval** command.

Variable	Value
<1-120>	Specifies the PSNP interval in seconds. This is a system level parameter that applies for level 1 PSNP generation on all interfaces. A longer interval reduces overhead, while a shorter interval speeds up convergence. The default value is 2. Use the no or default options to set this parameter to the default value of 2.

The following table defines parameters for the **retransmit-lsp-interval** command.

Variable	Value
<1-300>	Specifies the minimum time between retransmission of an LSP. This defines how fast the switch resends the same LSP. This is a system level parameter that applies for Level1 retransmission of LSPs. The default value is 5 seconds. Use the no or default options to set this parameter to the default value of 5.

The following table defines parameters for the **spf-delay** command.

Variable	Value
<0-5000>	Configures the delay, in milliseconds, to pace successive Shortest Path First (SPF) runs. The timer prevents more than two SPF runs from being scheduled back-to-back. The mechanism for pacing SPF allows two back-to-back SPF runs. The default value is 100 milliseconds. Use the no or default options to set this parameter to the default value of 100 milliseconds.

The following table defines parameters for the **sys-name** command.

Variable	Value
WORD<0-255>	Specifies the IS-IS system name. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system uses the global system name as the IS-IS system name, until you manually configure it. Also, the system does not save the global system name in the running configuration on the switch. Note: The system does not display any consistency checks when you edit sys-name.

The following table defines parameters for the **system-id** command.

Variable	Value
<xxxx.xxxx.xxxx>	Specifies the IS-IS system ID for the switch. Use the no or default options to set this parameter to the default value (node BMAC).

Job aid

**Important**

After you have configured the SPBM nickname and enabled IS-IS. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.
4. Enable IS-IS.
5. Wait up to 20 minutes for the LSPs with the original system ID to age out.

**Note**

To check the age out time, use the **show isis lsdb sysid <original-sys-id>** command on any of the other SPB nodes in the network. When there is no output from this command, proceed to the next step. The time left (in seconds) for the LSPs to age out is shown under the column **LIFETIME**.

6. Disable IS-IS.
7. Change the nickname to the original nickname.
8. Enable IS-IS.

Configure Optional IS-IS Interface Level 1 Parameters

Use the following procedure to configure optional IS-IS interface level 1 parameters.

**Important**

Save your configuration using **save config** for the updates to be available after reboot. Saving the configuration also ensures that any authentication keys (passwords) specified during the configuration are properly encrypted.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface mlt <1-512>
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure optional IS-IS interface level 1 parameters:

- a. Specify the authentication type used for IS-IS hello packets on the interface:

```
isis hello-auth type {none|simple|hmac-md5|hmac-sha-256}
```

- b. If you select `simple` as the `hello-auth` type, you must also specify a key value but the key-id is optional:

```
isis hello-auth type simple key WORD<1-16> [key-id <1-255>]
```

- c. If you select `hmac-md5` or `hmac-sha-256`, you must also specify a key value. The key-id is optional:

```
isis hello-auth type hmac-md5 key WORD<1-16> [key-id <1-255>]]
```

```
isis hello-auth type hmac-sha-256 key WORD<1-16> [key-id <1-255>]]
```

- d. Configure the IS-IS Interface level 1 designated router priority:

```
isis [l1-dr-priority <0-127>]
```



Note

This parameter is not used for SPBM because SPBM only runs on point-to-point interfaces. This parameter is for designated router election on a broadcast LAN segment, which is not supported.

- e. Configure the IS-IS Interface level 1 hello interval:

```
isis [l1-hello-interval <1-600>]
```

- f. Configure the IS-IS Interface level 1 hello multiplier:

```
isis [l1-hello-multiplier <1-600>]
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config):1#interface gigabitethernet 1/1
Switch:1(config-if):1#isis
Switch:1(config-if):1#isis hello-auth type hmac-md5 key test
Switch:1(config-if):1#isis l1-dr-priority 100
Switch:1(config-if):1#isis l1-hello-interval 20
Switch:1(config-if):1#isis l1-hello-multiplier 10
Switch:1(config):1#save config
```

Variable Definitions

The following table defines parameters for the **isis** command.

Variable	Value
<pre>hello-auth type {none simple hmac-md5 hmac-sha-256}} [key [key WORD<1-16>] [key-id <1- 255>]</pre>	<p>Specifies the authentication type used for IS-IS hello packets on the interface. type can be one of the following:</p> <ul style="list-style-type: none"> • none • simple: If selected, you must also specify a key value but the key id is optional. Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5: If selected, you must also specify a key value but the key-id is optional. MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. There is an optional key ID. • hmac-sha-256: If selected, you must also specify a key value but the key-id is optional. With SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. If the digests match, the packet is accepted. If the digests do not match, the receiving switch discards the packet. There is an optional key ID. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default is none. Use the no or default options to set the hello-auth type to none.</p>
<pre>ll-dr-priority <0-127></pre>	<p>Configures the IS-IS Interface level 1 designated router priority to the specified value. The default value is 64.</p> <p>Use the no or default options to set this parameter to the default value of 64.</p> <p>Note: This parameter is not used for SPBM because SPBM only runs on point-to-point interfaces. This parameter is for designated router election on a broadcast LAN segment, which is not supported.</p>

Variable	Value
<code>l1-hello-interval <1-600></code>	Configures the IS-IS interface level 1 hello interval. The default value is 9 seconds. Use the no or default options to set this parameter to the default value of 9 seconds.
<code>l1-hello-multiplier <1-600></code>	Configures the IS-IS interface level 1 hello multiplier. The default value is 3 seconds. Use the no or default options to set this parameter to the default value of 3 seconds.

Display IS-IS Interface Parameters

Use the following procedure to display the IS-IS interface parameters.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display IS-IS interface configuration, status parameters (including adjacencies), and MLT port members:
`show isis interface [l1|l2|l12] [home|remote] [verbose]`
3. Display IS-IS interface authentication configuration:
`show isis int-auth [home|remote]`
4. Display IS-IS interface timers:
`show isis int-timers [home|remote]`
5. Display IS-IS circuit level parameters:
`show isis int-ckt-level [home|remote]`

Examples

```
Switch:1>show isis interface verbose
=====
ISIS Interfaces
=====
```

INTERFACE	TYPE	LEVEL	OPER STATE	ADM STATE	ADJ ADJ	UP ADJ	ADM-SPBM L1-METRIC	OPER-SPBM L1-METRIC	ORIGIN	AREA	AREA-NAME	PORT MEMBERS
Mlt2	pt-pt	Level 1	UP	UP	1	1	10	(A) 10	CONFIG	HOME	area-9.00.02	1/16, 1/21
Port1/21	pt-pt	Level 1	UP	UP	1	1	10	(A) 10	CONFIG	HOME	area-9.00.02	

```
-----
Home: 2 out of 2 Total Num of ISIS Interfaces
```

```
Switch:1>show isis int-auth home
```

```
=====
ISIS Interface Auth
=====
```

IFIDX	AUTH-TYPE	AUTH-KEYID	AUTH-KEY	ORIGIN	AREA	AREA-NAME
Mlt2	none	0		CONFIG	HOME	
area-9.00.02						
Port1/21	none	0		CONFIG	HOME	area-9.00.02

```
Switch:1>show isis int-timers home
```

```
=====
ISIS Interface Timers
=====
```

IFIDX	LEVEL	HELLO INTERVAL	HELLO MULTIPLIER	HELLO DR	AREA	AREA-NAME

```

-----
Mlt2          Level 1      9      3      3      HOME      area-9.00.02
Port1/21      Level 1      9      3      3      HOME      area-9.00.02

Switch:1>show isis int-ckt-level home
-----
                        ISIS Circuit level parameters
-----
IFIDX          LEVEL      DIS      CKTID      AREA      AREA-NAME
-----
Mlt2           Level 1                1      HOME      area-9.00.02
Port1/21       Level 1                2      HOME      area-9.00.02

```

Variable Definitions

The following table defines parameters for the **show isis interface** command.

Variable	Value
<i>home</i>	Displays the IS-IS interface information that the system configures in the home area.
<i>l1</i>	Displays the interface information for Level 1.
<i>l2</i>	Displays the interface information for Level 2.
<i>l12</i>	Displays the interface information for Level 1 and Level 2.
<i>remote</i>	Displays the IS-IS interface information that the system configures in the remote area.
<i>verbose</i>	Displays additional information.

The following table defines parameters for the **show isis ini-auth** command.

Variable	Value
<i>home</i>	Displays the IS-IS interface authentication information that the system configures in the home area.
<i>remote</i>	Displays the IS-IS interface authentication information that the system configures in the remote area.

The following table defines parameters for the **show isis ini-timer** command.

Variable	Value
<i>home</i>	Displays the IS-IS interface timer information that the system configures in the home area.
<i>remote</i>	Displays the IS-IS interface timer information that the system configures in the remote area.

The following table defines parameters for the **show isis ini-ckt-level** command.

Variable	Value
<i>home</i>	Displays the IS-IS interface circuit level parameters that the system configures in the home area.
<i>remote</i>	Displays the IS-IS interface circuit level parameters that the system configures in the remote area.

Display the IP Unicast FIB, Unicast FIB, and Unicast Tree

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the SPBM IP unicast FIB:

- For IPv4:

```
show isis spbm ip-unicast-fib [all] [id <1-16777215>] [spbm-nh-as-mac] [home|remote]
```

- For IPv6:

```
show isis spbm ipv6-unicast-fib [all] [id <1-16777215>] [spbm-nh-as-mac] [home|remote]
```

3. Display the SPBM unicast FIB:

```
show isis spbm unicast-fib [b-mac <0x00:0x00:0x00:0x00:0x00:0x00>] [vlan <1-4059>] [summary] [home|remote]
```

4. Display the SPBM unicast tree:

```
show isis spbm unicast-tree <1-4059> [destination <xxxx.xxxx.xxxx>]
```

Examples

```
Switch:1>show isis spbm ip-unicast-fib
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF  DEST      OUTGOING  SPBM    PREFIX  PREFIX  IP ROUTE
VRF  ISID ISID Destination  NH BEB  VLAN INTERFACE COST  COST  TYPE  PREFERENCE AREA  AREA-NAME
-----
GRT  -    -    10.133.136.0/24  4K3(*)  4058 1/3      10    1    Internal 7      HOME  area-9.00.02
GRT  -    -    10.133.136.0/24  4K3(*)  4059 1/3      10    1    Internal 7      HOME  area-9.00.02
GRT  -    -    10.133.136.0/24  4K4(*)  4058 MLT-1  10000 1    Internal 7      HOME  area-9.00.02
GRT  -    -    10.133.136.0/24  4K4(*)  4059 MLT-1  10000 1    Internal 7      HOME  area-9.00.02
-----
Home : Total number of SPBM IP-UNICAST FIB entries 4
Remote: Total number of SPBM IP-UNICAST FIB entries 0
=====
```

```
Switch:1>show isis spbm unicast-fib
```

```
=====
SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION      BVLAN SYSID      HOST-NAME      OUTGOING      COST  AREA  AREA-NAME
ADDRESS          ADDRESS          ADDRESS          INTERFACE
-----
00:16:ca:23:73:df 1000 0016.ca23.73df SPBM-1        1/21        10    HOME  area-9.00.02
00:16:ca:23:73:df 2000 0016.ca23.73df SPBM-1        1/21        10    HOME  area-9.00.02
```

```

00:18:b0:bb:b3:df 1000 0018.b0bb.b3df SPBM-2 MLT-2 10 HOME area-9.00.02
00:14:c7:e1:33:e0 1000 0018.b0bb.b3df SPBM-2 MLT-2 10 HOME area-9.00.02
00:18:b0:bb:b3:df 2000 0018.b0bb.b3df SPBM-2 MLT-2 10 HOME area-9.00.02

```

```

-----
Home:   Total number of SPBM UNICAST FIB entries 5
Remote: Total number of SPBM UNICAST FIB entries 0
-----

```

Variable Definitions

The following table defines parameters for the **show isis spbm ip-unicast-fib** command.

Variable	Value
<i>all</i>	Displays entries for the Global Routing Table (GRT) and all Virtual Routing and Forwarding (VRF) instances. Note: If you use the command show isis spbm ip-unicast-fib the device displays only GRT entries. The command shows IP routes from remote Backbone Edge Bridges (BEBs).
<i>home</i>	Displays the IS-IS SPBM IP unicast Forwarding Information Base (FIB) information that the system configures in the home area.
<i>id <1-16777215></i>	Displays IS-IS SPBM IP unicast FIB information by Service Instance Identifier (I-SID) ID.
<i>remote</i>	Displays the IS-IS SPBM IP unicast FIB information that the system configures in the remote area.
<i>spbm-nh-as-mac</i>	Displays the next hop BMAC of the IP unicast FIB entry.

The following table defines parameters for the **show isis spbm ipv6-unicast-fib** command.

Variable	Value
<i>all</i>	Displays entries for the Global Routing Table (GRT) and all Virtual Routing and Forwarding (VRF) instances. Note: If you use the command show isis spbm ipv6-unicast-fib the device displays only GRT entries. The command shows IPv6 routes from remote Backbone Edge Bridges (BEBs).
<i>home</i>	Displays the IS-IS SPBM IPv6 unicast Forwarding Information Base (FIB) information that the system configures in the home area.
<i>id <1-16777215></i>	Displays IS-IS SPBM IPv6 unicast FIB information by Service Instance Identifier (I-SID) ID.
<i>remote</i>	Displays the IS-IS SPBM IPv6 unicast FIB information that the system configures in the remote area.
<i>spbm-nh-as-mac</i>	Displays the next hop as MAC of the IPv6 unicast FIB entry.

The following table defines parameters for the **show isis spbm unicast-fib** command.

Variable	Value
<i>b-mac</i> <0x00:0x00:0x00:0x00:0x00:0x00>	Displays the FIB for the specified BMAC.
<i>home</i>	Displays the IS-IS SPBM unicast Forwarding Information Base (FIB) information that the system configures in the home area.
<i>remote</i>	Displays the IS-IS SPBM unicast FIB information that the system configures in the remote area.
<i>vlan</i> <1-4059>	Displays the FIB for the specified SPBM VLAN.
<i>summary</i>	Displays a summary of the FIB.

The following table defines parameters for the **show isis spbm unicast-tree** command.

Variable	Value
<1-4059>	Specifies the SPBM B-VLAN ID.
<i>destination</i> <xxxx.xxxx.xxxx>	Displays the unicast tree for the specified destination.

Display IS-IS LSDB and Adjacencies

Use the following procedure to display the IS-IS LSDB and adjacencies.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IS-IS LSDB:

```
show isis lsdb [level {l1|l2|l12}] [sysid <xxxx.xxxx.xxxx>] [lspid <xxxx.xxxx.xxxx.xx-xx>] [tlv <1-236>] [detail] [home|remote]
```
3. Display IS-IS adjacencies:

```
show isis adjacencies [home|remote] [verbose]
```

Example

```
Switch:1>show isis lsdb
```

```
=====
                        ISIS LSDB
=====
LSP ID                LEVEL  LIFETIME  SEQNUM  CHKSUM  HOST-NAME  AREA
-----
0014.c7e1.33df.00-00   1      545      0xb1    0xed28   NewYork    HOME
0016.ca23.73df.00-00   1     1119      0x9f    0x9c9d   Switch-Lab2 HOME
0018.b0bb.b3df.00-00   1      708      0xb9    0xcb1a   Switch-Lab1 HOME
=====
Level-1 HOME AREA: 3 out of 3 Total Num of LSP Entries
```

```

Level-1 REMOTE AREA: 0 out of 3 Total Num of LSP Entries
Level-2 HOME AREA: 0 out of 0 Total Num of LSP Entries
Level-2 REMOTE AREA: 0 out of 3 Total Num of LSP Entries

```

```
Switch:1>show isis adjacencies verbose
```

```

=====
ISIS Adjacencies
=====
INTERFACE L STATE  UPTIME    PRI  HOLDDTIME  SYSID          HOST-NAME    STATUS  AREA  AREA-NAME  PORT MEMBERS
-----
Port1/11  1  UP      05:02:18 127   22    beb0.0000.7204 Switch-Lab1  ACTIVE  HOME  area-9.00.02  1/51-1/52
Port1/12  1  UP      05:00:18 127   25    beb0.0000.7204 Switch-Lab2  BACKUP  HOME  area-9.00.02  1/49-1/50
Port1/16  1  UP      05:00:25 127   24    beb0.0000.7204 Switch-Lab3  BACKUP  HOME  area-9.00.02
-----
Home:    3 out of 3 interfaces have formed an adjacency
Remote:  0 out of 0 interfaces have formed an adjacency
=====

```

```
Switch:1> show isis lsdb detail
```

```

=====
ISIS LSDB (DETAIL)
=====

Level-1 LspID: 0001.bcb0.0003.00-001      SeqNum: 0x00000522      Lifetime: 1144
        Chksum: 0x32f7  PDU Length: 312
        Host_name: C0
        Attributes:    IS-Type 1
TLV:1   Area Addresses: 1
        c1.3000.0000.00

TLV:22  Extended IS reachability:
        Adjacencies: 7
        TE Neighbors: 7
                0000.beb1.0007.01 (Switch0)      Metric:10
                SPBM Sub TLV:
                        port id: 640 num_port 1
                        Metric: 10
                0000.beb1.00b1.01 (Switch1)      Metric:10
                SPBM Sub TLV:
                        port id: 643 num_port 1
                        Metric: 10
                0000.bcb1.0004.01 (C1)  Metric:10
                SPBM Sub TLV:
                        port id: 6144 num_port 1
                        Metric: 10
                0000.beb1.00ca.01 (Switch2)      Metric:10
                SPBM Sub TLV:
                        port id: 6156 num_port 1
                        Metric: 10
                0000.beb1.00a5.01 (VSS0)      Metric:10
                SPBM Sub TLV:

```

```

        port id: 651 num_port 1

        Metric: 10

0000.beb1.00b2.01 (VSS1)          Metric:10

        SPBM Sub TLV:

        port id: 645 num_port 1

        Metric: 10

0000.beb1.0008.01 (Switch1)      Metric:10

        SPBM Sub TLV:

        port id: 652 num_port 1

        Metric: 10

TLV:129 Protocol Supported: SPBM

TLV:137 Host_name: C0#

TLV:144 SUB-TLV 1      SPBM INSTANCE:
        Instance: 0
        bridge_pri: 0
        OUI: 00-33-33
        num of trees: 2
        vid tuple : u-bit 1 m-bit 1 ect-alg 0x80c201 base vid 1000
        vid tuple : u-bit 1 m-bit 1 ect-alg 0x80c202 base vid 1001
TLV:144 SUB-TLV 3      ISID:
        Instance: 0
        Metric: 0
        B-MAC: 00-00-bc-b1-00-03
        BVID:1000
        Number of ISID's:8
        3001 (Both), 3002 (Rx), 3003 (Both), 3004 (Rx), 4001 (Both), 4002 (
Rx), 4003 (Both), 4004 (Rx)

        Instance: 0
        Metric: 0
        B-MAC: 00-00-bc-b1-00-03

--More-- (q = quit)

```

Variable Definitions

The following table defines parameters for the **show isis lsdb** command.

Variable	Value
<i>detail</i>	Displays detailed information.
<i>home</i>	Displays the IS-IS LSDB information that the system configures in the home area.
<i>level {l1 l2 l12}}</i>	Displays the LSDB for the specified level: l1, l2, or l12.
<i>local</i>	Displays IS-IS local LSDB information.

Variable	Value
<i>remote</i>	Displays the IS-IS LSDB information that the system configures in the remote area.
<i>sysid</i> <xxxx.xxxx.xxxx>	Displays the LSDB for the specified system ID.
<i>lspid</i> <xxxx.xxxx.xxxx.xx-xx>	Displays the LSDB for the specified LSP ID.
<i>tlv</i> <1-236>	Displays the LSDB by TLV type.

The following table defines parameters for the **show isis adjacencies** command.

Variable	Value
<i>home</i>	Displays the IS-IS adjacencies that the system configures in the home area.
<i>remote</i>	Displays the IS-IS adjacencies that the system configures in the remote area.
<i>verbose</i>	Displays additional information.

Display IS-IS Statistics and Counters

Use the following procedure to display the IS-IS statistics and counters.

Procedure

1. Display IS-IS system statistics:
`show isis statistics`
2. Display IS-IS interface counters:
`show isis int-counters [home|remote]`
3. Display IS-IS level 1 control packet counters:
`show isis int-l1-ctrl-pkts [home|remote]`



Note

The switch uses level 1 IS-IS. The switch does not support level 2 IS-IS. The command **show isis int-l2-ctrl-pkts** is not supported because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.

4. Clear IS-IS statistics:
`clear isis stats [error-counters] [packet-counters]`

Example

```
Switch:1>show isis statistics
=====
                        ISIS System Stats
=====
LEVEL   CORR  AUTH  AREA  MAX SEQ  SEQ NUM  OWN LSP  BAD ID  PART  LSP DB  AREA  AREA-NAME
LSPs    FAILS  DROP  EXCEEDED  SKIPS  PURGE   LEN      CHANGES OLOAD
-----
Level-1  0      0      0      1      0      0      0      0      0      HOME  area-9.00.02
Level-1  0      0      0      1      0      0      0      0      0      REMOTE area-9.00.02

Switch:1>show isis int-counters
=====
```

```

=====
ISIS Interface Counters
=====
IFIDX      LEVEL  AUTH   ADJ   INIT  REJ   ID LEN  MAX AREA  LAN DIS  AREA  AREA-NAME
          FAILS  CHANGES  FAILS  ADJ
-----
Mlt2      Level 1  0       1       0       0       0       0       0       HOME  area-9.00.02
Port1/21  Level 1  0       1       0       0       0       0       0       HOME  area-9.00.02

Switch:1>show isis int-l1-ctrl-pkts
=====
ISIS L1 Control Packet counters
=====
IFIDX      DIRECTION    HELLO    LSP     CSNP     PSNP     AREA  AREA-NAME
-----
Mlt2      Transmitted   13346    231     2        229     HOME  area-9.00.02
Mlt2      Received     13329    230     1        230     HOME  area-9.00.02
Port1/21  Transmitted   13340    227     2        226     HOME  area-9.00.02
Port1/21  Received     13335    226     1        227     HOME  area-9.00.02

```

Variable Definitions

The following table defines parameters for the **show isis int-counters** command.

Variable	Value
<i>home</i>	Displays the IS-IS interface counters that the system configures in the home area.
<i>remote</i>	Displays the IS-IS interface counters that the system configures in the remote area.

The following table defines parameters for the **show isis int-l1-ctrl-pkts** command.

Variable	Value
<i>home</i>	Displays the IS-IS L1 control packet counters that the system configures in the home area.
<i>remote</i>	Displays the IS-IS L1 control packet counters that the system configures in the remote area.

The following table defines parameters for the **clear isis stats** command.

Variable	Value
<i>error-counters</i>	Clears IS-IS stats error-counters.
<i>packet-counters</i>	Clears IS-IS stats packet-counters.

Suspend Duplicate System ID Detection When Replacing a Switch

When a switch is replaced and the original system ID and nickname is used, you must wait up to 20 minutes for the LSPs with the original system to age out. This is due to duplicate system ID and nickname detection. However, you can suspend duplicate detection on the replacement switch so that you can bring the switch into the network immediately.

About This Task

To temporarily disable duplicate detection on the replacement switch, perform the following steps:

Procedure

1. Copy the configuration file of the original switch to the replacement switch.
2. Power up the replacement switch while it is not connected to the SPB network, that is, network-to-network interface (NNI) ports are not connected.
3. Disable IS-IS on the original switch, or remove the switch from the network.
4. On the replacement switch, enter the following Global Configuration command to suspend duplicate detection for up to 21 minutes:
`isis dup-detection-temp-disable`
5. To check the remaining time, use the `show isis dup-detection-temp-disable remaining time` command.
6. Remove the original switch from the network.
7. Connect the replacement switch to the network.

Configure Dynamic Nickname Assignment

About This Task

Use this procedure to specify the nickname allocation prefix and enable Dynamic Nickname Assignment.



Note

You must disable Dynamic Nickname Assignment before you can change the nickname prefix.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the Dynamic Nickname Assignment nickname allocation:
`spbm nick-name server prefix x.xx.xx`
3. Enable Dynamic Nickname Assignment:
`spbm nick-name server`
4. Verify the configuration:
`show spbm`

Examples

Configure a nickname allocation prefix and enable Dynamic Nickname Assignment:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#spbm nick-name server prefix C.30.00
Switch:1(config)#spbm nick-name server
```


Dynamic Nickname Assignment configuration values and their associated behavior are shown in the following output from the **show spbm** command:

```
Switch:1>show spbm
                        spbm : enable
                        ethertype : 0x8100
                        nick-name server : enable
                        nick-name allocation : static
                        nick-name server range : C.30.00-C.3F.FF
```

Variable Definitions

The following table defines parameters for the **spbm nick-name server** command.

Variable	Value
<i>prefix x.xx.xx</i>	Specifies the nickname server allocation prefix. x.xx.xx uses the form X.X0.00 from 0.00.00 to F.F0.00. A group, X.X0.00 to X.XF.FF, can provide up to 4,096 nicknames. The default nickname allocation range is A.00.00-A.0F.FF.

Display Dynamic Nickname Assignment

About This Task

Use this procedure to display the current status and values for Dynamic Nickname Assignment.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display Dynamic Nickname Assignment configuration values:

```
show spbm
```

Example

The following example displays Dynamic Nickname Assignment configuration values and their associated behavior:

```
Switch:1>show spbm
                        spbm : enable
                        ethertype : 0x8100
                        nick-name server : disable
                        nick-name allocation : static
                        nick-name server range : A.00.00-A.0F.FF
```

Enable MSTP-Fabric Connect Multi Homing

Before You Begin

- You must configure a nickname for the specific SPBM instance on which you enable MSTP-Fabric Connect Multi Homing.
- Configure the MSTP global version as *mstp*.

About This Task

Perform this procedure to enable MSTP-Fabric Connect Multi Homing for a specific SPBM instance.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```
2. Enable MSTP-Fabric Connect Multi Homing on a specified SPBM instance:

```
spbm <1-100> stp-multi-homing enable
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 stp-multi-homing enable
```

Variable definitions

The following table defines parameters for the **spbm** command.

Variable	Value
<1-100>	Specifies the IS-IS SPBM instance ID to create an SPBM instance.
<i>stp-multi-homing enable</i>	Enables MSTP-Fabric Connect Multi Homing on the specific SPBM instance. The default is disabled.

Determine the Root Bridge in an MSTP-Fabric Connect Multi Homing Configuration

Identify the root bridge by determining where the Common and Internal Spanning Tree (CIST) regional root MAC address is learned for the STP-reserved I-SID. Check which MAC address has the same first five octets as the CIST regional root MAC address.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Determine the CIST regional root:

```
show spanning-tree mstp status
```
3. Determine which MAC address has the same first five octets as seen in the CIST regional root MAC address:

```
show i-sid mac-address-entry 16777003
```

Example

In the following example, bold text identifies the relevant information in the command output. In the output of the second command, the DEST-MAC shows 10:cd:ae:6e:d8:84, which is the system ID of the CIST regional root BEB, and the system name is BEB-1000.

```
Switch:1>show spanning-tree mstp status
```

```
=====
                        MSTP Status
=====
```

```

Bridge Address           : b0:ad:aa:4d:b8:00
Cist Root                : 80:00:10:cd:ae:6e:d8:00
Cist Regional Root      : 80:00:10:cd:ae:6e:d8:00
Cist Root Port          : fabric
Cist Root Cost           : 0
Cist Regional Root Cost : 2000000
Cist Instance Vlan Mapped : 1-1024
Cist Instance Vlan Mapped2k : 1025-2048
Cist Instance Vlan Mapped3k : 2049-3072
Cist Instance Vlan Mapped4k : 3073-4050,4053-4059
Cist Max Age             : 20 seconds
Cist Forward Delay      : 15 seconds

```

```
Switch:1>show i-sid mac-address-entry 16777003
```

I-SID Fdb Table									
I-SID	STATUS	MAC-ADDRESS	INTERFACE	TYPE	DEST-MAC	BVLAN	DEST-SYSNAME	AREA-ROLE	AREA-NAME
16777003	learned	10:cd:ae:6e:d8:82	Port-1/9	NON-LOCAL	10:cd:ae:6e:d8:84	4051	BEB-1000	HOME	
area-20.0020									
16777003	learned	10:cd:ae:db:a4:83	Port-1/40	NON-LOCAL	10:cd:ae:db:a4:84	4051	7208	HOME	area-20.0020
16777003	learned	b0:ad:aa:40:14:82	Port-1/40	NON-LOCAL	b0:ad:aa:40:14:84	4051	6222	REMOTE	area-20.0020

Increase the Number of SPB Nodes per Area



Note

This procedure only applies to 5320 Series (not 5320-16P-2MXT-2X or 5320-24T-4X-XT) and 5420 Series.

Perform this procedure to increase the number of SPB nodes per area that the switch supports. For more information about maximum scaling numbers, see *Fabric Engine Release Notes*.

About This Task

If you enable this boot config flag, it decreases the scaling limits for the following features:

- Switched UNI (S-UNI) endpoints
- Layer 2 and Layer 3 I-SIDs
- IP Multicast over Fabric Connect local streams
- Private VLANs

The scaling numbers for the following items are limited by the use of a shared hardware resource, and thus they affect one another:

- the number of SPB nodes per area
- the number of SPB nodes that send IP Multicast streams that the local BEB receives
- local SPB services:
 - Layer 2 I-SIDs
 - Layer 3 I-SIDs
 - IP Multicast over Fabric Connect remote IP Multicast stream data I-SIDs
 - I-SID mirroring instances

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Increase the number of supported SPB nodes per area:
`boot config flags spbm-node-scaling`
3. Verify the configuration:
`show boot config flags`
4. Save the configuration:
`save config`
5. Restart the switch for the change to take effect:
`reset`

Example

Enable the boot config flag to increase the number of SPB nodes per area that switch supports.

```
Switch:1>en
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags spbm-node-scaling
Switch:1(config)#save config
Switch:1(config)#reset

Warning: Please save the configuration and reboot the switch
         for this configuration to take effect.
```



Note

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
```

```
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

View Detailed FAN Transit Information

About This Task

Use this procedure to verify detailed FAN Transit information of a switch that acts as a transit node in a FAN. Transit nodes are not members of the FAN and only forward FAN traffic.

This procedure also includes verification on the FAN member nodes.


Procedure

1. Verify information on the FAN member nodes:
 - a. Verify that the node is a FAN member:
`show isis`
 - b. Verify that the node signals FAN membership on TLV 147:
`show isis lsdb local tlv 147 detail`
 - c. Verify that the node creates FAN multicast FIB entries for itself and the other member nodes:
`show isis spbm multicast-fib`
 - d. Verify that the transit node is a part of the FAN tree:
`l2 tracetree-fan`
 - e. Verify that the transit node is in the SPB path between the FAN member nodes:
`show isis spbm unicast-tree <2-4059>`
2. Verify information on the FAN transit node:
 - a. Verify that the transit node is not a member of the FAN and does not signal FAN membership on TLV 147:
`show isis`

`show isis lsdb local tlv 147 detail`
 - b. Verify that the transit node creates FAN multicast FIB entries for the FAN member nodes:
`show isis spbm multicast-fib`

Example

Use the following example to verify FAN Transit information on a transit node between two end-point SPB nodes, BEB-A and BEB-B.



Verification on BEB-A:



Note

You can run the following commands on either BEB-A or BEB-B. The following example displays the sample outputs for BEB-A.

Verify that BEB-A is a member of the FAN. The `FAN Member` attribute displays as `Yes`.

```
BEB-A:1#show isis
```

```
=====
                        ISIS General Info
=====
AdminState : enabled
RouterType : Level 1
System ID : b0ad.aa41.9c84
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name : BEB-A
ip source-address :
ipv6 source-address :
ip tunnel source-address :
Tunnel vrf :
ip tunnel mtu :
Num of Interfaces : 1
Num of Area Addresses : 1
inband-mgmt-ip :
backbone : disabled
Dynamically Learned Area : 00.0000.0000
FAN Member : Yes
```

```
Multi-Area OperState : disabled
Hello Padding : enabled
```

Verify that BEB-A signals FAN membership. FAN membership is signaled when TLV 147 displays the FAN multicast address.

```
BEB-A:1#show isis lsdb local tlv 147 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----
Level-1 LspID: b0ad.aa41.9c84.00-00      SeqNum: 0x00000010      Lifetime: 412
      Chksum: 0xac95 PDU Length: 122
      Host_name: BEB-A
      Attributes: IS-Type 1
TLV:147 Chassis MAC: b0:ad:aa:41:9c:00

TLV:147 FAN Mcast Addr: b1:ad:aa:41:9c:84
```

Verify that BEB-A creates FAN multicast FIB entries for itself and other FAN nodes on I-SID 16777001:

```
BEB-A:1#show isis spbm multicast-fib
```

```
=====
SPBM MULTICAST FIB ENTRY DETAIL INFO
=====
-----
MCAST DA      ISID      BVLAN SYSID      HOST-NAME      OUTGOING-INTERFACES      INCOMING      AREA      AREA-NAME
                                INTERFACE
-----
65:6a:52:ce:04:84  16777001  4058  646a.52ce.0484  BEB-B                                1/20      HOME      area-10.0010
65:6a:52:ce:04:84  16777001  4059  646a.52ce.0484  BEB-B                                1/20      HOME      area-10.0010
b1:ad:aa:41:9c:84  16777001  4058  b0ad.aa41.9c84  BEB-A      1/20      cpp      HOME      area-10.0010
-----
Total number of SPBM MULTICAST FIB entries 3
```

Verify that the transit node is in the SPB path between the FAN end-points BEB-A and BEB-B. 4058 and 4059 are the SPB B-VIDs.



Note

You can view the SPB B-VIDs by executing the command **show isis spbm**.

```
BEB-A:1#show isis spbm unicast-tree 4058
Node:646a.52ce.0484 (BEB-B) -> Node:f873.a202.53df (TRANSIT) -> ROOT
Node:f873.a202.53df (TRANSIT) -> ROOT

BEB-A:1#show isis spbm unicast-tree 4059
Node:646a.52ce.0484 (BEB-B) -> Node:f873.a202.53df (TRANSIT) -> ROOT
Node:f873.a202.53df (TRANSIT) -> ROOT
```

Verify that the transit node is a part of the FAN tree:

```
BEB-A:1#l2 tracetree-fan
```

```
Please wait for l2tracetree to complete or press any key to abort
```

```
l2tracetree to b1:ad:aa:41:9c:84, vlan 4058 i-sid 16777001 nickname 0.11.03 hops 64
1  BEB-A      b0:ad:aa:41:9c:84 -> TRANSIT      f8:73:a2:02:53:df
2  TRANSIT    f8:73:a2:02:53:df -> BEB-B      64:6a:52:ce:04:84
```

Verification on the Transit node:

Verify that the transit node is not a member of the FAN. The `FAN Member` parameter does not display.

```
TRANSIT:1#show isis
=====
                        ISIS General Info
=====
                        AdminState : enabled
                        RouterType : Level 1
                        System ID : f873.a202.53df
Max LSP Gen Interval : 900
                        Metric : wide
Overload-on-startup : 20
                        Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
                        spf-delay : 100
Router Name : TRANSIT
ip source-address :
ip tunnel source-address :
Tunnel vrf :
ip tunnel mtu :
Num of Interfaces : 2
Num of Area Addresses : 1
inband-mgmt-ip :
backbone : disabled
Multi-Area OperState : disabled
Hello Padding : enabled
```

Verify that the transit node does not signal FAN membership on TLV 147. The `FAN Mcast Addr` parameter is not displayed.

```
TRANSIT:1#show isis lsdb local tlv 147 detail
=====
                        ISIS LSDB (DETAIL)
=====
-----
Level-1 LspID: f873.a202.53df.00-00      SeqNum: 0x00000013      Lifetime: 1149
      Chksum: 0x424d PDU Length: 135
      Host_name: TRANSIT
      Attributes: IS-Type 1
TLV:147 Chassis MAC: f8:73:a2:02:50:00
```

Verify the transit node creates the FAN Multicast FIB entries for the FAN end-points on I-SID 16777001.

```
TRANSIT:1#show isis spbm multicast-fib
=====
                        SPBM MULTICAST FIB ENTRY DETAIL INFO
=====
=====
```

MCAST DA	ISID	BVLAN	SYSID	HOST-NAME	OUTGOING-INTERFACES	INCOMING INTERFACE	AREA	AREA-NAME
65:6a:52:ce:04:84	16777001	4058	646a.52ce.0484	BEB-B	1/20	1/9	HOME	area-10.0010
65:6a:52:ce:04:84	16777001	4059	646a.52ce.0484	BEB-B	1/20	1/9	HOME	area-10.0010
b1:ad:aa:41:9c:84	16777001	4058	b0ad.aa41.9c84	BEB-A	1/9	1/20	HOME	area-10.0010
b1:ad:aa:41:9c:84	16777001	4059	b0ad.aa41.9c84	BEB-A	1/9	1/20	HOME	area-10.0010

```
-----
Total number of SPBM MULTICAST FIB entries 4
=====
```


IS-IS external metric configuration using the CLI

This section provides procedures for IS-IS external metric configuration.

Match Metric Type for IS-IS Routes

Before You Begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must log on to the route-map configuration mode in the CLI.

About This Task

Use this procedure to match the external metric-type by using a route-map for any of the following cases:

- accepting a remote IS-IS route with the help of IS-IS accept policies.
- redistributing IS-IS routes into other protocols.

Procedure

1. Enter Route-Map Configuration mode:
enable

configure terminal

route-map WORD<1-64> <1-65535>
2. Match IS-IS metric type:
match metric-type-isis {any|internal|external}
3. Permit the route policy action:
permit
4. Enable the route policy
enable

Examples

Match metric type for IS-IS routes:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
```

Match metric type for IS-IS routes in accept policies:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router isis
Switch:1(config-isis)# accept route-map rol
Switch:1(config-isis)# exit
Switch:1(config)# isis apply accept
```

Match metric type to redistribute IS-IS routes into some other protocol (OSPF,RIP,BGP)

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router bgp
Switch:1(router-bgp)# redistribute isis route-map rol
Switch:1(router-bgp)# exit
Switch:1(config)# ip bgp apply redistribute
```

Variable Definitions

The following table defines parameters for the **match metric-type-isis** command.

Variable	Value
metric-type-isis {any internal external}	<p>Specifies the IS-IS metric type:</p> <ul style="list-style-type: none"> • internal – permits or denies routes that are internal to the IS-IS domain. • external – permits or denies routes that originate from an external routing protocol domain. • any – permits or denies both internal routes as well as external routes.

Configure the Metric Type for IS-IS Routes

Before You Begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must log on to the route-map configuration mode in the CLI.

About This Task

Use this procedure to configure the IS-IS external metric-type by using a route-map for any of the following cases:

- accept a remote IS-IS route with the help of IS-IS accept policies
- redistribute routes from other protocols into IS-IS



Caution

To use IS-IS Route Tagging on GRT IS-IS routes, you must also configure the metric-type as external. If you want to use IS-IS tags on GRT as internal routes, all Fabric nodes must be above a minimum software version. Any switch in the SPB Fabric that runs earlier software versions triggers an exception if you use metric type internal. To ensure this does not occur, if you attempt to configure a tag and the metric-type is not external, the switch reminds you to upgrade the software on all devices. You must ensure all devices in the network run the minimum required software.

Procedure

1. Enter Route-Map Configuration mode:
enable

configure terminal

route-map *WORD*<1-64> <1-65535>
2. Configure the IS-IS metric type:
set metric-type-isis {any|internal|external}
3. Permit the route policy action:
permit
4. Enable the route policy
enable

Examples

Configure the metric type for IS-IS routes:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# set metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
```

Set metric type for IS-IS routes in accept policies:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# set metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router isis
Switch:1(config-isis)# accept route-map rol
Switch:1(config-isis)# exit
Switch:1(config)# isis apply accept
```

Set metric type to redistribute routes from other protocols into IS-IS:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# route-map rol 10
Switch:1(route-map)# match metric-type-isis internal
Switch:1(route-map)# permit
Switch:1(route-map)# enable
Switch:1(route-map)# exit
Switch:1(config)# router isis
Switch:1(config-isis)# redistribute bgp route-map rol
Switch:1(config-isis)# exit
Switch:1(config)# isis apply redistribute
```

Variable Definitions

The following table defines parameters for the **set metric-type-isis** command.

Variable	Value
metric-type-isis {any internal external}	<p>Specifies the IS-IS metric type:</p> <ul style="list-style-type: none">• internal – permits or denies routes that are internal to the IS-IS domain.• external – permits or denies routes that originate from an external routing protocol domain.• any – permits or denies both internal routes as well as external routes.

Configure the Metric Type for IS-IS Routes Using Global Redistribute Command

Before You Begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must log on to the IS-IS router configuration mode in the CLI.

About This Task

Use this procedure to set the IS-IS external metric-type using the global redistribute command for the following cases redistributing routes from other protocols into IS-IS.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```
2. Set IS-IS metric type using global redistribute command:

```
redistribute direct metric-type {internal|external}
```
3. Enable the route policy

```
redistribute direct enable
```

Example

Set metric type for IS-IS routes using global redistribute command:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# router isis
Switch:1(config-isis)# redistribute direct metric-type internal
Switch:1(config-isis)# redistribute direct enable
```

Variable Definitions

The following table defines parameters for the **redistribute direct metric-type** command.

Variable	Value
metric-type {internal external}	Specifies the IS-IS metric type: <ul style="list-style-type: none"> internal – permits or denies routes that are internal to the IS-IS domain. external – permits or denies routes that originate from an external routing protocol domain.

SPBM and IS-IS infrastructure configuration using EDM

This section provides procedures to configure basic SPBM and IS-IS infrastructure using Enterprise Device Manager (EDM).



Important

The **EnableSpbmConfigMode** boot flag must be enabled (default) before you can configure SPBM or IS-IS. To verify the setting, navigate to **Configuration > Edit > Chassis** and select the **Boot Config** tab.

Configure Required SPBM and IS-IS Parameters

About This Task

Perform this procedure to configure the minimum required SPBM and IS-IS parameters to enable SPBM to operate on the switch. SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol to provide a loop free Ethernet topology that creates a shortest path topology from every node to every other node in the network based on node MAC addresses.

For additional information about the tabs referenced in this procedure, see the following sections:

- **VLAN > VLANs > Basic** tab: [Basic Field Descriptions](#) on page 4047
- **Fabric > SPBM > Globals** tab: [Globals Field Descriptions](#) on page 1292
- **Fabric > SPBM > SPBM** tab: [SPBM Field Descriptions](#) on page 1995
- **Fabric > IS-IS > Manual Area** tab: [Manual Area or Manual Area Remote Field Descriptions](#) on page 1292
- **Fabric > IS-IS > Globals** tab: [Globals Field Descriptions](#) on page 1324
- **Fabric > IS-IS > Interfaces** tab: [Interfaces Field Descriptions](#) on page 1284
- **Fabric > IS-IS > Interfaces > SPBM** tab: [Interfaces SPBM Field Descriptions](#) on page 1274

Procedure

1. Configure SBPM B-VLANs:



Note

Always configure two B-VLANs in the core to enable load distribution over both B-VLANs.

- a. In the navigation pane, expand **Configuration > VLAN**.
- b. Select **VLANs**.
- c. Select the **Basic** tab.
- d. Select **Insert**.
- e. In **type**, select **spbm-bvlan**.
- f. Select **Insert**.

2. Enable SPBM globally:

- a. In the navigation pane, expand **Configuration > Fabric**.
- b. Select **SPBM**.
- c. Select the **Globals** tab.
- d. In **GlobalEnable**, select **enable** to enable SPBM globally.
- e. Select **Apply**.

3. Create an SPBM instance:



Note

Only one SBPM instance is supported.

- a. In the navigation pane, expand **Configuration > Fabric**.
 - b. Select **SPBM**.
 - c. Select the **SPBM** tab.
 - d. Select **Insert** to create an SPBM instance.
 - e. In **Id**, specify the SPBM instance ID.
 - f. In **NodeNickName**, specify the node nickname (valid value is 2.5 bytes in the format <x.xx.xx>).
 - g. In **Vlans**, specify the IDs of the SPBM B-VLANs to add to the SPBM instance.
 - h. In **PrimaryVlan**, specify which of the SPBM B-VLANs specified in the previous step is the primary B-VLAN.
 - i. Select **Insert**.
- ### 4. Create a manual area:



Note

Only one manual area is supported.

- a. In the navigation pane, expand **Configuration > Fabric**.
- b. Select **IS-IS**.
- c. Select the **Manual Area** tab.
- d. Select **Insert**.

- e. Specify the Manual Area address (a valid value is 1–13 bytes in the format <xx.xxxx.xxxx...xxxx>).
 - f. Select **Insert**.
5. Update the default IS-IS system ID to a recognizable address:
- a. In the navigation pane, expand **Configuration > Fabric**.
 - b. Select **IS-IS**.
 - c. Select the **Globals** tab.
 - d. In **SystemId**, update the default BMAC value to a recognizable address.

**Note**

Although it is not strictly required for SPBM operation, you must change the IS-IS system ID from the default BMAC value to a recognizable address to easily identify a switch. This helps to recognize the source and destination addresses for troubleshooting purposes.

- e. In **AdminState**, select **on**.
 - f. Select **Apply**.
6. Create an IS-IS circuit and enable SPBM on the circuit:
- a. In the navigation pane, expand **Configuration > Fabric**.
 - b. Select **IS-IS**.
 - c. Select the **Interfaces** tab.
 - d. Select **Insert** to create an IS-IS circuit.
 - e. In **IfIndex**, specify the port or MLT on which to create the IS-IS circuit.
 - f. Select **Insert**.
 - g. Select the newly created IS-IS circuit entry, and select **SPBM**.
 - h. In the **Interfaces SPBM** tab, select **Insert**.
 - i. In **State**, select **enable**.
 - j. Select **Insert** to enable the SPBM instance on the IS-IS circuit.
 - k. Navigate back to the **Interfaces** tab.
 - l. In **AdminState** for the IS-IS circuit entry, select **on** to enable the IS-IS circuit.
 - m. Select **Apply**.

Job aid

**Important**

After you have configured the SPBM nickname and enabled IS-IS. To maintain the same nickname with a different system ID, perform the following steps:

1. Disable IS-IS.
2. Change the system ID.
3. Change the nickname to a temporary one.
4. Enable IS-IS.
5. Wait up to 20 minutes for the LSPs with the original system ID to age out.

**Note**

To check the age out time, use the **show isis lsdb sysid <original-sys-id>** command on any of the other SPB nodes in the network. When there is no output from this command, proceed to the next step. The time left (in seconds) for the LSPs to age out is shown under the column **LIFETIME**.

6. Disable IS-IS.
7. Change the nickname to the original nickname.
8. Enable IS-IS.

*Display SPBM and IS-IS Summary Information***Procedure**

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Protocol Summary** tab.

Protocol Summary Field Descriptions

Use the data in the following table to use the **Protocol Summary** tab.

Name	Description
Globals ISIS	
AdminState	Indicates the global status of IS-IS on the switch.
SystemId	Indicates the IS-IS system ID for the switch. Valid value is a 6-byte value in the format <xxxx.xxxx.xxxx>
HostName	Indicates IS-IS system name. This may be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, IS-IS system name uses global system name as it is synchronized with the IS-IS system name until the IS-IS sytem name is manually configured.
Globals SPBM	

Name	Description
GlobalEnable	Indicates whether SPBM is enabled or disabled at the global level.
NodeNickName	Indicates the nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Indicates the primary VLAN ID for this SPBM instance.
SmltSplitBEB	Indicates whether the switch is the primary or secondary IST peer.
ISIS Interfaces	
Circuit Index	Displays the identifier of this IS-IS circuit, unique within the Intermediate System. This is for SNMP Indexing purposes only and need not have any relation to any protocol value.
IfIndex	Indicates the interface to which this circuit corresponds.
AdminState	Indicates the administrative state of the circuit: on or off.
OperState	Indicates the operational state of the circuit: up or down.
PortMembers	Displays the members of an MLT.
ISIS Adjacency View	
Circuit Index	Displays the identifier of this IS-IS circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
AdjIndex	Displays a unique value identifying the IS adjacency from all other such adjacencies on this circuit. This value is automatically assigned by the system when the adjacency is created
AdjIfIndex	Indicates the interface to which this circuit corresponds.
AdjState	Indicates the state of the adjacency: <ul style="list-style-type: none"> • down • initializing • up • failed
AdjNeighSysID	Indicates the system ID of the neighboring Intermediate System.
AdjHostName	Indicates the host name listed in the LSP, or the IS-IS system name. The IS-IS system name uses global system name, if the host name is not configured.
PortMembers	Displays the members of an MLT.

View the SPBM I-SID Information

Use the following procedure to display the SPBM Service Instance Identifier (I-SID) information. The SPBM BMAC header includes an I-SID with a length of 24 bits. This I-SID can be used to identify and transmit any virtualized traffic in an encapsulated SPBM frame.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **I-SID** tab.

I-SID Field Descriptions

Use the data in the following table to use the **I-SID** tab.

Name	Description
SysId	Indicates the system identifier.
Vlan	Indicates the B-VLAN where this I-SID was configured or discovered.
Isid	Indicates the IS-IS SPBM I-SID identifier.
NickName	Indicates the nickname of the node where this I-SID was configured or discovered.
HostName	Indicates the host name listed in the LSP, or the IS-IS system. The IS-IS system name uses global system name, if the host name is not configured.
Type	Indicates the SPBM I-SID type; either configured or discovered.

View Level 1 Area Information

Use the following procedure to display Level 1 area information. IS-IS provides support for hierarchical routing, which enables you to partition large routing domains into smaller areas. IS-IS uses a two-level hierarchy, dividing the domain into multiple Level 1 areas and one Level 2 area. The Level 2 area serves as backbone of the domain, connecting to all the Level 1 areas.

**Important**

The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **L1 Area** tab.

L1 Area Field Descriptions

Use the data in the following table to use the **L1 Area** tab.

Name	Description
AreaAddr	Specifies an area address reported in a Level 1 link-state packets (LSP) generated or received by this Intermediate System.

Configure SMLT Parameters for SPBM

Use the following procedure to configure the required Split MultiLink Trunking (SMLT) parameters to enable SPBM to interoperate with SMLT on the switch.



Note

- The assignment of primary and secondary roles to the vIST peers is automatic. The switch with the lower system ID (between the two vIST peers) is primary, and the switch with the higher system ID is secondary when default system-id values are being used.
- SMLT peer system ID is part of the required configuration. You must configure the SMLT peer system ID as the nodal MAC of the peer device. In the IS-IS network, the nodal MAC of devices should be eight apart from each other.
- When using the default hardware assigned system-id value, the SMLT Virtual BMAC is automatically derived by comparing the system-id values of the two vIST peers. A value of 0x01 plus the lower of the two system-id values is used as the SMLT Virtual BMAC.

When using a manually configured system-id value, the SMLT Virtual BMAC must also be manually configured.

- An I-SID must be assigned to every VLAN that is a member of an Layer 2 VSN. Also if an Layer 2 VSN is created on one vIST Peer, it must also be created on the other vIST peer.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **SPBM** tab.
4. Use **SmltSplitBEB** to see whether the switch is the primary or secondary vIST peer. This field cannot be modified.
5. In **SmltVirtualBmac**, specify a virtual MAC address that can be used by both peers.
6. In **SmltPeerSysId**, specify the vIST peer BMAC address.
7. Select **Apply**.

Enable or Disable SPBM at the Global Level

Use the following procedure to enable or disable SPBM at the global level. SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol to

provide a loop free Ethernet topology that creates a shortest path topology from every node to every other node in the network based on node MAC addresses.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **Globals** tab.
4. To enable or disable SPBM, select **enable** or **disable** in **GlobalEnable**.
5. To configure the global ethertype value, select the desired option in **GlobalEtherType**.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
GlobalEnable	Enables or disables SPBM globally. The default is disabled. To ensure proper cleanup of MAC tables after you disable SPBM, save the configuration, and then reboot the switch.
GlobalEtherType	Specifies the global ethertype value as 0x8100 or 0x88a8. The default value is 0x8100.
NicknameServerEnable	Enables or disables the nickname server. The default is disabled.
NicknameDynamicAllocationStatus	Displays the Dynamic Nickname Allocation service operational status.
NicknameServerPrefix	Specifies the nickname server allocation prefix. x.xx.xx uses the form X.X0.00 from 0.00.00 to F.F0.00. A group, X.X0.00 to X.XF.FF, can provide up to 4,096 nicknames. The default is A.00.00.

Configure SPBM parameters

Use the following procedure to configure SPBM global parameters. SPBM uses the Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol to provide a loop free Ethernet topology that creates a shortest path topology from every node to every other node in the network based on node MAC addresses.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **SPBM** tab.
4. To create an SPBM instance, select **Insert**.
5. Configure the SPBM parameters.

6. Select **Apply**.

SPBM Field Descriptions

Use the data in the following table to use the **SPBM** tab.

Name	Description
Id	Specifies the SPBM instance ID. Only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Specifies the primary SPBM B-VLANs to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.
LsdbTrap	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.
IpShortcut	Enables or disables SPBM IP shortcut state. The default is disable.
SmltSplitBEB	Specifies whether the switch is the primary or secondary vIST peer. The default is primary.
SmltVirtualBmac	Specifies a virtual MAC address that can be used by both peers.
SmltPeerSysId	Specifies the system ID of the SPBM SMLT for this SPBM instance.
Mcast	Specifies if IP multicast over SPBM is enabled. The default is disabled.
McastFwdCacheTimeout	Specifies the global forward cache timeout in seconds. The default is 210 seconds.
Ipv6Shortcut	Enables or disables SPBM IPv6 shortcut state. The default is disable.
McastSpbPimGwControllerEnable	Enables or disables IS-IS PLSB Multicast SPB PIM Gateway controller. Disabled by default.
McastSpbPimGwGatewayEnable	Enables or disables IS-IS PLSB Multicast SPB PIM Gateway. Disabled by default.
StpMultiHoming	Enables or disables MSTP-Fabric Connect Multi Homing. The default is disabled (false).

Name	Description
BVlanOrigin	Shows how the B-VLAN was created. The values can be config for manual configuration using CLI or SNMP, or dynamic through Zero Touch Fabric Configuration and Auto-sense. The default is dynamic.
RoutedMulticastTtl1Bridging	Enables or disables routed SPB IP Multicast TTL 1 bridging. The default is disabled.

Displaying SPBM nicknames

Use the following procedure to display SPBM nicknames.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **Nick Names** tab.

Nickname Field Descriptions

Use the data in the following table to use the **Nickname** tab.

Name	Description
Level	Indicates the level at which the system displays this LSP.
ID	Indicates the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number.
LifetimeRemain	Indicates the remaining lifetime in seconds for the LSP.
NickName	Indicates the nickname for the SPBM node.
HostName	Indicates the hostname listed in the LSP, or the IS-IS system name. The IS-IS system name uses global system name, if the host name is not configured.

Configure Interface SPBM Parameters

Use the following procedure to configure the SPBM interface parameters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **Interfaces SPBM** tab.
4. Configure the SPBM interface parameters.
5. Select **Apply**.

Interfaces SPBM Field Descriptions

Use the data in the following table to use the **Interfaces SPBM** tab.

Name	Description
Index	Specifies an Index value for the SPBM interface.
SpbmId	Specifies the SPBM ID.
State	Specifies whether the SPBM interface is enabled or disabled.
Type	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT: ptpt or bcast. Only the point-to-point (ptpt) interface type is supported.
L1Metric	Configures the IS-IS Interface level 1 metric on the specified port or MLT. The default value is 10. You can also configure the L1 metric as auto so the network route is determined by summing the lowest value metrics. The metrics are inversely proportional to port speed.
Origin	Specifies the source of the SPBM instance configuration, either manually configured through CLI or EDM, or dynamically configured through Auto-sense.

Configure SPBM on an Interface

Use the following procedure to configure SPBM on an interface.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select **SPBM**.
6. Select the **Interfaces SPBM** tab.
7. Select **Insert**.
8. Enter the SPBM ID.
9. Select **Insert**.

Interfaces SPBM Field Descriptions

Use the data in the following table to use the **Interfaces SPBM** tab.

Name	Description
Index	Specifies an Index value for the SPBM interface.
SpbmId	Specifies the SPBM ID.
State	Specifies whether the SPBM interface is enabled or disabled.
Type	Configures the SPBM instance interface-type on the IS-IS interface located on the specified port or MLT: ptpt or bcast. Only the point-to-point (ptpt) interface type is supported.
L1Metric	Configures the IS-IS Interface level 1 metric on the specified port or MLT. The default value for L1 metric is 10 for any link, despite the port speed. Configure to auto to use the inverse of the port speed as the metric.
Origin	Specifies the source of the SPBM instance configuration, either manually configured through CLI or EDM, or dynamically configured through Auto-sense.

View the IP Unicast FIB

About This Task

Use the following procedure to display the IP unicast Forwarding Information Base (FIB). The tab shows IP routes from remote Backbone Edge Bridges (BEBs)

The **IP Unicast FIB** tab displays all of the IS-IS routes in the IS-IS LSDB. The Preference column in the **IP Unicast FIB** tab displays the IP route preference.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **IP Unicast FIB** tab.

IP Unicast FIB Field Descriptions

Use the data in the following table to use the **IP Unicast FIB** tab.

Name	Description
VrfId	Specifies the VRF ID of the IP unicast FIB entry, 0 indicates NRE.
DestinationIpAddrType	Specifies the address type of the destination IP address.

Name	Description
DestinationIpAddr	Specifies the destination IP Address of the IP unicast FIB entry.
DestinationMask	Specifies the destination IP mask of the IP unicast FIB entry
NextHopBmac	Specifies the nexthop BMAC of the IP unicast FIB entry.
DestId	Specifies the destination I-SID of the IP unicast FIB entry.
Vlan	Specifies the VLAN of the IP unicast FIB entry.
Isid	Specifies the I-SID of the IP unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the IP unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the IP unicast FIB entry.
PrefixCost	Specifies the prefix cost of the IP unicast FIB entry.
SpbmCost	Specifies the BMAC cost of the IP unicast FIB entry.
Preference	Specifies the IP Route preference of the IP unicast FIB entry
MetricType	Specifies the IP Metric Type of the IP unicast FIB entry.

View the IPv6 Unicast FIB

About This Task

Use the following procedure to display the IPv6 unicast Forwarding Information Base (FIB). The tab shows IPv6 routes from remote Backbone Edge Bridges (BEBs)

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **IPv6 Unicast FIB** tab.

IPv6 Unicast FIB Field Descriptions

Use the data in the following table to use the **IPv6 Unicast FIB** tab.

Name	Description
VrfId	Specifies the VRF ID of the IPv6 unicast FIB entry, 0 indicates NRE.
DestinationIpAddrType	Specifies the address type of the destination IPv6 address.
DestinationIpAddr	Specifies the destination IPv6 Address of the IPv6 unicast FIB entry.

Name	Description
DestinationMask	Specifies the destination IPv6 mask of the IPv6 unicast FIB entry.
NextHopBmac	Specifies the nexthop BMAC of the IPv6 unicast FIB entry.
DestIsid	Specifies the destination I-SID of the IPv6 unicast FIB entry.
Vlan	Specifies the VLAN of the IPv6 unicast FIB entry.
Isid	Specifies the I-SID of the IPv6 unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the IPv6 unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the IPv6 unicast FIB entry.
PrefixCost	Specifies the prefix cost of the IPv6 unicast FIB entry.
SpbmCost	Specifies the BMAC cost of the IPv6 unicast FIB entry.
MetricType	Specifies the Metric Type of the IPv6 unicast FIB entry.

View the Unicast FIB

About This Task

Use the following procedure to display the unicast FIB.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **Unicast FIB** tab.

Unicast FIB Field Descriptions

Use the data in the following table to use the **Unicast FIB** tab.

Name	Description
SysId	Specifies the system ID of the node where the unicast FIB entry originated.
Vlan	Specifies the VLAN of the unicast FIB entry.
DestinationMacAddr	Specifies the destination MAC Address of the unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the unicast FIB entry.
HostName	Specifies the host name of the node where unicast FIB entry originated.
Cost	Specifies the cost of the unicast FIB entry.

View LSP Summary Information

Use the following procedure to display link-state packet (LSP) summary information. Link State Packets (LSP) contain information about the state of adjacencies or defined and distributed static routes. Intermediate System to Intermediate System (IS-IS) exchanges this information with neighboring IS-IS routers at periodic intervals.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **LSP Summary** tab.

LSP Summary Field Descriptions

Use the data in the following table to use the **LSP Summary** tab.

Name	Description
Level	Specifies the level at which the system displays this LSP.
ID	Specifies the 8 byte LSP ID, consisting of the SystemID, Circuit ID, and Fragment Number.
Seq	Specifies the sequence number for this LSP.
Checksum	Specifies the 16 bit Fletcher Checksum for this LSP.
LifetimeRemain	The remaining lifetime in seconds for this LSP.
HostName	The hostname listed in LSP, or the IS-IS system name. The IS-IS system name uses global system name, if host name is not configured.

View IS-IS Adjacencies

Use the following procedure to display IS-IS adjacency information. The platform sends IS-IS Hello (IIH) packets to discover IS-IS neighbors and establish and maintain IS-IS adjacencies. The platform continues to send IIH packets to maintain the established adjacencies. For two nodes to form an adjacency the B-VLAN pairs for the primary B-VLAN and secondary B-VLAN must match.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Adjacency** tab.

Adjacency Field Descriptions

Use the data in the following table to use the **Adjacency** tab.

Name	Description
Interface	Specifies the IS-IS interface on which the adjacency is found.
Level	Indicates the level of the IS-IS interface (Level 1 [default] or Level 2).
State	Specifies the state of the adjacency: <ul style="list-style-type: none"> • down • initializing • up • failed
LastUpTime	Indicates when the adjacency most recently entered the state up , measured in hundredths of a second since the last re-initialization of the network management subsystem. Displays 0 if the adjacency has never been in state up .
NeighPriority	Specifies the priority of the neighboring Intermediate System for becoming the Designated Intermediate System.
HoldTimer	Specifies the holding time in seconds for this adjacency. This value is based on received IS-IS Hello (IIH) PDUs and the elapsed time since receipt.
NeighSysID	Specifies the system ID of the neighboring Intermediate System.
AdjHostName	Specifies the host name listed in the LSP, or the IS-IS system name. The IS-IS system name uses global system name, if host name is not configured.
ParallelActive	Specifies if the current adjacency among all the parallel adjacencies between two nodes is active. This field also indicates if the adjacency is a backup.
PortMembers	Displays the members of an MLT.

Configure IS-IS Global Parameters

Use the following procedure to configure IS-IS global parameters. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Globals** tab.
4. Configure the global IS-IS parameters.

5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminState	Specifies the global status of IS-IS on the switch: on or off. The default is off .
LevelType	Sets the router type globally: <ul style="list-style-type: none"> level1: Level-1 router type level1and2: Level-1/2 router type is not supported. The default value is level1.
SystemId	Specifies the IS-IS system ID for the switch. Valid value is a 6-byte value in the format <xxxx.xxxx.xxxx>. <p>Important: After you have configured the SPBM nickname and enabled IS-IS, you must change the nickname when you change the system ID. If you want to maintain the same nickname with a different system ID, see Job aid on page 1266.</p>
MaxLspGenInt	Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for RxmtLsplnt. <p>The default value is 900 seconds.</p>
Csnplnt	Specifies the Complete Sequence Number Packet (CSNP) interval in seconds. This is a system level parameter that applies for L1 CSNP generation on all interfaces. <p>The default value is 10.</p>
RxmtLsplnt	Specifies the minimum time between retransmission of an LSP, in seconds. This defines how fast the switch resends the same LSP. This is a system level parameter that applies for L1 retransmission of LSPs. <p>The default value is 5 seconds.</p>
PSNPInterval	Specifies the Partial Sequence Number Packet (PSNP) interval in seconds. This is a system level parameter that applies for L1 PSNP generation on all interfaces. <p>The default value is 2.</p>

Name	Description
SpfDelay	Specifies the SPF delay in milliseconds. This value is used to pace successive SPF runs. The timer prevents two SPF runs from being scheduled very closely. The default value is 100 milliseconds.
HostName	Specifies the IS-IS system name. This can be used as the host name for dynamic host name exchange in accordance with RFC 2763. By default, the system uses the global system name as the IS-IS system name, until you manually configure it. Also, the system does not save the global system name in the running configuration on the switch.
IpSourceAddress	Specifies IP source address for SPBM IP shortcuts.
Ipv6SourceAddress	Specifies IPv6 source address for SPBM IP shortcuts.
IpTunnelSourceAddress	Specifies the IS-IS IPv4 tunnel source address.
IpTunnelVrf	Specifies the VRF name associated with the IP tunnel.
IpTunnelOverlay	Permits the configuration of the tunnel source address even though it belongs to a VRF with an attached I-SID. The default is disabled.
DesignatedBNHostName Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the hostname of the designated boundary node.
DesignatedBNHomeSystemId Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the system ID of the designated boundary node in the home area.
DesignatedBNRemoteSystemId Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the system ID of the designated boundary node in the remote area.
MgmtCliIpAddr	Specifies the DvR management IP address for this node, in the DvR domain.

Name	Description
BackboneEnable	Select to enable this node to join the DvR backbone so that it can receive redistributed DvR host routes from all DvR Controllers in the network.
FanMember	Specifies whether the node is a member of the Fabric Area Network (FAN) .
DynamicallyLearnedArea	For FAN members, specifies the IS-IS area that is dynamically learned from the neighbor's Hello PDU if the node does not have the IS-IS manual area configured.
MAOperState Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the Multi-area SPB operational state. The default is disabled.
AreaName	Configures the system area name.
HelloPadding	Configures IS-IS hello padding on all IS-IS network-to-network interface (NNI) links. IS-IS hello padding is enabled by default. IS-IS logical-interfaces with a configured IS-IS MTU do not perform IS-IS Hello Padding, regardless of the global IS-IS Hello Padding configuration.
MAHomeAlwaysUp Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Forwards traffic from UNI ports on the Multi-area SPB designated boundary node into the remote-area, and receive traffic from the remote area that is destined for the UNIs ports, without requiring an established IS-IS adjacency in the home area. Important: To avoid packet loops back into the home area, this feature does not operate on non-designated boundary nodes. The default is disabled.

Configure System-Level IS-IS Parameters

Use the following procedure to configure system-level IS-IS parameters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric > IS-IS**.
2. Select the **System Level** tab.
3. Configure the IS-IS system level parameters.
4. Select **Apply**.

System Level Field Descriptions

Use the data in the following table to use the **System Level** tab.

Name	Description
Index	Specifies the level: I1 or I2. Only I1 is supported.
State	Specifies the state of the database at this level. The value 'off' indicates that IS-IS is not active at this level. The value 'on' indicates that IS-IS is active at this level, and not overloaded. The value 'waiting' indicates a database that is low on an essential resources, such as memory. The administrator may force the state to 'overloaded' by setting the object SetOverload . If the state is 'waiting' or 'overloaded', you originate LSPs with the Overload bit set.
SetOverload	Sets or clears the overload condition. The possible values are true or false. The default value is false.
SetOverloadUntil	Sets the IS-IS overload-on-startup value in seconds. The overload-on-startup value is used as a timer to control when to send out LSPs with the overload bit cleared after IS-IS startup. Note: If you configure SetOverloadUntil to a number other than zero, then the overload bit is set at this level when the AdminState variable goes to the state 'on' for this Intermediate System. After the SetOverloadUntil seconds elapse, the overload flag remains set if the implementation runs out of memory or if you configured it manually using SetOverload to true. If SetOverload is false, the system clears the overload bit after SetOverloadUntil seconds elapse, if the system has not run out of memory. The default value is 20.
MetricStyle	Specifies the IS-IS metric type. Available values are narrow, wide or both. Only wide is supported.

View IS-IS System Statistics

Use the following procedure to view the Intermediate-System-to-Intermediate-System (IS-IS) system statistics.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Stats**.
3. Select the **System Stats** tab.

System Stats Field Descriptions

Use the data in the following table to use the **System Stats** tab.

Name	Description
CorrLSPs	Indicates the number of corrupted in-memory link-state packets (LSPs) detected. LSPs received from the wire with a bad checksum are silently dropped and not counted.
AuthFails	Indicates the number of authentication key failures recognized by this Intermediate System.
LSPDbaseOloads	Indicates the number of times the LSP database has become overloaded.
ManAddrDropFromAreas	Indicates the number of times a manual address has been dropped from the area.
AttmptToExMaxSeqNums	Indicates the number of times the IS has attempted to exceed the maximum sequence number.
SeqNumSkips	Indicates the number of times a sequence number skip has occurred.
OwnLSPPurges	Indicates the number of times a zero-aged copy of the system's own LSP is received from some other node.
IDFieldLenMismatches	Indicates the number of times a PDU is received with a different value for ID field length to that of the receiving system.
PartChanges	Indicates partition changes.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.

Configure IS-IS Interfaces

Use the following procedure to configure the IS-IS interfaces. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Interfaces** tab.
4. Configure the IS-IS interface parameters.

5. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Index	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
IfIndex	Specifies the interface on which the circuit is configured (port or MLT).
Type	Specifies the IS-IS circuit type. Only the point-to-point (PtToPt) interface type is supported.
AdminState	Specifies the administrative state of the circuit: on or off.
OperState	Specifies the operational state of the circuit.

Name	Description
AuthType	<p>Specifies the authentication type:</p> <ul style="list-style-type: none"> • none • simple: If selected, you must also specify a key value but the key id is optional. Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5: If selected, you must also specify a key value, but the key-id is optional. MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. There is an optional key ID. • hmac-sha-256: If selected, you must also specify a key value, but the key-id is optional. With SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. If the digests match, the packet is accepted. If the digests do not match, the receiving switch discards the packet. There is an optional key ID. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default is none.</p>
AuthKey	Specifies the authentication key.
KeyId	Specifies the authentication key ID.

Name	Description
LevelType	Specifies the router type globally: <ul style="list-style-type: none"> level1: Level-1 router type level 1and2: Level-1/2 router type. This type is not supported. The default value is level1.
NumAdj	Specifies the number of adjacencies on this circuit.
NumUpAdj	Specifies the number of adjacencies that are up.
AutoNniEnable	Enable to have the node create an IS-IS interface, attach the interface to an SPBM instance, and then enable IS-IS on the port interface. This field displays on the Insert Interfaces dialog box and applies to port interfaces only.
Origin	Specifies the origin of the IS-IS circuit configuration on the port, either manually configured through CLI or EDM or dynamically configured through Auto-sense.
PortMembers	Displays the members of an MLT.

Configure IS-IS Interface Level Parameters

Use the following procedure to configure IS-IS interface level parameters. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Interfaces Level** tab.
4. Configure the IS-IS interface level parameters.
5. Select **Apply**.

Interfaces Level Field Descriptions

Use the data in the following table to use the **Interfaces Level** tab.

Name	Description
Index	The identifier of this circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
Level	Specifies the router type globally: <ul style="list-style-type: none"> • l1: Level1 router type • l12: Level1/Level2 router type. This type is not supported. The default value is l1.
ISPriority	Specifies an integer sub-range for IS-IS priority. The default is 64.
HelloTimer	Configures the level 1 hello interval. Specifies the maximum period, in milliseconds, between IS-IS Hello Packets (IIH) PDUs on multiaccess networks at this level for LANs. The value at Level1 is used as the period between Hellos on Level1/Level2 point to point circuits. Setting this value at Level 2 on an Level1/Level2 point-to-point circuit results in an error of InconsistentValue. The default value is 9000 milliseconds or 9 seconds.
HelloMultiplier	Configures the level 1 hello multiplier. The default value is 3 seconds.
DRHelloTimer	Indicates the period, in milliseconds, between Hello PDUs on multiaccess networks when this Intermediate System is the Designated Intermediate System. The default is 3.

View IS-IS Interface Counters

Use the following procedure to view the Intermediate-System-Intermediate-System (IS-IS) interface counters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Stats**.
3. Select the **Interface Counters** tab.

Interface Counters Field Descriptions

Use the data in the following table to use the **Interface Counters** tab.

Name	Description
Index	Shows a unique value identifying the IS-IS interface.
Level	Shows the type of circuit that discovered the interface counters. The point to point Hello PDU includes both Layer 1 and Layer 2, and IS from a single adjacency on point to point links, therefore combining counts on point to point links into one group.
AdjChanges	Shows the number of times an adjacency state change has occurred on this circuit.
InitFails	Shows the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures. Failures to form an adjacency are counted by isisCircRejAdjs.
RejAdjs	Shows the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Shows the number of times an IS-IS control PDU with an ID field length different to that for this system has been received.
MaxAreaAddrMismatches	Shows the number of times an IS-IS control PDU with a max area address field different to that for this system has been received.
AuthFails	Shows the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Shows the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.

View IS-IS interface control packets

Use the following procedure to view the Intermediate-System-to-Intermediate-System (IS-IS) interface control packets.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Stats**.
3. Select the **Interface Control Packets** tab.

Interface Control Packets Field Descriptions

Use the data in the following table to use the **Interface Control Packets** tab.

Name	Description
Index	Shows a unique value identifying the Intermediate-System-to-Intermediate-System (IS-IS) interface.
Direction	Indicates whether the switch is sending or receiving the PDUs.

Name	Description
Hello	Indicates the number of IS-IS Hello frames seen in this direction at this level.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packets (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNP) frames seen in this direction at this level.

View Statistical Graph of IS-IS Interface Counters

Use the following procedure to view statistical graph of the IS-IS interface counters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select **Graph**.

Interface Counters Field Descriptions

Use the data in the following table to use the **Interface Counters** tab.

Name	Description
InitFails	Indicates the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures.
RejAdjs	Indicates the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Indicates the number of times an Intermediate-System-to-Intermediate-System (IS-IS) control PDU with an ID field length different from that for this system has been received.
MaxAreaAddrMismatches	Indicates the number of times an IS-IS control PDU with a max area address field different from that for this system has been received.
AuthFails	Indicates the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Indicates the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.
AbsoluteValue	Displays the counter value.

Name	Description
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

View Statistical Graph of IS-IS Interface Sending Control Packet

Use the following procedure to view the statistical graph of the IS-IS interface sending control packet.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select **Graph**.
6. Select the **Interface Sending Control Packets** tab.

Interface Sending Control Packets Field Descriptions

Use the data in the following table to use the **Interface Sending Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello (IIH) PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on Layer 1 or L1L2 circuits, and at Layer 2 otherwise.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

View Statistical Graph of IS-IS Interface Receiving Control Packets

Use the following procedure to view statistical graph of the IS-IS interface receiving control packets.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select the **Graph** button.
6. Select the **Interface Receiving Control Packets** tab.

Interface Receiving Control Packets Field Descriptions

Use the data in the following table to use the **Interface Receiving Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on Layer 1 or L1L2 circuits, and at Layer 2 otherwise.
LSP	Indicates the number of IS-IS link-state packet (LSP) frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Configure an IS-IS Manual Area

Use the following procedure to configure an IS-IS manual area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Manual Area** tab.
4. Select **Insert**.

5. In **AreaAddr**, type the area address.
6. Select **Insert**.

Manual Area or Manual Area Remote Field Descriptions

Use the data in the following table to use the **Manual Area** or **Manual Area Remote** tab.

Name	Description
AreaAddr	Specifies the IS-IS manual area. Valid value is 1-13 bytes in the format <xx.xxxx.xxxx...xxx>. Only one manual area is supported. Use the same manual area across the entire SPBM cloud. For IS-IS to operate, you must configure at least one manual area.

Configure Dynamic Nickname Assignment

About This Task

Use this procedure to enable Dynamic Nickname Assignment and specify a nickname allocation range.



Note

You must disable Dynamic Nickname Assignment before you can change the nickname allocation range.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **Globals** tab.
4. To enable the nickname server, select **enable** for **NicknameServerEnable**.
5. In **NicknameServerPrefix**, type a prefix.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
GlobalEnable	Enables or disables SPBM globally. The default is disabled. To ensure proper cleanup of MAC tables after you disable SPBM, save the configuration, and then reboot the switch.
GlobalEtherType	Specifies the global ethertype value as 0x8100 or 0x88a8. The default value is 0x8100.
NicknameServerEnable	Enables or disables the nickname server. The default is disabled.

Name	Description
NicknameDynamicAllocationStatus	Displays the Dynamic Nickname Allocation service operational status.
NicknameServerPrefix	Specifies the nickname server allocation prefix. x.xx.xx uses the form X.X0.00 from 0.00.00 to F.F0.00. A group, X.X0.00 to X.XF.FF, can provide up to 4,096 nicknames. The default is A.00.00.

SPBM Configuration Examples


This section provides configuration examples to configure basic SPBM and IS-IS infrastructure.

Examples and network illustrations in this document may not include your specific platform. Unless otherwise noted, the concept illustrated applies to all supported platforms. For information about high-level feature support, see *Fabric Engine and VOSS Feature Matrix*.

Basic SPBM configuration example

The following figure shows a sample greenfield deployment for SPBM.

Figure 63: Greenfield SPBM deployment



Note

For migration purposes, SPBM can coexist with existing SMLT configurations.

Ethernet and MLT configuration

The following sections show the steps required to configure the Ethernet and MLT interfaces in this example.

SwitchC

```
PORT CONFIGURATION - PHASE 1

interface GigabitEthernet 1/30
encapsulation dot1q
exit
```

SwitchG

```

PORT CONFIGURATION - PHASE 1

interface GigabitEthernet 1/5
encapsulation dot1q
exit

MLT CONFIGURATION

mlt 1 enable
mlt 1 member 1/21-1/22
mlt 1 encapsulation dot1q

```

SwitchD

```

MLT CONFIGURATION

mlt 1 enable
mlt 1 member 1/20,1/30
mlt 1 encapsulation dot1q

```

IS-IS SPBM Global Configuration

The following figure shows the IS-IS area information added to the network.




Figure 64: IS-IS SPBM global

The following sections show the steps required to configure the global IS-IS SPBM parameters in this example.

SwitchC

```

enable
configure terminal
prompt SwitchC

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.13
spbm 1 b-vid 20

```

```
ISIS CONFIGURATION

is-type 11
manual-area 30.0000
sys-name SwitchC
exit
router isis enable

VLAN CONFIGURATION

vlan create 20 name "B-VLAN" type spbm-bvlan
```

SwitchG

```
enable
configure terminal
prompt SwitchG

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.10
spbm 1 b-vid 20

ISIS CONFIGURATION

is-type 11
manual-area 30.0000
sys-name SwitchG
exit
router isis enable

VLAN CONFIGURATION

vlan create 20 name "B-VLAN" type spbm-bvlan
```

SwitchD

```
enable
configure terminal
prompt SwitchD

BOOT CONFIGURATION

spbm
spbm ethertype 0x8100

ISIS SPBM CONFIGURATION

router isis
spbm 1
spbm 1 nick-name f.30.14
spbm 1 b-vid 20

ISIS CONFIGURATION
```

```

is-type l1
manual-area 30.0000
sys-name SwitchD
exit
router isis enable

VLAN CONFIGURATION


vlan create 20 name "B-VLAN" type spbm-bvlan

```

IS-IS SPBM Interface Configuration

The following figure shows the IS-IS area information and interfaces in the network.

Figure 65: IS-IS SPBM interface



The following sections show the steps required to configure the IS-IS SPBM interfaces in this example.

SwitchC

```

PORT CONFIGURATION - PHASE II

interface GigabitEthernet 1/30
isis
isis spbm 1
isis enable
exit

```

SwitchG

```

PORT CONFIGURATION - PHASE II

interface GigabitEthernet 1/5
isis
isis spbm 1
isis enable
exit

MLT INTERFACE CONFIGURATION

interface mlt 1
isis
isis spbm 1
isis enable
exit

```

SwitchD**MLT INTERFACE CONFIGURATION**

```

interface mlt 1
isis
isis spbm 1
isis enable
exit

```

Verify SPBM Operations

The following sections show the output from verifying the sample IS-IS SPBM configuration.

Checking Operation — SwitchC

```
SwitchC:1# show isis interface
```

```

=====
ISIS Interfaces
=====
IFIDX   TYPE   LEVEL  OP-STATE  ADM-STATE  ADJ  UP-ADJ  SPBM-L1  OP-SPBM-  ORIGIN  AREA  AREA-NAME
-METRIC  L1-METRIC
-----
Port1/30 pt-pt  Level 1   UP        UP          1    1      10      (A) 10    CONFIG  HOME  area-9.00.02
Home: 1 out of 1 Total Num of ISIS Interfaces
=====

```

```

Legend:
(A): l1 metric is automatically updated based on detected interface speed.

```

```
SwitchC:1# show isis adjacencies
```

```

=====
ISIS Adjacencies
=====
INTERFACE L STATE  UPTIME    PRI  HOLDDTIME  SYSID          HOST-NAME  STATUS  AREA  AREA-NAME
-----
Port1/30  1  UP    1d 19:19:52  22    beb0.0000.7204  SwitchC    ACTIVE  HOME  area-9.00.02
Home: 1 out of 1 interfaces have formed an adjacency
Remote: 0 out of 0 interfaces have formed an adjacency
=====

```

```
SwitchC:1# show isis spbm unicast-fib
```

```

=====
SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION          BVLAN  SYSID          HOST-NAME  OUTGOING  COST  AREA  AREA-NAME
ADDRESS              INTERFACE
-----
00:0e:62:25:a3:df    4000   0016.ca23.73df  SwitchG    1/30      10    HOME  area-9.00.02
00:14:0d:a0:13:df    4000   0014.0da0.13df  SwitchD    1/30      10    HOME  area-9.00.02
=====

```

```

Home: Total number of SPBM UNICAST FIB entries 2
Remote: Total number of SPBM UNICAST FIB entries 0
=====

```

```
SwitchC:1# show isis spbm unicast-tree 4000
```

```

Node:000e.6225.a3df.00 (SwitchG) -> ROOT
Node:0014.0da0.13df.00 (SwitchD) -> Node:000e.6225.a3df.00 (SwitchG) -> ROOT

```

Checking Operation — SwitchG

```
SwitchG:1# show isis interface verbose
```

```

=====
ISIS Interfaces
=====

```

```

INTERFACE TYPE  LEVEL  OPER   ADM   ADJ   UP   ADM-SPBM  OPER-SPBM  ORIGIN  AREA  AREA-NAME  PORT
                STATE  STATE  ADJ  ADJ  L1-METRIC L1-METRIC  L1-METRIC  AREA  AREA-NAME  MEMBERS
-----
Port1/5 pt-pt  Level 1  UP    UP    1    1    10      (A) 10      CONFIG HOME  area-9.00.02
Mtl1  pt-pt  Level 1  UP    UP    1    1    10      (A) 10      CONFIG HOME  area-9.00.02  1/4-1/5
-----
Home: 2 out of 2 Total Num of ISIS Interfaces
-----
Legend:
(A): l1 metric is automatically updated based on detected interface speed.

SwitchG:1# show isis adjacencies
=====
ISIS Adjacencies
=====
INTERFACE L STATE  UPTIME      PRI  HOLDDTIME  SYSID          HOST-NAME  STATUS  AREA  AREA-NAME
-----
Port1/30  1  UP    1d 19:19:52  127  26          beb0.0000.7204  SwitchC   ACTIVE  HOME  area-9.00.02
Mtl1      1  UP    04:57:34    127  20          0014.0da0.13df SwitchD   ACTIVE  HOME  area-9.00.02
-----
Home: 2 out of 2 interfaces have formed an adjacency
Remote: 0 out of 0 interfaces have formed an adjacency
-----

SwitchG:1# show isis spbm unicast-fib
=====
SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION      BVLAN SYSID          HOST-NAME  OUTGOING  COST  AREA  AREA-NAME
ADDRESS
-----
00:14:0d:a0:13:df 4000 0014.0da0.13df  SwitchD   MLT-1     10    HOME  area-9.00.02
00:15:e8:9f:e3:df 4000 0015.e89f.e3df  SwitchC   1/5       10    HOME  area-9.00.02
-----
Home: Total number of SPBM UNICAST FIB entries 2
Remote: Total number of SPBM UNICAST FIB entries 0
-----

SwitchG:1# show isis spbm unicast-tree 4000
Node:0015.e89f.e3df.00 (SwitchC) -> ROOT
Node:0014.0da0.13df.00 (SwitchD) -> ROOT

```

Checking Operation — SwitchD

```

SwitchD:1# show isis interface verbose
=====
ISIS Interfaces
=====
INTERFACE TYPE  LEVEL  OPER   ADM   ADJ   UP   ADM-SPBM  OPER-SPBM  ORIGIN  AREA  AREA-NAME  PORT
                STATE  STATE  ADJ  ADJ  L1-METRIC L1-METRIC  L1-METRIC  AREA  AREA-NAME  MEMBERS
-----
Mtl1  pt-pt  Level 1  UP    UP    1    1    10      (A) 10      CONFIG HOME  area-9.00.02  1/1-1/2
-----
Home: 1 out of 1 Total Num of ISIS Interfaces
-----
Legend:
(A): l1 metric is automatically updated based on detected interface speed.

SwitchD:1# show isis adjacencies
=====
ISIS Adjacencies
=====
INTERFACE L STATE  UPTIME      PRI  HOLDDTIME  SYSID          HOST-NAME  STATUS  AREA  AREA-NAME
-----
Mtl1      1  UP    05:03:59    127  21          000e.6225.a3df SwitchG   ACTIVE  HOME  area-9.00.02
-----
Home: 1 out of 1 interfaces have formed an adjacency
Remote: 0 out of 0 interfaces have formed an adjacency
-----

SwitchD:1# show isis spbm unicast-fib
=====
SPBM UNICAST FIB ENTRY INFO
=====

```



```

=====
DESTINATION          BVLAN SYSID          HOST-NAME    OUTGOING          COST   AREA   AREA-NAME
ADDRESS
-----
00:0e:62:25:a3:df    4000  000e.6225.a3df    SwitchG          MLT-1         10     HOME   area-9.00.02
00:15:e8:9f:e3:df    4000  0015.e89f.e3df    SwitchC          MLT-1         10     HOME   area-9.00.02
-----

Home:   Total number of SPBM UNICAST FIB entries 2
Remote: Total number of SPBM UNICAST FIB entries 0
-----

SwitchD:1# show isis spbm unicast-tree 4000
Node:000e.6225.a3df.00 (SwitchG) -> ROOT
Node:0015.e89f.e3df.00 (SwitchC) -> Node:000e.6225.a3df.00 (SwitchG) -> ROOT

```

Fabric Extend

Table 125: Fabric Extend product support

Feature	Product	Release introduced
Fabric Extend	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 125: Fabric Extend product support (continued)

Feature	Product	Release introduced
Fabric Extend over IPsec	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 8.7 Supported on 5720-24MXW and 5720-48MXW Supported using Fabric IPsec Gateway
	7520 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7720 Series	Fabric Engine 8.10 Supported using Fabric IPsec Gateway
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3 Supported on VSP4900-12MXU-12XE and VSP4900-24XE Supported using Fabric IPsec Gateway
	VSP 7400 Series	VOSS 8.2 Supported using Fabric IPsec Gateway

You can use these switches in a main office of a hub and spoke deployment or to connect one Data Center to another Data Center.

Fabric Extend enables Enterprises to extend the Fabric Connect technology over Layer 2 or Layer 3 core networks. The *logical IS-IS interface* is the mechanism that enables Fabric Extend to connect SPB fabric nodes. When connecting to a Layer 3 core network, logical IS-IS interfaces create virtual tunnels and encapsulate SPB traffic by adding a VXLAN header to SPB packets.

The following figure illustrates two Fabric Connect “islands” separated by a third-party core IP network. The IP network could be third-party equipment in an enterprise or a service provider’s infrastructure such as an MPLS VPN service.




Figure 66: Fabric Connect networks connected by an IP network

The following figure illustrates how Fabric Extend enables you to connect the fabric islands to create ONE Fabric Connect network. This figure shows a Layer 3 core network where Fabric Extend uses IP tunneling by adding a VXLAN header to the SPBM packets. This can be over a third party IPv4 transport network such as MPLS IP-VPN or in a Campus IP backbones.




Figure 67: Single Fabric Connect Domain realized using Fabric Extend

The following figure shows a Layer 2 core network where Fabric Extend can transport SPBM packets over a Layer 2 MPLS VPLS or PBB E-LINE service by creating layer 3 tunnels over a Layer 2 third party network.




Figure 68: Fabric Extend over VLAN tunnels

Advantages of Fabric Extend

Fabric Connect is an Ethernet-based, industry-standard (IEEE 802.1aq) networking virtualization solution. With Fabric Connect, you can have thousands of virtualized service instances at any point in the network. Other Fabric Connect advantages include rapid time to service, Layer 2 and Layer 3 Unicast and IP Multicast virtualization, and scalable IP multicast. But the most significant advantage of Fabric Connect is that you provision services at the network edge only, not the core.

The Fabric Extend feature enables you to extend the Fabric Connect model over Layer 2 and Layer 3 core networks. The interconnection of Fabric Connect deployments can be over any IP-based network whether it's a campus backbone, Data Center, or a MAN/WAN IP MPLS network.

With Auto-sense enabled, Fabric Extend leverages LLDP to establish Fabric Extend tunnels automatically between a Fabric switch and the ExtremeCloud SD-WAN Appliance. Based on the LLDP standard, Fabric Extend information is transmitted using TLVs with LLDP Protocol Data Units (PDU) automatically.

For more information, see [Auto-sense Port States](#) on page 28.

Logical IS-IS Interface

The logical IS-IS interface is the mechanism that enables Fabric Extend to connect SPB fabric nodes.

Logical IS-IS interfaces perform the following functions depending on the type of core network:

- In a Layer 2 core network, logical IS-IS interfaces do not use VXLAN. The tunnels are point-to-point VLAN connections so there is no need to encapsulate a VXLAN header to SPB packets. The logical IS-IS interfaces translate the Backbone VLAN IDs (B-VIDs) and map them to each of the branch provider VIDs. This use case is also called Fabric Extend VID (FE-VID).
- In a Layer 3 core network, logical IS-IS interfaces create virtual IP tunnels and encapsulate SPB traffic by adding a Virtual Extensible LAN (VXLAN) header to SPB packets. This use case is also called Fabric Extend IP (FE-IP).

Fabric Extend uses virtual tunnels in Layer 3 core solutions to connect SPB fabric nodes. These nodes can stretch over IP routed campus networks, service provider Layer 2 core networks, or service provider Layer 3 core networks such as IP MPLS VPNs.



Note

You cannot use VLACP on logical IS-IS interface connections.

Layer 2 Core Network (FE-VID)

If the service provider has a Layer 2 core network, note the following points:

- The list of VLANs should have two VLANs, not more than two or less than two. You do not have to configure the VLAN IDs (VIDs) as platforms VLANs.
- Include a primary VLAN ID in the list of VLANs.

- Each logical interface must have a unique set of VIDs for each port or MLT. The same VIDs however, can be reused across a different set of ports or MLTs.
- Logical interface VIDs and B-VLANs cannot be the same.
- Configuring the same VIDs as primary and secondary is not allowed.
- The port or MLT on which the Layer 2 core IS-IS logical interface is configured cannot be part of any other user configured VLANs.
- You cannot delete an MLT that is configured as a logical interface tunnel MLT.
- A logical interface consists of a port/MLT and a list of VLANs, where port/MLT is the physical connectivity to the Layer 2 core network and VLANs are the list of VLANs used to transport/bridge IS-IS control packets and Mac-in-Mac data traffic.
- VXLAN headers are not used in Layer 2 core Fabric Extend solutions.
- IS-IS control packets are not encapsulated before they are sent over a logical interface. Instead, the VLAN in the outer Ethernet header (SPB primary bvid) is replaced by the user configured logical interface VLAN.
- Spanning tree is disabled by default on **port/MLT** on which a Layer 2 core logical IS-IS interface is configured.

Layer 3 Core Network (FE-IP)

If the service provider has a Layer 3 core network, note the following points:

- A logical IS-IS interface points to a remote BEB destination IP address.
- Port and VLAN ID are not needed to create a logical IS-IS interface, instead they can be retrieved from the next hop of destination IP address.
- IS-IS control packets (IS-IS hello, LSDB, CSNP, PSNP) are encapsulated with a VXLAN header and sent over a logical IS-IS interface.

Parallel Tunnel Support for Fabric Extend

Table 126: Fabric Extend - Parallel Tunnel product support

Feature	Product	Release introduced
Parallel Tunnel Support for Fabric Extend	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0
	5420 Series	Fabric Engine 9.0
	5520 Series	Fabric Engine 9.0
	5720 Series	Fabric Engine 9.0
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8

You can create parallel Fabric Extend IP tunnels between the same two nodes to create a backup Fabric Extend adjacency. Parallel tunnels are active simultaneously.

The switch supports the following parallel tunnel configuration:

- The primary tunnel configuration uses the global IP tunnel source address that you configure for IS-IS by using the **ip-tunnel-source-address** CLI command.
- The backup tunnel configuration uses an additional source address and optional VRF that you configure for the IS-IS logical interface by using **logical-intf isis** CLI command parameters.

**Note**

The backup tunnel can be either a native Fabric Extend tunnel or an IPsec-encrypted Fabric Extend tunnel for products that support Fabric IPsec Gateway.

IS-IS adjacencies operate the same through parallel tunnels as existing parallel IS-IS adjacency support:

- If multiple adjacencies exist between the same two nodes, only one adjacency is active and the rest are backup.
- The adjacency with the lowest metric is the active adjacency.
- If multiple interfaces use the same metric, the one with the lowest logical interface ID from the node with the lowest system ID is the active adjacency.

**Note**

These are the restrictions to configure Fabric Extend:

- Only two different Fabric Extend tunnel IP sources are supported.
- The same IP source on different VRFs is not supported.
- The same IP destination on different VRFs is not supported.

For configuration steps, see the following tasks:

- CLI configuration steps: [Configure Fabric Extend](#) on page 1315
- EDM configuration steps: [Configure Fabric Extend Logical Interfaces](#) on page 1327

IS-IS Hello Padding

Table 127: IS-IS Hello Padding

Feature	Product	Release introduced
IS-IS hello padding (ability to disable)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4.2
	5520 Series	VOSS 8.2.7
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2.7
	VSP 7400 Series	VOSS 8.2.7

To detect maximum transmission unit (MTU) mismatches, Intermediate System-to-Intermediate System (IS-IS) pads hello packets to the full interface MTU. All hello packets on Fabric Extend network-to-network interface (NNI) links are padded. On non-Fabric Extend point-to-point NNI links, hello packets are padded until a hello packet is received from the other side.

Adjusting the TCP Maximum Segment Size

Table 128: TCP Maximum Segment Size (MSS) Adjustment for Fabric Extend

Feature	Product	Release introduced
Ability to adjust the maximum segment size (MSS)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3.1
	VSP 7400 Series	VOSS 8.3.1

You can adjust the TCP maximum segment size (MSS) to improve the throughput for the TCP session over a Fabric Extend (FE) adjacency.

TCP MSS Overview

When a client initiates a connection with a server, it uses TCP SYN packets to negotiate the MSS to avoid fragmentation. The client and server use the outgoing maximum transmission unit (MTU) to advertise the MSS.

If a tunnel exists between the client and server, the encapsulation consumes more room in the outer IP header. As a result, the router that performs the tunnel encapsulation fragments the packet to fit over the tunnel. Adjust the MSS to modify the value in the TCP SYN packet so the client and server negotiate a lower number and leave headroom for tunnel encapsulation.

This adjustment functionality applies to IPv4 only.



Important

If you enable this functionality and port mirroring simultaneously, the switch does not mirror CP-generated packets.

Support

TCP MSS adjustment applies unidirectionally when a packet is forwarded from a UNI interface to any other interface. To use this functionality, you must enable TCP MSS adjustment on both sides of the FE tunnel.

For 7520 Series and 7720 Series TCP MSS adjustment applies unidirectionally when a packet is forwarded into an FE tunnel. To use this functionality, you must enable TCP MSS adjustment on both sides of the FE tunnel.

Types of Fabric Extend Deployments

As the number of Fabric Connect networks increased, the need to connect those networks became more and more desirable. Fabric Extend solves the problem of going beyond the Ethernet Fabric Connect connections to include the following IP routed wide area network (WAN) and campus solutions:

1. Fabric Extend over an MPLS IP-VPN provider WAN
2. Fabric Extend over an MPLS Virtual Private LAN Service (VPLS) or Provider Backbone Bridging (PBB) Ethernet LAN (ELAN) provider network
3. Fabric Extend over an IP campus network
4. Fabric Extend over an MPLS Pseudo-Wire or Ethernet Virtual Private Line (E-Line) provider network

Fabric Extend over an MPLS IP-VPN Provider WAN

The most common Fabric Extend deployment is a hub and spoke topology that connects the Main office over a service provider's MPLS IP VPN to multiple Branch offices. The following figure illustrates how the hub device on the main site establishes

virtual tunnels with all of the spoke devices in the same domain. In this scenario, the traffic flows are bidirectional: from hub-to-spoke and spoke-to-hub.

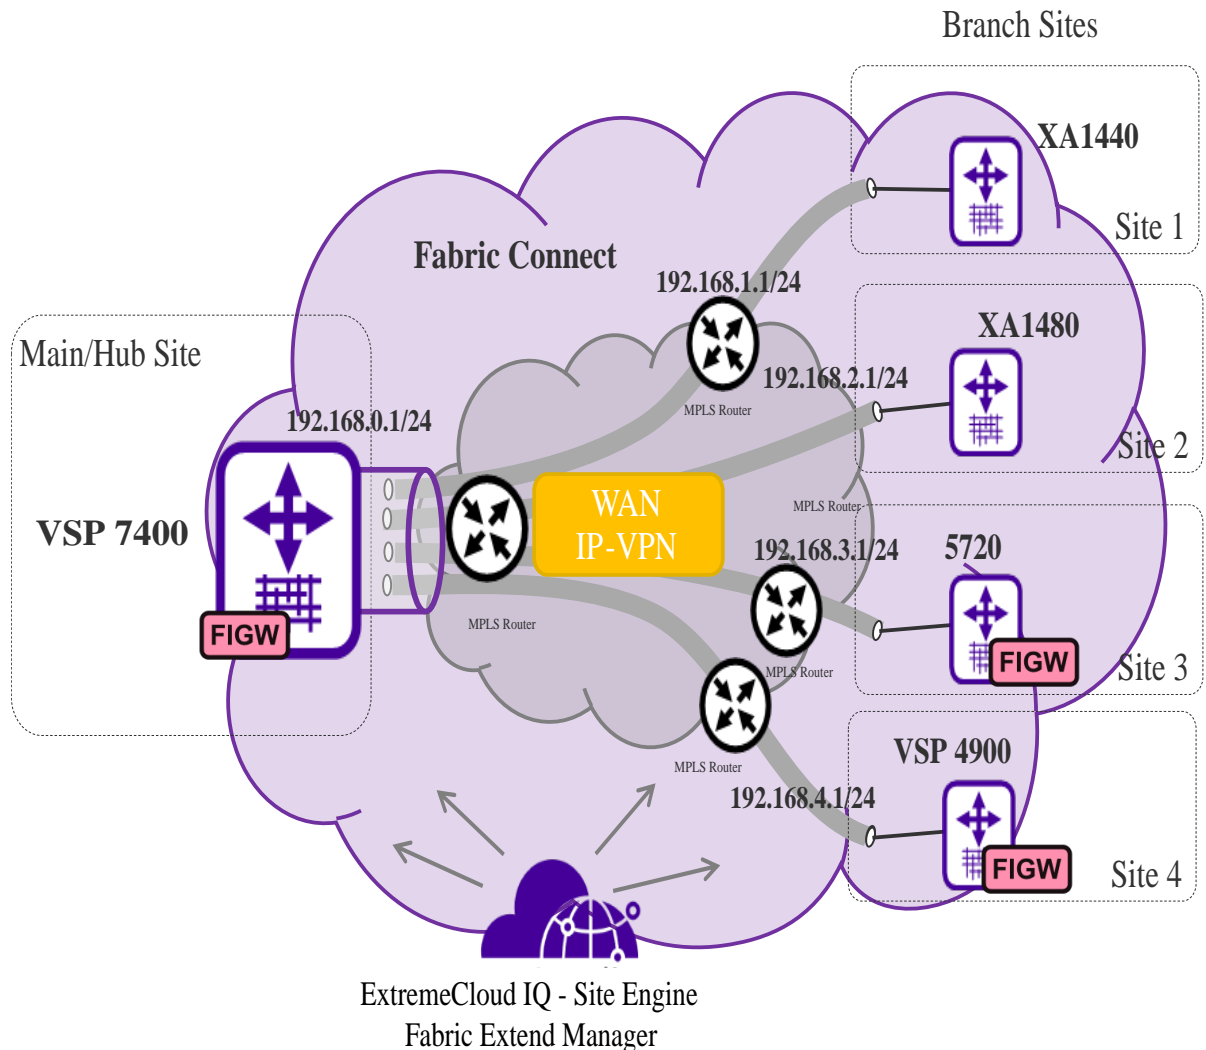


Figure 69: Fabric Extend IP VPN Deployment Option



Note

If Fabric Extend with IPsec or fragmentation and reassembly is a requirement, depending on your requirements, you can use a mix of VSP 7400 Series, VSP 4900 Series, or 5720 Series with Fabric IPsec Gateway and XA1400 Series at the main and branch sites.

Fabric Extend over an MPLS VPLS/P2P-VPLS/E-LINE/P2P-VLAN Provider Network

Where the preceding hub and spoke deployment is over a Layer 3 MPLS IP-VPN, the following VPLS deployment is over a Layer 2 segment. This type of hub and spoke deployment extends the fabric over an MPLS Virtual Private LAN Service (VPLS) or Provider Backbone Bridging (PBB) Ethernet LAN (E-LINE) network. In this scenario, the SPB nodes are connected with a point-to-point Ethernet link.

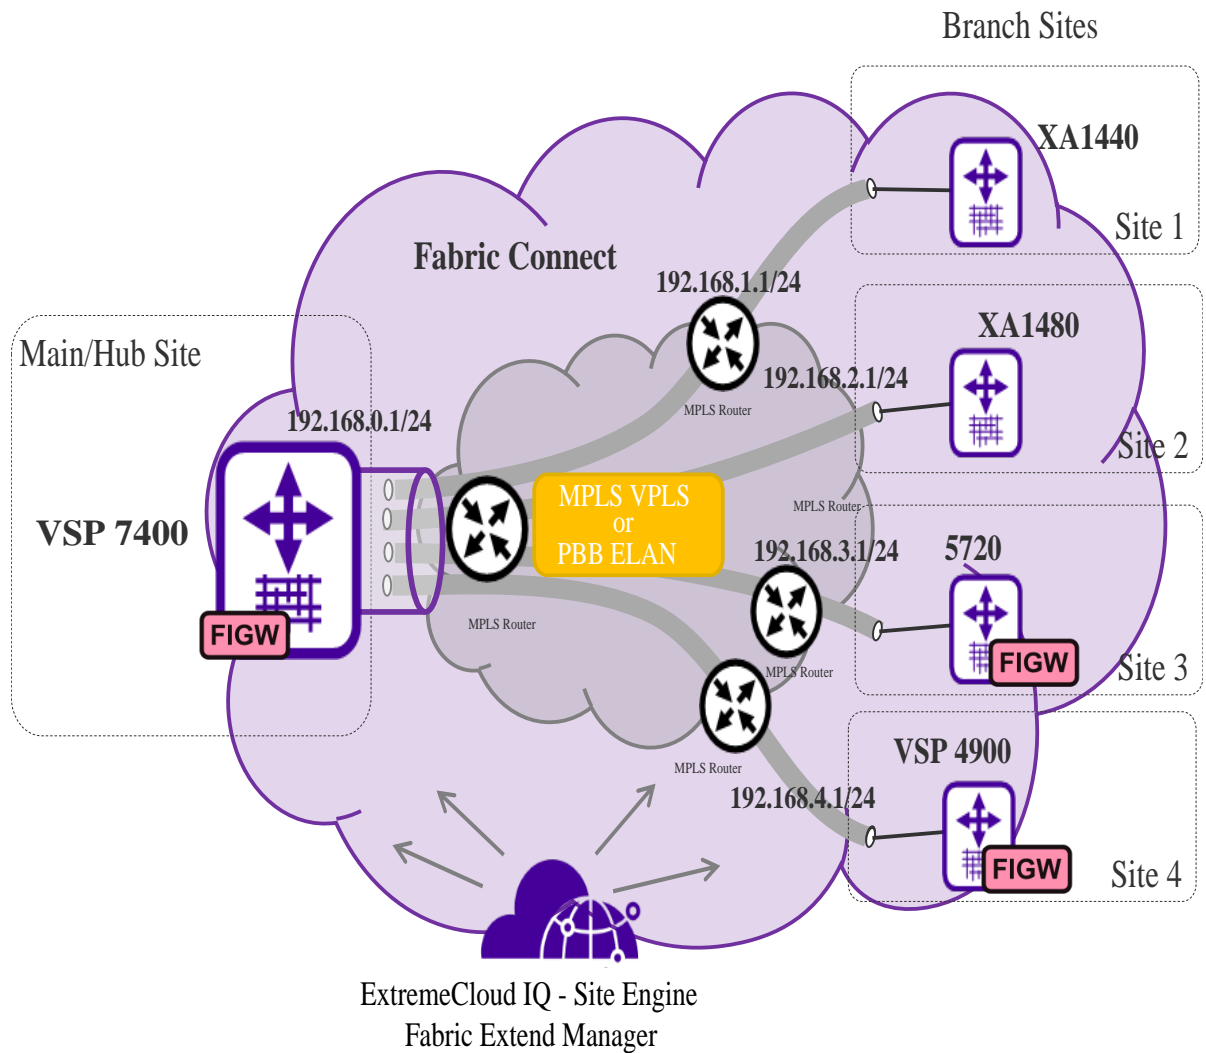


Figure 70: Fabric Extend VPLS Deployment Option

Fabric Extend over an IP Campus Network

Some customers do not want to migrate their infrastructures to SPB immediately. They want to keep their existing IP core network and deploy SPB on the edge. In this scenario, Fabric Extend supports a fabric overlay on top of the existing campus infrastructure.

The following figure illustrates how this deployment supports any-to-any traffic with full-mesh tunnels between fabric nodes. The fabric nodes serve as campus switches, support routing into the IP infrastructure, and provide an overlay fabric that enables all fabric benefits.




Figure 71: Fabric Extend Full Mesh Campus Deployment Option

Fabric Extend over an MPLS PWE3/E-Line Provider Network

The following hub and spoke deployment over an MPLS Pseudowire or Ethernet Virtual Private Line (E-Line) uses service provider VLAN tunnels. Because you can map many (VID, port/mlt list) sets to an I-SID, this gives Service Providers the flexibility to let more than one customer use the same VLAN with different I-SIDs.

The following figure illustrates how two dedicated Backbone VLAN IDs (B-VIDs) are mapped from the hub to spoke sites. Logical IS-IS interfaces translate the B-VIDs and maps them to each of the branch provider VIDs.

For a detailed configuration example showing logical interfaces using B-VID translation to two different logical VLAN IDs, see *Shortest Path Bridging (802.1aq) Technical Configuration Guide*.




Figure 72: Fabric Extend Pseudowire Deployment Option

Fabric Extend Tunnel IS-IS MTU

Table 129: Fabric Extend Tunnel IS-IS MTU Configuration product support

Feature	Product	Release introduced
Fabric Extend Tunnel IS-IS MTU configuration	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.10.1
	5420 Series	Fabric Engine 8.10.1
	5520 Series	Fabric Engine 8.10.1
	5720 Series	Fabric Engine 8.10.1
	7520 Series	Fabric Engine 8.10.1
	7720 Series	Fabric Engine 8.10.1
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.10.1
	VSP 7400 Series	VOSS 8.10.1

You can configure a unique IS-IS MTU value for each Fabric Extend (FE) tunnel. The IS-IS MTU gets applied to IS-IS packets that get sent through the interface.

You can configure each IS-IS logical interface with a unique IS-IS MTU value for each FE tunnel to improve fragmentation and reassembly in WAN connectivity. This configuration can be applied to standard WAN connections such as MPLS IP WAN as well as SD-WAN. To configure the MTU, see [Configure Fabric Extend](#) on page 1315. For more information about Fabric Extend options for SD-WAN, see [SD-WAN](#) on page 3304.



Note

- FE Tunnel IS-IS MTU is an optional configuration that applies to IS-IS packets only. The MTU for data packets is configured separately. For example, with SD-WAN deployments, the MTU for data packets is managed using the SD-WAN Appliance.
- The FE Tunnel IS-IS MTU configuration is supported only for Fabric Extend tunnels that run over a Layer 3 core network (FE-IP). This configuration is not supported for Fabric Extend tunnels over a Layer 2 core network.

IS-IS Packet Fragmentation and Reassembly

Fragmentation and reassembly of IS-IS packets is based on the IS-IS MTU value configured for the IS-IS logical interface that connects to the FE tunnel. For Layer 3 IS-IS logical interfaces, you can change the IS-IS MTU configuration at any time. The supported range, in bytes, is 750 to 1600 with a default of 1600.

For example, if you configure an FE tunnel with an IS-IS MTU of 900, and a packet size of 1950 is received on UNI with the destination on the FE tunnel, the system fragments the original 1950-sized packet into the three packets (900, 900, 150) with a packet size equal to or less than 900. The system transmits the three fragmented packets over the

IS-IS logical interface of the FE tunnel. After the fragmented packets are received at the destination, the system reassembles the original 1950-sized packet.

Note that for reassembly to succeed, the fragments must arrive at the destination switch in the correct order (based on the offset). Otherwise, the fragments are discarded.

Fabric Extend Tunnel IS-IS MTU Considerations

Consider the following interactions between route MTU and FE Tunnel IS-IS MTU configurations:

- If route MTU is not configured, the MTU value for each FE tunnel is applicable to ingress and egress traffic on the tunnel.
- If route MTU is configured, the MTU value for each FE tunnel is applicable for ingress traffic on the tunnel. The route MTU value applies to all egress traffic.



Note

System MTU maximum is a separate configuration. You can configure a system maximum MTU size of 1522, 1950, or 9600. The default value is 1950.

Fabric Extend Considerations

Review the following restrictions, limitations, and behavioral characteristics that are associated with Fabric Extend.

Tunnel Source IP

Fabric Extend supports the tunnel source IP address using a brouter port interface, a CLIP IP, or a VLAN IP.

On the 4220 Series and the 5320 Series switches that support a single active VRF, you must configure a route-map policy to suppress IS-IS redistribution of the FE tunnel subnet:

- Configure route-maps to not permit redistribution of the local route used as the tunnel source address (**ip-tunnel-source-address** command).
- Configure an accept policy to deny IS-IS routes that overlap with the destination tunnel IP address.

For more information about single active VRF switches, see [VRF Lite Configuration Rules](#) on page 4078 .

Tunnel Failover Time

With IS-IS interface default values, tunnel failure detection can take up to 27 seconds. You can reduce the IS-IS interface hello timers to speed up logical link failure detection, but be careful to avoid link flapping due to values that are too low.



Note

If the number of IS-IS interfaces on a node is greater than 100, it is a good practice to set the hello timer not lower than 5 seconds.

ACL Filters over VXLAN

IP filters configured to match IP header fields in the headers of VXLAN encapsulated packets, work only when the switch acts as a transit router and does not participate in the initiation or termination of VXLAN traffic.

VLACP

VLACP is not supported over logical IS-IS interfaces.

CFM CCM

CFM Continuity Check Messages are not supported over logical IS-IS interfaces.

CFM traceroute and tracemroute

If CFM packets transit over a layer 3 tunnel (that is the CFM packets ingress a Fabric Extend layer 3 core tunnel and egress through another layer 3 core tunnel), the transit SPBM nodes do not display as intermediate hops in the output for CFM **12 traceroute** and **12 tracemroute**.

This is because the CFM packets are encapsulated in the outer layer 3 header as part of VXLAN encapsulation, and the transit SPBM nodes cannot look into the payload of the VXLAN packet and send a copy of the CFM packet to local CPU for processing.

CFM L2 Ping

CFM Layer 2 ping to MCoSPB source mac is not supported and can fail if they are reachable through a Fabric Extend tunnel.

MACsec

Switch-based MAC Security (MACsec) encryption is Layer 2 so it cannot be used with Fabric Extend IP, which is Layer 3.

MTU Minimum in Layer 2 Pseudowire Core Networks

Service provider Layer 2 connections must be at least 1544 bytes. In this type of deployment the tunnels are point-to-point VLAN connections that do not require VXLAN encapsulation. The default MTU value is 1950.

Logical IS-IS Interfaces

Layer 2 core and Layer 3 core logical IS-IS interfaces are not supported on the same switch at the same time.

Fragmentation and Reassembly

There is no fragmentation and reassembly support in Layer 2 core solutions.

Layer 2 Logical IS-IS Interfaces

Layer 2 logical IS-IS interfaces are created using VLANs. Different Layer 2 network Service Providers can share the same VLAN as long as they use different ports or MLT IDs.

MTU Minimum in Layer 3 Core Networks

Service provider IP connections must be at least 1594 bytes to establish IS-IS adjacency over FE tunnels. The 1594 bytes includes the actual maximum frame size with MAC-in-MAC and VXLAN headers. If this required MTU size is not available, a log message reports that the IS-IS adjacency was not established. MTU cannot be auto-discovered over an IP tunnel so the tunnel MTU will not be automatically set. The default MTU value is 1950.

IP Shortcuts

The tunnel destination IP cannot be reachable through an IP Shortcuts route.



Important

If you enable IP Shortcuts and you are using the GRT as the tunnel source VRF, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you enable IP Shortcuts and you are using a VRF as the tunnel source VRF, this is not an issue.

Layer 3 over Layer 2 Limitation

- The switch requires a single next hop (default gateway) for all tunnels.
 - Over a Layer 3 core network, on a given outgoing port or MLT, there is no issue as the one router next hop can support multiple VXLAN tunnels to one or more remote sites.
 - For Layer 3 tunneling over a Layer 2 core, the switch without any specific configuration supports only one Fabric Extend tunnel to one remote site. The workaround for this single next hop issue is to create an additional VRF, VLAN, and loopback interface.
- You cannot establish a Virtual IST (vIST) session over a logical IS-IS interface. IST hellos cannot be processed or sent over a logical IS-IS interface if that is the only interface to reach BEBs in vIST pairs.

Assume that vIST is established over a regular network-to-network interface (NNI) and the NNI goes down. If the vIST pairs are reachable through a logical IS-IS interface, then the vIST session goes down in up to 240 seconds (based on the IST hold down timer). During this time, the error message is logged:

```
IST packets cannot be sent over Fabric Extend tunnels, vist session  
may go down
```



Caution

Expect traffic loss when the vIST session is down or when the error message is being logged.

Port Mirroring Resources

Port mirroring resources are limited to four ports simultaneously (where each mirroring direction counts as one). For example, if two mirroring ports are designated to mirror both ingress and egress traffic then all four mirroring ports are consumed.

Port mirroring shares these four resources with other applications such as port mirroring RSPAN, Fabric Extend, Application Telemetry, IPFIX, and ACL with mirror action. Each one of these applications consumes at least one port mirroring resource. (port mirroring RSPAN consumes two if you configure both Ingress and Egress modes.)

**Important**

To enable any one of the preceding applications, you must have at least one free mirroring resource. If all four port mirroring resources are already in use, the switch displays a `Resource not available` error message when you try to enable the application.

Fabric Extend configuration using the CLI

The following sections provide procedural information you can use to configure Fabric Extend (FE) using the Command Line Interface (CLI).

Configure Fabric Extend

Use the following procedure to configure Fabric Extend (FE). A typical Fabric Extend deployment would be between a Main office and a Branch office.

To complete the Fabric Extend tunnel, you must configure Fabric Extend at both ends of the tunnel.

Optionally, you can configure a parallel tunnel between the same two nodes to create a backup Fabric Extend adjacency.

Before You Begin

If using the tunnel originating address on a VRF, configure a CLIP and tunnel source IP address on the VRF.

If using the tunnel originating address, the logical interface address, on the GRT, Fabric Extend has the following requirements:

- The tunnel source IP address must be on the GRT, not on a VRF.

**Note**

To configure Fabric Extend using a tunnel source IP address on the GRT, use separate IP addresses for the SPBM IP Shortcuts (**ip source-address**) and the Fabric Extend **ip-tunnel-source** address. Exclude the **ip-source-address** address with an IS-IS accept policy if you want these IP addresses to be the same. You cannot use the **redistribute** command with a route-map exclusion.

Specify a CLIP interface to use as the **ip source-address** for SPBM IP shortcuts.

- If you enable IP Shortcuts, you must configure an IS-IS accept policy or use an exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

About This Task

The tunnel source IP address can be a brouter port IP, a CLIP IP, or a VLAN IP.

For information about product support, see [Fabric Extend Considerations](#) on page 1312.



Note

VRF is an optional parameter. If you do not configure a VRF, then FE uses the GRT.



Important

Switches that support a single active VRF have feature interactions with Fabric Extend. For more information, see [VRF Lite Configuration Rules](#) on page 4078. To assist with the single-active VRF restrictions, an *overlay* parameter exists for the *ip-tunnel-source-address* command.

For a logical IS-IS interface, Layer 2 and Layer 3 refer to the following use cases:

- Layer 2 — Fabric Extend VID (FE-VID)
- Layer 3 — Fabric Extend IP (FE-IP)

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```

2. Configure the IP tunnel source address:

```
ip-tunnel-source-address <A.B.C.D> [vrf WORD<1-16>] [overlay]
```

3. Enter Global Configuration mode:

```
exit
```

4. Create a logical IS-IS interface. Use the command that applies to your interface type:

- For FE-VID interfaces (layer 2 core network), enter:

```
logical-intf isis <1-255> vid {vlan-id[-vlan-id][,...]} primary-vid
<2-4059> port <slot/port> mlt PT_MLT<1-512> [name WORD<1-64>]
```



Note

The primary VLAN ID (**primary-vid** must be one of the VIDs in the **vid {vlan-id[-vlan-id][,...]}**.

- For FE-IP (Layer 3 core network), enter:

```
logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-64>]
```

5. Configure the IS-IS MTU for the FE-IP (Layer 3) logical interface.

```
isis mtu <750-1600>
```

6. Configure IS-IS on the logical interface that you just created:
 - a. Create an IS-IS circuit and interface:


```
isis
```
 - b. Enable the SPBM instance on the IS-IS interface:


```
isis spbm <1-100>
```
 - c. Enable the IS-IS circuit/interface:


```
isis enable
```
 - d. Exit logical interface configuration mode:


```
exit
```
7. (Optional) If you want to deploy parallel tunnels, configure a backup IS-IS logical interface:
 - a. Configure an additional source address and optional VRF to use for the backup tunnel:


```
logical-intf isis <1-255> dest-ip <A.B.C.D> src-ip <A.B.C.D> [vrf WORD<1-16>]
```
 - b. (Optional) Configure the IS-IS MTU for the backup tunnel interface. This option is available for FE-IP only.


```
isis mtu <750-1600>
```
 - c. Configure the backup tunnel interface as an IS-IS logical interface. For the backup tunnel, enter the IS-IS commands that you entered for the primary tunnel in step 6 on page 1317.

What to Do Next

Configure Fabric Extend on the switch at the other end of the tunnel.

Variable Definitions

The following table defines parameters for the **ip-tunnel-source-address** command.

Variable	Value
<code><A.B.C.D></code>	Specifies the IS-IS IPv4 tunnel source address, which can be a brouter interface IP, a CLIP IP, or a VLAN IP.
<code>overlay</code>	Permits the configuration of the tunnel source address even though it belongs to a VRF with an attached I-SID.
<code>vrf WORD<1-16></code>	Specifies the VRF name associated with the IP tunnel.

The following tables define parameters for the **logical-intf isis** command, depending on whether you have a Layer 2 or Layer 3 core.

Table 130: Layer 2 core (FE-VID)

Variable	Value
<1-255>	Specifies the index number that uniquely identifies this logical interface.
<i>mlt PT_MLT</i> <1-512>	Specifies the MLT ID that the logical interface is connected to in a Layer 2 network.
<i>name WORD</i> <1-64>	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
<i>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]}</i>	Specifies the physical port that the logical interface is connected to in a Layer 2 network.
<i>primary-vid</i> <2-4059>	Specifies the primary tunnel VLAN ID associated with this Layer 2 IS-IS logical interface.
<i>vid {vlan-id [-vlan-id] [,...] }</i>	Specifies the list of VLANs that are associated with this logical interface.

Table 131: Layer 3 core (FE-IP)

Variable	Value
<1-255>	Specifies the index number that uniquely identifies this logical interface.
<i>dest-ip</i> <A.B.C.D>	Specifies the tunnel destination IP address of the remote BEB.
<i>name WORD</i> <1-64>	Specifies the administratively-assigned name of this logical interface.
<i>src-ip</i> <A.B.C.D> <i>[vrf WORD</i> <1-16>]	Configures an additional source address and optional VRF to use as the parallel tunnel to create a backup adjacency. The VRF is the next-hop VRF to reach the logical tunnel destination IP associated with the parallel tunnel. To use an IPsec-encrypted tunnel as the parallel tunnel, ensure that you configure the same source IP address on the logical IS-IS interface and in the Fabric IPsec Gateway virtual machine.

The following table defines parameters for the **isis mtu** command.

Variable	Definition
<750-1600>	Specifies the Maximum Transmission Unit (MTU) size in bytes for each IS-IS packet that uses this logical interface. The default is 1600. This command is available for Layer 3 core (FE-IP) interfaces only.

Configure IS-IS Hello Padding

Perform this procedure to dynamically configure IS-IS hello padding on all IS-IS network-to-network interface (NNI) links. IS-IS hello padding is enabled by default.

About This Task

Disable hello padding if the WAN-line MTU is less than 1596 bytes and fragmentation and reassembly functionality is enabled.

IS-IS logical-interfaces with a configured IS-IS MTU do not perform IS-IS Hello Padding, regardless of the global IS-IS Hello Padding configuration.

Procedure

1. Enter IS-IS Router Configuration mode:


```
enable

configure terminal

router isis
```
2. Perform one of the following actions:
 - Disable hello padding: `no hello-padding`
 - Enable hello padding, if previously disabled: `hello-padding`
3. Verify the configuration:


```
show isis
```

Example

Verify IS-IS hello padding status:

```
Switch:1>show isis
=====
                        ISIS General Info
=====
                        AdminState : enabled
                        RouterType : Level 1
                        System ID : 0014.c7e1.33df
Max LSP Gen Interval : 900
                        Metric : wide
Overload-on-startup : 20
                        Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
                        spf-delay : 100
Router Name : Switch1
ip source-address : 41.41.41.100
ipv6 source-address : 41:0:0:0:0:0:100
ip tunnel source-address : 11.11.12.11
                        Tunnel vrf : spboip
ip tunnel mtu : 1950
Num of Interfaces : 2
Num of Area Addresses : 1
inband-mgmt-ip :
                        backbone : disabled
Dynamically Learned Area : 00.0000.0000
FAN Member : No
```

```

Hello Padding : enabled
Multi-Area OperState : disabled
Multi-Area Flags : home-always-up
Designated BN Hostname : BN4435
Designated BN Home System ID : 0010.7400.4435
Designated BN Remote System ID : 0020.7400.4435

```

Adjust the TCP Maximum Segment Size

Adjust the TCP maximum segment size (MSS) to improve the throughput for the TCP session over a Fabric Extend (FE) adjacency.

About This Task

The default value, when enabled, is 1300.

Procedure

1. Enter Global Configuration mode:


```
enable
```

```
configure terminal
```
2. Use one of the following commands to configure the MSS adjustment functionality as required:
 - a. Configure an explicit MSS adjust value:


```
ip tcp adjust-mss <500-1900> [enable]
```
 - b. Disable MSS adjustment explicitly:


```
no ip tcp adjust-mss enable
```
3. Verify the configuration:


```
show ip tcp adjust-mss
```

Examples

Configure an MSS value of 1100 and verify the configuration.

```

Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip tcp adjust-mss 1200
Switch:1>show ip tcp adjust-mss
=====
                        IP TCP Adjust MSS
=====
ENABLE          STATUS      TCP MSS      TCP MSS
                        TYPE      VALUE
-----
TRUE            ACTIVE      MANUAL-CONFIG  1200

```

Configure BFD on a Fabric Extend Tunnel

About This Task

Use the following procedure to configure BFD on a Fabric Extend Tunnel.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Enable BFD:
router bfd enable
3. In the VLAN Interface Configuration mode, you can enable BFD:
ip bfd enable
4. In the Loopback Interface Configuration mode, you can enable BFD:
ip bfd enable
5. Enable BFD on an IS-IS Logical Interface:
logical-intf isis <1-255> bfd enable

Example

Enable BFD on a Fabric Extend tunnel:

```
Switch:1>enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router bfd enable
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip bfd enable
Switch:1(config-if)#logical-intf isis 1
Switch:1(config-isis-1-1.2.3.5)#bfd enable
```

Display IS-IS Logical Interfaces

Use the following procedure to display the Intermediate-System-to-Intermediate-System (IS-IS) logical interfaces configured on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IS-IS logical interface:
show isis logical-interface [name]

Examples

Example of a Layer 2 Core:

```
Switch:1>show isis logical-interface
=====
ISIS Logical Interfaces
=====
```

IFIDX	NAME	ENCAP TYPE	L2_INFO PORT/MLT	VIDS (PRIMARY)	TUNNEL DEST-IP	L3_TUNNEL PORT/MLT	NEXT_HOP_INFO VLAN	VRF	BFD STATUS	TUNNEL SRC-IP	ORIGIN	ISIS MTU	SDWAN OPER STATE
1	SD-WAN-1	IP	--	--	192.0.2.3	Port1/44	4047	sd-wan	disabled	192.0.2.1	ZTF	1400	UP
2	--	L2-P2P-VID	Port2/1	101,201(101)	--	--	--	--	disabled	--	config		N/A
3	--	L2-P2P-VID	Port1/3	102,202(102)	--	--	--	--	disabled	--	config		N/A

```
-----
3 out of 3 Total Num of Logical ISIS interfaces
-----
```

Example of a Layer 3 Core:

```
Switch:1>show isis logical-interface
=====
```

```

=====
ISIS Logical Interfaces
=====
IFIDX NAME          ENCAP L2_INFO VIDS      TUNNEL  L3_TUNNEL NEXT_HOP_INFO  BFD  TUNNEL  ORIGIN  ISIS  SDWAN
TYPE  PORT/MLT (PRIMARY) DEST-IP  PORT/MLT VLAN    VRF    STATUS SRC-IP          MTU  OPER STATE
-----
1 SD-WAN-1          IP    --    --      192.0.2.3  Port1/44  4047  sd-wan  disabled 192.0.2.1  ZTF    1400  UP
2 SPBoIP_T1         IP    --    --      192.0.2.15 Port1/25  500   vrf23   disabled 192.0.2.16 CONFIG 1000  N/A
3 SPBoIP_T2         IP    --    --      192.0.2.224 MLT10    2     vrf24   disabled 192.0.2.22 CONFIG 1600  N/A
4 Virtual-link-4    IP    --    --      4.4.4.4   PortRX-NNI 4051  GlobalRouter disabled 2.2.2.2  CONFIG 1600  N/A
5 Virtual-link-5    IP    --    --      5.5.5.5   Null      0     GlobalRouter disabled 2.2.2.2  CONFIG 1600  N/A
-----
5 out of 5 Total Num of Logical ISIS interfaces
=====

```

The command **show isis logical-interface** truncates the IS-IS logical interface name to the first 16 characters. To view the entire name (up to a maximum of 64 characters), use the command **show isis logical-interface name**.

```

Switch:1>show isis logical-interface name
=====
ISIS Logical Interface name
=====
ID      NAME
-----
1       SD-WAN-1
2       SPBoIP_T1
3       SPBoIP_T2
4       Virtual-link-4
5       Virtual-link-5
-----
4 out of 4 Total Num of Logical ISIS interfaces
=====

```

Variable Definitions

The following table defines parameters for the **show isis logical-interface** command.

Variable	Value
name	Displays the full name of the IS-IS logical interface (up to a maximum of 64 characters).

Display BFD Fabric Extend Neighbor Information

About This Task

Use this procedure to display BFD Fabric Extend neighbors.

Procedure

- To enter User EXEC mode, log on to the switch.
- Display the BFD configurations:

```
show ip bfd neighbors vrf WORD<1-16>
```

Example

```

Switch:1>show ip bfd neighbors vrf vrf30
=====
BFD Session - VRF vrf30
=====
MY_DISC    YOUR_DISC  NEXT_HOP      STATE    MULTI  MIN_TX  MIN_RX  ACT_TX  DETECT_TIME  REMOTE_STATE  APP    RUN

```



```

1      1      192.0.2.11      UP      3      200      200      1000      600      UP      ISIS      ISIS
-----
1 out of 1 BFD session displayed
-----

```

Fabric Extend Configuration using EDM

The following sections provide procedural information you can use to configure Fabric Extend (FE) using Enterprise Device Manager (EDM).

Complete the configurations in the following table to configure Fabric Extend.

Table 132: Fabric Extend Configuration

Step	Procedure	Description
1	Configure Fabric Extend Tunnels on page 1323	Configure the tunnel source address on the Globals tab.
2	Configure Fabric Extend Logical Interfaces on page 1327	Configure the IS-IS logical interface for the Fabric Extend tunnel.
3	Configure Fabric Extend on the switch at the other end of the tunnel	

Configure Fabric Extend Tunnels

Use the following procedure to configure the tunnel source address on a Fabric Extend (FE) tunnel. A typical deployment would be between a Main office and a Branch office.

You can also configure an optional parallel tunnel between the same two nodes to create a backup Fabric Extend adjacency. For more information, see [Configure Fabric Extend Logical Interfaces](#) on page 1327.

Before You Begin

- If using the tunnel originating address, the logical interface address, on the GRT, Fabric Extend has the following requirements:
 - The tunnel source IP address must be on the GRT, not on a VRF.



Note

To configure Fabric Extend using a tunnel source IP address on the GRT, use separate IP addresses for the SPBM IP Shortcuts (**IpSourceAddress**) and the Fabric Extend **IpTunnelSourceAddress** address. Exclude the **IpSourceAddress** address with an IS-IS accept policy if you want these IP addresses to be the same. You cannot use the **redistribute** command with a route-map exclusion.

Specify a CLIP interface to use as the **IpSourceAddress** for SPBM IP shortcuts.

- If you enable IP Shortcuts, you must configure an IS-IS accept policy or use an exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.
- If IP Shortcuts is enabled, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.
- If you are using the tunnel originating address on a VRF, configure a CLIP and tunnel source IP address on the VRF.

About This Task

The tunnel source IP address can be either a brouter port IP, a CLIP IP, or a VLAN IP.



Note

VRF is an optional parameter. If you do not configure a VRF, then FE uses the GRT.



Important

Switches that support a single active VRF have feature interactions with Fabric Extend. For more information, see the Configuration Rules in the VRF feature content. To assist with the single-active VRF restrictions, an overlay parameter is available for the IP tunnel source address configuration.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Globals** tab.
4. For **IpTunnelSourceAddress**, type the IP tunnel source address.
5. (Optional) For **IpTunnelVrf**, select the VRF.
6. Select **Apply**.

What to Do Next

[Configure Fabric Extend Logical Interfaces](#) on page 1327

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminState	Specifies the global status of IS-IS on the switch: on or off. The default is off .
LevelType	<p>Sets the router type globally:</p> <ul style="list-style-type: none"> • level1: Level-1 router type • level1and2: Level-1/2 router type is not supported. <p>The default value is level1.</p>

Name	Description
SystemId	<p>Specifies the IS-IS system ID for the switch. Valid value is a 6-byte value in the format <xxxx.xxxx.xxxx>.</p> <p>Important: After you have configured the SPBM nickname and enabled IS-IS, you must change the nickname when you change the system ID. If you want to maintain the same nickname with a different system ID, see Job aid on page 1266.</p>
MaxLspGenInt	<p>Specifies the maximum interval, in seconds, between generated LSPs by this Intermediate system. The value must be greater than any value configured for RxmtLspInt.</p> <p>The default value is 900 seconds.</p>
Csnplnt	<p>Specifies the Complete Sequence Number Packet (CSNP) interval in seconds. This is a system level parameter that applies for L1 CSNP generation on all interfaces.</p> <p>The default value is 10.</p>
RxmtLspInt	<p>Specifies the minimum time between retransmission of an LSP, in seconds. This defines how fast the switch resends the same LSP. This is a system level parameter that applies for L1 retransmission of LSPs.</p> <p>The default value is 5 seconds.</p>
PSNPInterval	<p>Specifies the Partial Sequence Number Packet (PSNP) interval in seconds. This is a system level parameter that applies for L1 PSNP generation on all interfaces.</p> <p>The default value is 2.</p>
SpfDelay	<p>Specifies the SPF delay in milliseconds. This value is used to pace successive SPF runs. The timer prevents two SPF runs from being scheduled very closely.</p> <p>The default value is 100 milliseconds.</p>
HostName	<p>Specifies the IS-IS system name. This can be used as the host name for dynamic host name exchange in accordance with RFC 2763.</p> <p>By default, the system uses the global system name as the IS-IS system name, until you manually configure it. Also, the system does not save the global system name in the running configuration on the switch.</p>

Name	Description
IpSourceAddress	Specifies IP source address for SPBM IP shortcuts.
Ipv6SourceAddress	Specifies IPv6 source address for SPBM IP shortcuts.
IpTunnelSourceAddress	Specifies the IS-IS IPv4 tunnel source address.
IpTunnelVrf	Specifies the VRF name associated with the IP tunnel.
IpTunnelOverlay	Permits the configuration of the tunnel source address even though it belongs to a VRF with an attached I-SID. The default is disabled.
DesignatedBNHostName Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the hostname of the designated boundary node.
DesignatedBNHomeSystemId Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the system ID of the designated boundary node in the home area.
DesignatedBNRemoteSystemId Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the system ID of the designated boundary node in the remote area.
MgmtCliIpAddr	Specifies the DvR management IP address for this node, in the DvR domain.
BackboneEnable	Select to enable this node to join the DvR backbone so that it can receive redistributed DvR host routes from all DvR Controllers in the network.
FanMember	Specifies whether the node is a member of the Fabric Area Network (FAN) .
DynamicallyLearnedArea	For FAN members, specifies the IS-IS area that is dynamically learned from the neighbor's Hello PDU if the node does not have the IS-IS manual area configured.
MAOperState Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the Multi-area SPB operational state. The default is disabled.
AreaName	Configures the system area name.

Name	Description
HelloPadding	Configures IS-IS hello padding on all IS-IS network-to-network interface (NNI) links. IS-IS hello padding is enabled by default. IS-IS logical-interfaces with a configured IS-IS MTU do not perform IS-IS Hello Padding, regardless of the global IS-IS Hello Padding configuration.
MAHomeAlwaysUp Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Forwards traffic from UNI ports on the Multi-area SPB designated boundary node into the remote-area, and receive traffic from the remote area that is destined for the UNIs ports, without requiring an established IS-IS adjacency in the home area. Important: To avoid packet loops back into the home area, this feature does not operate on non-designated boundary nodes. The default is disabled.

Configure Fabric Extend Logical Interfaces

Use this procedure to configure the Fabric Extend logical interface.

Configuring Fabric Extend (FE) consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Before You Begin

Configure the Fabric Extend tunnel with a tunnel source IP address. See [Configure Fabric Extend Tunnels](#) on page 1323.

About This Task

VRF is an optional parameter. If you do not configure a VRF, then FE uses the GRT.

For a logical IS-IS interface, Layer 2 and Layer 3 refer to the following use cases:

- Layer 2 — Fabric Extend VID (FE-VID)
- Layer 3 — Fabric Extend IP (FE-IP)

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Logical Interfaces** tab.
4. Select **Insert**.
5. For **Id**, type the index number that uniquely identifies this logical interface.
6. For **Name**, type the name of this logical interface.

7. For **Type**, select the type of core network that the tunnel will traverse:



Note

Different fields are available depending on the type of core network you select.

- If it is a Layer 2 Core, select **layer2**.
 - If it is a Layer 3 Core, select **ip**.
8. (Optional) For **BFDEnable**, select **enable** to enable BFD.
9. For a Layer 2 Core, configure the following parameters:
- a. For **DestIfIndex**, select the physical port that the logical interface connects to or enter the name of the MLT.
 - b. For **Vids**, type the list of VLANs for this logical interface.
 - c. For **PrimaryVid**, type the primary tunnel VLAN ID.



Note

The primary VLAN ID must be one of the VIDs listed in **Vids**.

10. For a Layer 3 Core, complete the following parameters:
- a. For **DestIPAddr**, type the destination IP address for the logical interface.
 - b. For **ISIS Mtu**, enter the Maximum Transmission Unit value, in bytes.
11. (Optional) Configure a parallel tunnel for Fabric Extend:
- a. For **SrcIPAddr**, type the source address for the parallel tunnel.
 - b. For **NextHopVrf**, select the VRF name to reach the logical tunnel destination IP associated with the parallel tunnel.
 - c. Select **Apply**.
12. Select **Insert**.

What to Do Next

Configure Fabric Extend on the switch at the other end of the tunnel.

Logical Interfaces Field Descriptions

Use the data in the following table to use the **Logical Interfaces** tab and the **Insert Logical Interfaces** dialog. The available fields in the dialog differ depending on the type of core you select: **layer 2** or **ip**.

Name	Description
Id	Specifies the index number that uniquely identifies this logical interface. This field displays on the Insert Logical Interfaces dialog only.
IfIndex	Specifies the index number that uniquely identifies this logical interface. This field is read-only. This field displays on the Logical Interfaces tab only.

Name	Description
Name	<p>Specifies a name associated with this logical interface, which can be up to 64 characters.</p> <p>Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link-x</code>, where <code>x</code> is the logical interface ID.</p>
Type	<p>Specifies the type of logical interface to create:</p> <ul style="list-style-type: none"> Specify layer 2 for a Layer 2 core network that the tunnel will traverse. Specify ip for a Layer 3 core network that the tunnel will traverse.
DestIPAddr	Specifies the destination IP address for the IP-type logical interface.
DestIfIndex	Specifies the physical port or MultiLink Trunking (MLT) that the Layer 2 logical interface is connected to.
Vids	Specifies the list of VLANs that are associated with this logical interface.
PrimaryVid	Specifies the primary tunnel VLAN ID associated with this Layer 2 Intermediate-System-to-Intermediate-System (IS-IS) logical interface.
CircIndex	<p>Identifies the IS-IS circuit created under the logical interface.</p> <p>This field displays on the Logical Interfaces tab only.</p>
NextHopVrf	<p>Displays the next-hop VRF name to reach the logical tunnel destination IP.</p> <p>This field displays on the Logical Interfaces tab only.</p> <p>You can use this field to specify the VRF to reach the logical tunnel destination IP associated with a parallel tunnel.</p>
ISIS Mtu	Specifies the Maximum Transmission Unit (MTU) size in bytes for IS-IS packets that use this logical interface. The default value is 1600.
BfdEnable	Enables or disables BFD on an IS-IS Logical Interface.
SrcIPAddr	<p>Configures an additional source address to use as the parallel tunnel to create a backup adjacency.</p> <p>Note: To use an IPsec-encrypted tunnel as the parallel tunnel ensure that you configure the same source IP address on the logical IS-IS interface and in the Fabric IPsec Gateway virtual machine.</p>

Name	Description
Origin	Specifies the origin of the IS-IS logical interface configuration, either through Zero Touch Fabric Configuration (ZTF) or manual configuration (config) through CLI or EDM.
Multi-Area Virtual Link Note: Exception: only applies to 7520 Series and 7720 Series.	Enables (true) or disables (false) a Multi-area boundary node virtual NNI link on an IS-IS Logical Interface. The default is disabled.

Adjust the TCP Maximum Segment Size

Adjust the TCP maximum segment size (MSS) to improve the throughput for the TCP session over a Fabric Extend (FE) adjacency.

About This Task

By default, this functionality is disabled. The default value, when enabled, is 1300.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.
4. Select **TcpAdjustMssEnable** to enable this functionality.
5. (Optional) Enter a value in **TcpAdjustMssValue**.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Forwarding	Configures the system for forwarding (routing) or for dropping. The default value is forwarding. You cannot configure this option on a DvR Leaf node.
DefaultTTL	Configures the default time-to-live (TTL) value for a routed packet. TTL indicates the maximum number of seconds elapsed before a packet is discarded. Enter an integer from 1 to 255. The default value of 255 is used if a value is not supplied in the datagram header.

Name	Description
ReasmTimeout	Specifies the maximum number of seconds that received fragments are held while they wait for reassembly. The default value is 30 seconds.
ICMPUnreachableMsgEnable	<p>Enables the generation of Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this system. These messages help determine if the system is reachable over the network. The default is disabled.</p> <p>Important: As a best practice, enable icmp-unreach-msg only if it is absolutely required. If icmp-unreach-msg is enabled and a packet is received for which there is no route in the routing table, CPU utilization can dramatically increase.</p>
ICMPRedirectMsgEnable	Enables or disables the system sending ICMP destination redirect messages.
IcmpEchoBroadcastRequestEnable	Enables or disables IP ICMP echo broadcast request feature. The default is enabled.
IcmpDropFragmentsEnable	Enables or disables IPv4 Fragmented ICMP packet filtering globally. The default is disabled.
AlternativeEnable	<p>Globally enables or disables the Alternative Route feature.</p> <p>If the alternative-route parameter is disabled, all existing alternative routes are removed. After the parameter is enabled, all alternative routes are re-added. The default is enabled.</p>
RouteDiscoveryEnable	Enables the ICMP Router Discovery feature. The default is disabled (not selected). Use ICMP Router Discovery to enable hosts attached to multicast or broadcast networks to discover the IP addresses of neighboring routers.
AllowMoreSpecificNonLocalRouteEnable	Enables or disables a more-specific nonlocal route. If enabled, the system can enter a more-specific nonlocal route into the routing table. The default is disabled.

Name	Description
SuperNetEnable	Enables or disables supernetting. If supernetting is globally enabled, the system can learn routes with a route mask less than 8 bits. Routes with a mask length less than 8 bits cannot have ECMP paths, even if you globally enable the ECMP feature. The default is disabled.
UdpChecksumEnable	Enables or disables the UDP checksum calculation. The default is enable.
SourceRouteEnable	Enables or disables IP Source Routing globally. It is disabled by default.
SpbMulticastPolicyEnable	Enables the IP SPB Multicast Policy .
SpbMulticastPolicyRmap	Configures the name of the route-map for IP SPB Multicast Policy .
SpbMulticastPolicyApply	Applies the IP SPB Multicast Policy . The default value is noAction.
ARPLifeTime	Specifies the lifetime of an ARP entry within the system, global to the switch. The default value is 360 minutes.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled. After ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.
EcmpMaxPath	Globally configures the maximum number of ECMP paths. You cannot configure this feature unless ECMP is enabled globally.
Ecmp1PathList	Selects a preconfigured ECMP path.
Ecmp2PathList	Selects a preconfigured ECMP path.
Ecmp3PathList	Selects a preconfigured ECMP path.
Ecmp4PathList	Selects a preconfigured ECMP path.
Ecmp5PathList	Selects a preconfigured ECMP path.
Ecmp6PathList	Selects a preconfigured ECMP path.
Ecmp7PathList	Selects a preconfigured ECMP path.
Ecmp8PathList	Selects a preconfigured ECMP path.
EcmpPathListApply	Applies changes in the ECMP pathlist configuration, or in the prefix lists configured as the pathlists.

Name	Description
TcpAdjustMssEnable	Adjusts the TCP maximum segment size (MSS) to improve the throughput for the TCP session over a Fabric Extend (FE) adjacency. The default is disabled.
TcpAdjustMssStatus	Displays the activation status of the MSS adjustment functionality.
TcpAdjustMssType	Displays if the MSS adjustment value is manually configured or auto-derived. The software does not support auto-derived values for this feature.
TcpAdjustMssValue	Configures the MSS adjustment value. • The default value is 1300.

Display the Logical Interface Next Hop

Use the following procedure to display the next hop for the logical interface.

Procedure

1. In the navigation pane, expand **Configuration > Fabric > IS-IS**.
2. Select the **Logical Interfaces NextHop** tab.

Logical Interfaces NextHop Field Descriptions

Use the data in the following table to use the **Logical Interfaces NextHop** tab.

Name	Description
Id	Shows a unique value that identifies the logical interface tunnel.
Ip	Shows a unique value that identifies the next hop IP address of the logical interface tunnel.
DestIfIndex	Shows the next hop destination interface index to reach the next hop IP of the logical interface tunnel.
DestVid	Shows the next hop destination VLAN ID to reach the next hop IP of the logical interface tunnel.

Fabric Attach

Table 133: Fabric Attach product support

Feature	Product	Release introduced
Fabric Attach Server	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

With Fabric Attach, network edge devices that do not support Shortest Path Bridging (SPB), MAC-in-MAC encapsulation (802.1ah) or service identifiers (I-SIDs) can take advantage of SPB infrastructure. To attach to an SPB network, edge devices signal an SPB-aware FA Server to automatically configure the I-SIDs. The edge devices can then utilize existing SPB features across the fabric and leverage SPB infrastructure capabilities without manual configuration. Fabric Attach uses the IEEE 802.1AB Logical Link Discovery Protocol (LLDP) to signal a desire to join the SPB network.

FA uses the client-server model. An initial handshake occurs between the FA Server and the FA Client. After the discovery phase is complete, the FA Server accepts requests (from FA Clients) to add the C-VID (VLAN ID) and I-SID elements in the SPB network, and also automatically configures the necessary C-VID and I-SID. The FA Server then responds with an acknowledgement of whether the request succeeded. FA Clients can also be aggregated into a proxy device that handles the handshakes and requests on behalf of many clients, to the server. All of the discovery handshakes and I-SID mapping requests are then transferred using LLDP Type, Length, Value (TLV) fields.

FA leverages LLDP to discover directly connected FA peers and to exchange information associated with FA between those peers. Based on the LLDP standard, FA information is transmitted using organizational TLVs within LLDP Protocol Data Units (PDU).

FA Zero Touch Client Attachment

Table 134: Fabric Attach Zero Touch Client Attachment product support

Feature	Product	Release introduced
Fabric Attach Zero Touch Client Attachment	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

FA Zero Touch Client Attachment eases the configuration process on FA-capable devices by automating specific configuration tasks required for FA functionality.



Note

Only the base functionality of Zero Touch Client Attachment is supported.

After you initially configure Zero Touch Client Attachment on the FA Server, the settings are exported to receiving FA devices, where the required configuration tasks are automatically performed.

Base Zero Touch Client Attachment operation is tightly coupled with FA operation. Although you can enable or disable Zero Touch Client Attachment separately from FA, the feature is dependant on data that is only available during exchanges between the FA Server and FA Proxies, after a primary FA Server has been selected. By default, base Zero Touch Client Attachment support is enabled.

Base Zero Touch Client Attachment operation, when enabled, extracts management VLAN data from the primary FA Server advertisements and uses this data to update the in-use management VLAN if applicable. An FA Client can also utilize FA-provided management VLAN data after the FA Proxy or Server is discovered.

Zero Touch is active when the following criteria are met:

- On an FA Proxy:
 - Zero Touch Client Attachment is enabled
 - Fabric Attach is enabled
 - A primary FA Server is discovered and selected
- On an FA Server:
 - Zero Touch Client Attachment is enabled
 - FA is enabled

- FA Proxies or FA Clients are discovered

The switch supports configurable VLANs in the range of 1 to 4059. VLAN 0 is invalid. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. VLAN IDs on the switch range from 2 to 4094 but, by default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the *vrf-scaling* and *spbm-config-mode* boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.

**Note**

You must enable Base Zero Touch **auto-client attach** and define the target Fabric Attach client to initiate Zero Touch Client Attachment processing.

FA Signaling generated by an FA Proxy or Server contains management VLAN data. If the management VLAN advertised by the primary FA Server differs from the management VLAN currently configured on the FA Proxy, Zero Touch Client Attachment initiates the following:

- VLAN creation — If the FA Server-specified management VLAN does not exist on the FA Proxy, Zero Touch Client Attachment creates a port-based VLAN.
- Management VLAN update — The created port-based VLAN becomes the designated management VLAN for the FA Proxy. No operations related to the previous management VLAN, such as port membership updates or VLAN deletion, are performed.
- Port VLAN membership update (FA Proxy/Server) — If required, Zero Touch Client Attachment updates the port VLAN membership to ensure that the uplink port through which the primary FA Server is accessed is a member of the management VLAN, for network accessibility.
- Port Default VLAN (PVID) update — The port-based PVID is automatically updated based on the VLAN ID value.
- Port Default Priority update — The default 802.1p user priority for the port is updated based on the specified port priority value of the Zero Touch client (range is 0–7).
- Zero Touch Client Specification removal — All Zero Touch client-related settings are updated based on the FA client discovery. Deleting a Zero Touch client specification or disabling any related Zero Touch option does not result in the immediate removal of any previously applied settings.

**Note**

The FA Proxy does not update the acquired management VLAN if the primary FA Server is lost. This data is updated if the management VLAN advertised by the current primary FA Server changes or if another primary FA Server is selected and new management VLAN data is advertised by the server.

Management VLAN and port membership updates performed by Zero Touch are maintained in non-volatile storage and are restored following a system reset. You must remove or update these configuration settings if they are deemed unnecessary at a later time.

- IP Address Source Mode Update — Updates the IP address source mode of the receiving device to DHCP-When-Needed, to initiate DHCP-based IP address acquisition if necessary.
- Automation of the FA Client Port Mode — Automates the configuration of EAP port modes based on the type of discovered FA Clients. Applies to FA Proxy and FA Server devices. Automated configuration is applied only to FA-enabled ports.
- ZTC Installation — Initiates ZTC installation on applicable ports on the receiving device. Applies to FA Proxy and, in a limited manner, to FA Server devices. Automated configuration is applied only to FA-enabled ports.
- Auto PVID FA Client Port Mode — Initiates automatic port PVID, port management VLAN membership and port tagging mode based on the type of discovered FA device. Applies to FA Proxy and FA Server devices. Automated configuration is applied only to FA-enabled ports. This configuration is incompatible with the automatic FA Client Port Mode and ZTC Automatic attach options.

Fabric Attach Components

FA components dynamically communicate with each other using FA signaling.

FA Signaling

FA has defined organizational specific TLVs within the standard LLDP protocol, to exchange messages and data amongst components of an FA solution. The FA TLVs facilitate handshaking and authentication, processing of requests for the creation of services, and providing responses on whether the requests succeeded. In addition, these services are deleted when the service requests are terminated, or when the authentication criteria are no longer valid. All components that participate in FA must be able to send, receive, and interpret the FA TLVs.

FA Components

FA includes the following network elements as components:

- FA Server:

An SPB-capable switch at the edge of a Fabric Connect cloud.

An FA Server receives requests from FA Clients or FA Proxies to create services with specific I-SID-to-VLAN bindings. The FA Server completes the association between conventional networks and fabric-based virtual service networks. For more details on the operation of an FA Server, see [Fabric Attach Server](#) on page 1338.

- FA Proxy:

A network switch that supports the definition of I-SID-to-VLAN assignments and has the ability to advertise these assignments for possible use by an FA Server. FA Proxy switches also support the client mode for directly attached users or end devices. Typically, FA Proxies support downstream FA Client devices, while being directly connected to an upstream FA Server device.

- FA Client:

A network attached end-point device that advertises I-SID-to-VLAN binding requests for service creation, to an FA Proxy or an FA Server. FA Clients use FA signaling to automatically attach to fabric services.

Fabric Attach Server

FA Server operation

In an FA solution, the FA Server performs the role of connecting FA Clients and FA Proxies to the SPB fabric, with minimal configuration. As part of the discovery handshake between the FA Server and client or proxy devices, LLDP PDUs are exchanged. Using standard LLDP, the FA Server learns neighbors, that include the proxy and client devices. In addition, the FA Server transmits organizational-specific element-discovery TLVs that are used by the client or proxy device to recognize its attachment to the FA Server.




Figure 73: Fabric Attach Server connecting client or proxy devices to the Fabric network

After the initial discovery handshake is complete, the client or proxy device transmits I-SID-to-VLAN assignment mapping requests to the FA Server to join the SPB fabric. These requests include the C-VID (VLAN ID) and the I-SID that the client or proxy device needs to join. The FA Server then creates the requested C-VID and I-SID on its device. It then responds with a PDU (containing the FA-specific TLV) to indicate whether the request succeeded. The I-SID thus created is a ELAN I-SID with endpoints of type Switched UNI. After I-SID creation, the I-SID is also advertised to the SPB network by IS-IS.

The traffic that is sent to or received from the SPB cloud is MAC-in-MAC (MiM) encapsulated. The FA Server, being SPB-capable, decapsulates the MiM traffic. If the I-SID matches the I-SID created on behalf of the client or proxy, the FA Server sends the traffic to that client or proxy and passes it on the C-VID that it expects.

FA Server configuration

An FA Server can be configured at two levels—global and interface.

Configuration at the global level enables or disables FA on the entire switch. However, for attachment of clients or proxy devices, you must also configure FA at the interface

level. Interfaces can be ports (including channelized ports), MLTs, SMLT or LACP MLTs. Enabling FA on an interface also enables transmission of LLDP packets that contain the FA-specific TLVs.

When you disable FA on an interface, LLDP transmission automatically stops on that interface.

**Caution**

Disabling FA or IS-IS triggers a flush of FA information on the switch. Disabling FA at the global level flushes FA element-discovery information and mappings. Disabling at the interface level flushes element-discovery information and mappings associated with that interface.

**Important**

The only provisioning mode supported on the FA Server is SPB.

FA Proxies and FA Clients

The configuration mode of FA Proxies and FA Clients is not supported. However, in an FA solution, the FA Server interacts with FA Proxies and FA Clients by accepting LLDP PDUs (containing FA TLVs) and using them to automatically create Switched UNI I-SIDs and endpoints, based on the mapping requests contained in those TLVs. For more information, see [FA TLVs](#) on page 1339.

Fabric Attach operation

The following sections detail FA operation.

FA TLVs

FA leverages LLDP to discover directly connected FA peers and to exchange information associated with FA amongst those peers. FA information is transmitted using company-specific proprietary organizational Type, Length, Value (TLV) fields within LLDP Protocol Data Units (PDU). The following section describes the TLVs for FA.

FA uses two TLVs:

- FA Element TLV
- FA Assignment TLV

FA Element TLV

The FA Element TLV is used by FA elements to advertise Fabric Attach capabilities. This data forms the basis for FA element discovery and is used in the initial handshake between the FA Server and a client or proxy device.

TLV Type [127]	TLV Length [50 octets]	OUI [00-04-0D]	Subtype [11]	HMAC-SHA Digest	Element Type	State	Mgmt VLAN	Rsvd	System ID
7 bits	9 bits	3 octets	1 octet	32 octets	6 bits	6 bits	12 bits	1 octet	10 octets

Figure 74: FA Element TLV format**Table 135: FA Element TLV field descriptions**

Field	Description
TLV Type	Indicates whether the discovered element is a client or a proxy device.
OUI and Subtype	The information in these fields is used in LLDP packet handling.
HMAC-SHA Digest	<p>Data integrity and source validation is supported through the use of the HMAC-SHA256 message authentication. This field supports a digest exchange between the source and destination devices. Symmetric private keys are used for digest generation. The HMAC-SHA256 generated digest size is 32 octets.</p> <p>The HMAC-SHA256 digest is computed starting with the Element Type data, that is, it starts at zero-based byte 38 of the TLV. The digest is then placed in the HMAC-SHA256 Digest field in the TLV prior to transmission. Upon receipt, the digest is again computed and the resulting digest is compared against the received digest. If the received digest is the same as the newly computed digest, the TLV is considered valid and processing commences. If the comparison fails, the TLV is discarded and processing is terminated.</p> <p>Caution:</p> <p>If FA communication occurs between non-secure systems, the HMAC-SHA256 Digest data must always be zero. If one system operates in secure mode and the other operates in non-secure mode, the FA Element TLV is discarded before it is processed by the system operating in secure mode.</p>

Table 135: FA Element TLV field descriptions (continued)

Field	Description
Element Type	<p>Indicates the supported element type. The primary element types are the FA Server, FA Proxy, and FA Client.</p> <p>An FA Server is an SPB capable device that accepts externally generated I-SID-to-VLAN assignments. An FA Proxy is a non-SPBM device that supports I-SID-to-VLAN assignment definitions and advertises these assignments for possible use by an FA Server. An FA Client, also a non-SPBM device, advertises I-SID-to-VLAN assignments to a directly connected FA Proxy or an FA Server. Both tagged and untagged FA Client connections are supported.</p> <p>The list of supported element types and their values are:</p> <ul style="list-style-type: none"> • FA Element Type - Other (1) • FA Server (2) • FA Proxy (3) • FA Server No Authentication (4) • FA Proxy No Authentication (5) • FA Client - Wireless Access Point Type 1, which directly attaches to the SPBM network. • FA Client - Wireless Access Point Type 2, which is tunneled to a controller. • FA Client - Switch (8) • FA Client - Router (9) • FA Client - IP Phone (10) • FA Client - IP Camera (11) • FA Client - IP Video (12) • FA Client - Security Device (13) • FA Client - Virtual Switch (14) • FA Client - Server/Endpoint (15) • FA Proxy Ring (18)
State	<p>Indicates the link tagging requirements in FA Client-sourced frames. This field also indicates the current provisioning mode.</p> <p>The Link VLAN Tagging bit (bit 1) has one of the following values:</p> <ul style="list-style-type: none"> • 0 — indicates that all traffic on the link is tagged. In this case, all discovered FA Clients are treated as tagged. • 1 — indicates that traffic on the link is either tagged or untagged. Here, all discovered FA Clients are treated as untagged. <p>The automatic provisioning mode bits (bits 2 and 3) always have the value 1 for SPB provisioning. The switch only supports the SPB provisioning mode.</p>

Table 135: FA Element TLV field descriptions (continued)

Field	Description
Mgmt VLAN	When you configure a management VLAN on the FA Server, it is included in this field in FA Server or FA Proxy sourced frames, and is used to support management VLAN auto-configuration on the downstream proxy and client devices.
System ID	This field contains connection information that a TLV recipient can use to enforce connectivity restrictions. It contains the system MAC address (6 octets) for MLT configurations and the virtual BMAC address for vIST and SMLT configurations. It also contains information on the connection type such as MLT or SMLT.

Limitations

- The FA Element TLV exists only once in an LLDP PDU and is included in all PDUs when the FA service is enabled.
- The maximum length of the FA Element TLV is 56 bytes.

FA I-SID-to-VLAN Assignment TLV

The FA I-SID-to-VLAN Assignment TLV is used by FA Clients to distribute I-SID-to-VLAN assignments that need to be supported by an FA Proxy or an FA Server.

TLV Type [127]	TLV Length [41-506 octets]	OUI [00-04-0D]	Subtype [12]	HMAC-SHA Digest	Assignment Status	VLAN	I-SID
7 bits	9 bits	3 octets	1 octet	32 octets	4 bits	12 bits	3 octets

Figure 75: FA Assignment TLV format

Some fields are common to both the FA Element and FA Assignment TLVs. The following fields are specific only to the FA Assignment TLV.

Table 136: FA I-SID-to-VLAN Assignment TLV fields

TLV Field	Description
HMAC-SHA Digest	<p>The HMAC-SHA256 digest is computed for the series 1 to 94 of I-SID-to-VLAN assignments, that is, the data for the digest computation starts at zero-based byte 38 of the TLV. The digest is then placed in the HMAC-SHA256 Digest field in the TLV prior to transmission. Upon receipt, the digest is again computed for the series 1 to 94 of I-SID-to-VLAN assignments in the received TLV and the resulting digest is compared against the received digest. If the received digest is the same as the newly computed digest, the TLV is considered valid and processing can commence. If the comparison fails, the TLV is discarded and processing is terminated.</p> <p>Caution: If FA communication occurs between non-secure systems, the HMAC-SHA256 Digest data must always be zero. If one system operates in secure mode and the other operates in non-secure mode, the FA I-SID-to-VLAN Assignment TLV is discarded before it is processed by the system operating in secure mode.</p>
Assignment status	Indicates whether the FA Server accepted or rejected the I-SID-to-VLAN mapping request from a client or proxy device.
VLAN	Indicates the C-VID value advertised by the client or proxy device in the FA I-SID-to-VLAN mapping request.
I-SID	<p>Indicates the I-SID that is advertised by a client or proxy device in the FA I-SID-to-VLAN mapping request. This I-SID is used to create a Switched UNI (ELAN) I-SID.</p> <p>Note: This I-SID cannot be used by IPVPN, MVPN, SPBM dynamic multicast range, or Transparent Port UNI.</p>

Limitations

- The FA I-SID-to-VLAN Assignment TLV is included in an LLDP PDU only if the FA Server and proxy or client devices are directly connected to each other.
- This TLV can exist only once in an LLDP PDU.
- The size limit of this TLV is 511 bytes. This limits the maximum number of I-SID-to-VLAN assignments supported in an LLDP PDU to 94.
- For an FA I-SID-to-VLAN Assignment TLV to be processed, the FA Element TLV must also be present in the LLDP PDU.

FA Element Discovery

The first stage of establishing FA connectivity is element discovery.

On an FA Server, FA is enabled globally by default. However, you must explicitly enable FA on a desired port or MLT interface. After FA is enabled, the FA Server begins

transmitting LLDP PDUs that contain the element discovery TLVs. This information is received by FA Client and FA Proxy devices which in turn also transmit their FA capabilities and settings. After the element handshake completes, the FA Server receives I-SID-to-VLAN assignment mappings from the connected client or proxy devices.

An FA Server can communicate with multiple different FA Client and FA Proxy devices.

FA Data Processing

In this type of FA deployment, a client device (Client 1) attaches to the FA Server (FA Server 1) using a proxy device. Another client device (Client 2) attaches to the FA Server (FA Server 2) at the other edge of the network and also through a proxy device. The following section describes how data is processed when data traffic is transmitted from Client 1 to Client 2.

When Client 1 successfully attaches to FA Server 1, FA Server 1 creates a unique I-SID-to-VLAN mapping for Client 1 on its device. This mapping contains the I-SID and C-VID advertised by Client 1, using the FA Assignment TLV. For example, assume that Client 1 advertises I-SID 200 and C-VID 250.

Similarly, when Client 2 attaches to FA Server 2, FA Server 2 creates an I-SID-to-VLAN mapping for Client 2 on its device with, for example, I-SID 200 and C-VID 100.

When data traffic ingresses FA Server 1 at the FA-enabled port 1/1, it contains the C-VID of Client 1, which is 250. The data is VLAN-encapsulated at this stage. As traffic egresses FA Server 1 into the SPB cloud, it is encapsulated with the ELAN I-SID created on FA Server 1 on behalf of Client 1, that is I-SID 200. The traffic is now MiM encapsulated with I-SID 200.

The following figure depicts VLAN encapsulation of data traffic from the FA Client to the FA Server (at either end of the SPB cloud, labeled as 1) and its MiM encapsulation as it traverses the SPB cloud, labeled as 2.




Figure 76: Data encapsulation — VLAN encapsulation and MiM encapsulation

As traffic exits the SPB cloud and ingresses the remote FA Server 2, it continues to be MiM encapsulated with I-SID 200.

At FA Server 2, the MiM traffic is decapsulated. Because the I-SID in the data packet matches the I-SID created on its device on behalf of Client 2, FA Server 2 prepares to send traffic to Client 2. At this stage, to successfully transmit the data traffic to Client 2, FA Server 2 must use the C-VID that Client 2 expects traffic on. This information is obtained from the I-SID-to-VLAN mapping on FA Server 2 created on behalf of Client 2, which is C-VID 100. Thus FA Server 2 translates the C-VID in its data packets to this VLAN ID, and then passes it on to Client 2.

The following figure depicts the typical MiM encapsulation of a data packet. The B-DA and B-SA components indicated the system ID of the FA Server running SPB.

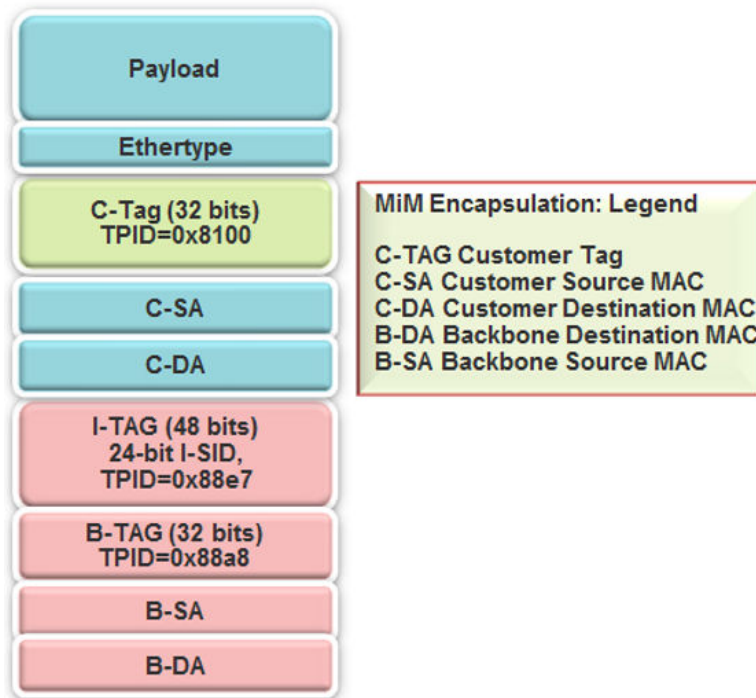


Figure 77: MiM encapsulation

FA Server and I-SID-to-VLAN Assignments

FA Client or FA Proxy devices advertise I-SID-to-VLAN assignments to be supported on the FA Server. These assignments can be accepted or rejected by the FA Server. All communication between FA Proxies or Clients and the FA Server is using LLDP. Successful assignments result in the creation of a Switched UNI I-SIDs and endpoints based on the mapping requests.

The FA Server rejects I-SID-to-VLAN assignment requests if:

- FA is not configured properly on the port or MLT.
- Router IS-IS is disabled.



Note

For Fabric Attach to operate properly and for the FA Server to accept I-SID-to-VLAN assignment requests, IS-IS must be enabled.

The following error message is logged immediately after IS-IS is disabled, and displays the error message only once in the log file. The system does not display it again when an assignment request is made from the FA Proxy.

```
CP1 [12/04/15 00:33:49.733:UTC] 0x00374589 00000000
```

```
GlobalRouter FA INFO Fabric Attach Assignments will be rejected since ISIS is disabled.
```

- The C-VID and I-SID are not within the supported range.

The value 4095 is not supported. The value 4096 indicates that the port is untagged. An I-SID value of 0 is not supported on the FA Server.

- The I-SID is already assigned to an IP VPN.

The system displays the error message `I-SID is already assigned to an IPVPN`.

- The I-SID is already in use for SPB multicast.

The system displays the error message `SPB Multicast is enabled, ISID 16000000 and greater reserved for dynamic data-isid's used to carry Multicast traffic over SPB`.

- The I-SID has a value that is reserved for internal use.
- The I-SID cannot be used in an IS-IS accept policy.
- The I-SID is associated with a platform VLAN and that VLAN is used as a private VLAN (that is, has a secondary VLAN specified).
- The I-SID is already in use for Transparent Port UNI.
- The port that receives the I-SID-to-VLAN assignment is a member of an MLT, but FA is not successfully enabled on that MLT interface.
- There is a resource error on the FA Server system, such as lack of memory.
- The number of I-SID-to-VLAN assignments on a port exceeds the maximum limit which is 94.
- The number of I-SIDs on the switch exceeds the maximum limit.
- The same endpoint is configured on more than one I-SID.
- The port or MLT is associated with more than one C-VID in the same I-SID.

When the FA Server rejects I-SID-to-VLAN assignments, aside from viewing the log file, you can use trace to troubleshoot the cause of rejection.

For an example on troubleshooting rejection of I-SID-to-VLAN assignments on the FA Server and for more information on using trace, see [Troubleshooting Fabric Attach](#) on page 1368.

FA management

You can configure a management I-SID on an FA-enabled port or MLT. This I-SID includes an optional C-VID parameter, which is a VLAN ID that is locally significant to the port or MLT and does **not** represent a platform VLAN.

Depending on whether the C-VID value is specified, the behavior is as follows:

- If the C-VID value is specified, the FA Server transmits this VLAN ID as the management VLAN in the FA Element TLV. A client or proxy receiving this TLV uses this VLAN ID for management traffic on the FA Server uplink.
- If the C-VID value is not specified, the FA Server transmits a management VLAN with a VLAN ID value of 4095 in the FA Element TLV. A client or proxy receiving this TLV uses **untagged** traffic for network management on the FA Server uplink.

If you do not configure a management I-SID, the FA Server transmits a management VLAN ID value of 0 in the FA Element TLV. A client/proxy that receives the FA Element TLV retains the initial management configuration (if any) on its device.

Limitations of FA management I-SIDs

- A management I-SID value of 0 is not supported on the FA Server.
- You cannot enable BPDU on a management I-SID.

FA management configuration considerations

A Switched UNI I-SID that is created when an FA assignment is learned on a port or MLT, is uniquely identified by a tuple comprising of one of the combinations of (port, I-SID and C-VID) or (MLT ID, I-SID and C-VID). When you configure FA management, similar tuples are used. You can configure FA management on an FA-enabled port or MLT on which FA assignment mappings are learned, as long as the FA management tuple exactly matches the tuple created by the learned FA mapping.

The following scenarios describe the behavior when you configure FA management on a port or MLT that also receives learned FA mappings, but the tuples do not match.

- **Scenario 1:** You attempt to configure FA management on a port or MLT where an FA assignment mapping is already learned.

For example, consider an FA-enabled port 1/1 on which an assignment mapping is learned, with I-SID 100 and C-VID 20. You can configure FA management on port 1/1 as long as the I-SID and C-VID values exactly match that of the learned FA mapping. However, if you attempt to configure FA management on the port with a different I-SID and C-VID value, the configuration is not successful and an error message displays.

- **Scenario 2:** An FA assignment mapping is learned on a port or MLT that already has FA management configured.

For example, consider that FA management is configured on port 1/1. If an FA assignment mapping is learned on the port with the same I-SID and C-VID values as that of the FA management configuration, then the mapping is accepted. Otherwise the mapping is rejected.

FA message authentication and integrity protection

For the security of FA communication in terms of data integrity and authenticity, a keyed-hash message authentication code can be transmitted within every FA TLV.

It protects the I-SID-to-VLAN assignment exchanges between the FA Server and FA Proxy. The standard HMAC-SHA256 algorithm calculates the message authentication code (digest) involving a cryptographic hash function (SHA-256) in combination with a shared secret key. The key is symmetric, that is, it is known by both the source and destination parties.

By default, on the FA Server, message authentication is enabled at the interface level and a default key is defined to provide secure communication.

You can configure a different authentication key on an interface (port or MLT) on the FA Server, to authenticate a client on that interface. The authentication key is stored in encrypted form when you save configuration on the FA Server. For an FA Client to authenticate and attach to the FA Server, the authentication key must match on both

the client and the server. In general, the FA authentication key must match between two FA components exchanging FA TLVs through LLDP.

When you enable FA message authentication, the message authentication key (default or configured) generates a Hash-based Message Authentication Code (HMAC) digest that is included in FA I-SID-to-VLAN Assignment TLV. Upon receipt, the HMAC digest is recomputed for the TLV data and compared against the digest included in the TLV. If the digests are the same, the data is valid. If the digests are not the same, the data is considered invalid and is ignored.

The FA secure communication setting (enabled/disabled) and the symmetric key data are maintained across resets and restored during FA initialization.

Fabric Attach Ring with ISW-Series Managed Industrial Ethernet Switches

Table 137: Fabric Attach Ring with ISW-Series Managed Industrial Ethernet Switch product support

Feature	Product	Release introduced
Fabric Attach Ring Topology Change Notification	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0
	5420 Series	Fabric Engine 9.0
	5520 Series	Fabric Engine 9.0
	5720 Series	Fabric Engine 9.0
	7520 Series	Fabric Engine 9.0
	7720 Series	Fabric Engine 9.0
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.0
	VSP 7400 Series	VOSS 9.0

Fabric Attach interfaces that connect to ISW-Series Managed Industrial Ethernet Switches (ISW) deploy in a ring topology.

You can enable Fabric Attach (FA) Ring Topology Change Notification (TCN) processing on an FA interface that connects to an ISW-Series switch either manually or dynamically through Auto-sense.

For more information about Auto-sense, see [Auto-sense Port States](#) on page 28.

The ISW-Series switch at the end of the ring generates a TCN BPDU packet if the connection within the ring topology is interrupted, and again, after the connection is restored. The switch FA interface that connects to the ISW switch processes the TCN BPDU packet. All MACs learnt on the FA interface age out. TCN BPDUs propagate across the SPBM network only if you enable MSTP-Fabric Connect Multi Homing on the switch.

If switches at the end of the ring are not virtual interswitch trunk (vIST) peers, the best practice is to enable MSTP-Fabric Connect Multi Homing on all switches that

participate in the same Layer 2 VSN. Switches at the end of the ring that use virtual interswitch trunk (vIST) synchronize MAC addresses between the vIST peers. In this case, you do not need to enable MSTP-Fabric Connect Multi Homing.

For information about configuring MSTP-Fabric Connect Multi Homing, see [Enable MSTP-Fabric Connect Multi Homing](#) on page 1251.

For information about ISW-Series with ring configuration, see ISW-Series Managed Industrial Ethernet Switch documentation at <https://supportdocs.extremenetworks.com/support/documentation/>.

Fabric Attach and Switched UNI

With the C-VLAN UNI feature, I-SID-to-VLAN mappings must be unique across the network. With the Transparent Port UNI (ELAN Transparent) feature, you can map an entire port or MLT to an I-SID.

With the Switched UNI feature, you can associate many different C-VID/port or C-VID/MLT list combinations to a single I-SID.

Switched UNI and FA

FA brings the capability of automatically creating Switched UNI I-SIDs on a switch, without manual intervention. The I-SIDs thus created are ELAN I-SIDs with endpoints of type Switched UNI, and are by default for Layer 2. MAC learning takes place and there is an any-to-any relationship. For Layer 3 participation, you must configure a platform VLAN with the same I-SID value as that of the I-SID in a learned FA mapping.



Note

The number of Switched UNI I-SIDs created are different for different product families. For more information, see *Fabric Engine Release Notes*.

Limitations of FA-created Switched UNI I-SIDs

- An FA-created Switched UNI I-SID is always ELAN.
- You cannot enable BPDU on an FA-created Switched UNI I-SID.
- The ELAN I-SIDs created are by default for Layer 2. For Layer 3 participation, you must manually configure a platform VLAN with the same I-SID value as that of the I-SID in a learned FA mapping. You can configure the platform VLAN with the same VLAN ID as that of the C-VID, or use a different value.
- The Switched UNI (ELAN) I-SID cannot be used by IPVPN, MVPN, SPBM dynamic multicast range, or a T-UNI.
- You cannot change from one UNI type to another dynamically. The I-SID must be deleted and created with the new UNI type (Customer VLAN (C-VLAN), Transparent Port user-network-interface (T-UNI), ELAN).
- If the port is a member of an MLT, you must add the entire MLT to the C-VID.
- The port is always in the forwarding state.
- You cannot associate a port or MLT with more than one C-VID in the same I-SID.

- The same C-VID, port or MLT cannot be a member of more than one I-SID. The value 4095 is not supported and cannot be configured. The value 4096 indicates that the port is untagged.
- An I-SID value of 0 is not supported on the FA Server.

Fabric Attach Deployment Scenarios

Fabric Attach is typically deployed in the access layer(s) of a Fabric Connect network.

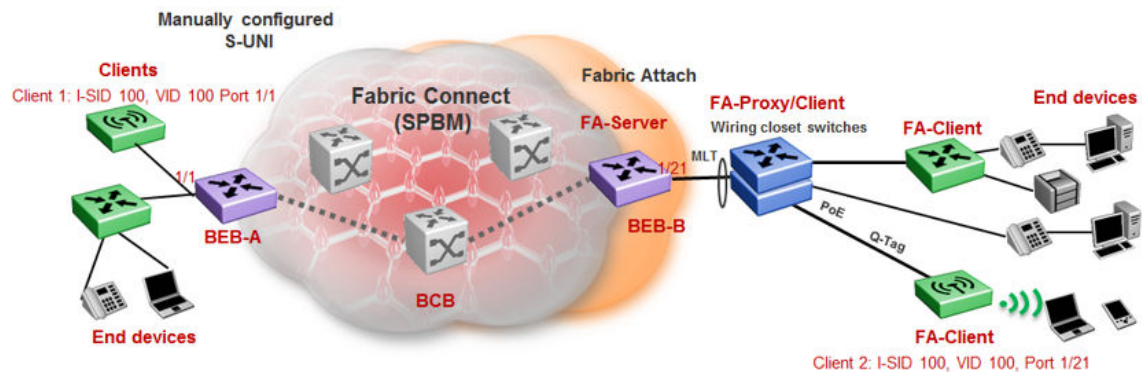
Fabric Attach, when used with a Fabric Connect solution, provides the same capabilities at the access layer, but those services and policies are mapped across the entire network end-to-end. FA makes user and end device attachment simple and creates network configuration and sets up resources only when needed.

An FA Server can be connected to FA Client or FA Proxy devices on three types of interfaces, namely, a port, MLT or an SMLT. The following sections discuss FA in SMLT and non-SMLT deployments.

FA and Switched UNI in non-SMLT Deployments

The following deployment shows an SPBM network in which one edge has manually configured Switched UNI I-SIDs and the other edge has Fabric Attach (FA). At the FA edge, the I-SIDs are learned using FA TLVs and are automatically created on the FA Server as ELAN I-SIDs with Switched UNI endpoints.

This deployment demonstrates that the FA-created I-SIDs can communicate with any other I-SID (manually created Switched UNI or a C-VLAN with an I-SID), on the local switch or across the SPBM fabric, as long as the I-SID values are the same.



BEB-B is a switch acting as the FA Server with a network-to-network interface (NNI) to the SPBM cloud. FA Client and FA Proxy devices send I-SID-to-VLAN mapping requests to the FA Server on the respective FA-enabled ports, using LLDP TLVs. This enables the I-SID endpoints to communicate with the SPB cloud.

If several clients are aggregated in an MLT, at least one of the ports must send the mapping requests for the FA Server to create the I-SID endpoints for that MLT. For example, let Client 2 be a wireless FA Client (such as a WLAN 9100 AP device) on port 1/21, that sends an FA mapping request for I-SID 100 and C-VID (VLAN ID) 100. The FA Server (BEB-B) creates the requested I-SID 100 on its device, and advertises it to the SPB cloud.

BEB-A has manually configured Switched UNI endpoints, one of which is Client 1 (connected at port 1/1) using the same I-SID value 100.

With this setup, data traffic can freely flow between Client 1 and Client 2 through the two BEBs and the BCB.

Thus the Switched UNI I-SIDs learned using FA TLVs on one edge of the Fabric Connect (SPBM) network can communicate with the manually created I-SIDs on the other edge, as long as they both have the same value.

FA and Switched UNI in SMLT Deployments

The following examples discuss FA in dual-homed and single-homed SMLT deployments.

The following section describes FA in a dual-homed SMLT deployment. A pair of switches that operate as IST peers act as the FA Server. An FA Proxy (typically a wiring closet switch or an access switch) is connected to FA Clients and in turn to end devices. The FA Clients or FA Proxies advertise I-SID-to-VLAN mappings namely the interface C-VID and the I-SID to the FA Server switches. Both switches receive the mapping information using LLDP TLVs. The switch that learns the mapping first from the LLDP TLV considers the I-SID endpoint to be discovered locally, and creates the I-SID on its device. It then sends the mapping information to its peer switch. When the peer switch receives the mapping across IST in a new SMLT message, it too creates the I-SID and endpoint on its device. This I-SID however, is considered to be discovered remotely, because the data was synchronized from its peer.



Note

- For the peer switches acting as the FA Server to transmit the same FA System ID (based on the virtual MAC), SMLT configuration must be the same on both peers.
- For successful FA operation, configuration of FA message authentication and the authentication key must be the same on both peers.
- For successful operation in Layer 3, a platform VLAN must be configured on both peers. This is necessary for proper MAC learning.




Figure 78: FA in a dual-homed SMLT deployment

In the previous example deployment, BEB-A and BEB-B are IST peers collectively acting as the FA Server. FA TLVs sent from the clients (through the proxy) are learned on

FA-enabled ports on BEB-A and BEB-B. When BEB-A learns the mapping for the first time on its port, it creates an I-SID on its device. This is considered locally discovered. In addition, it sends an SMLT message to its peer BEB-B, which also creates the I-SID on its device. This time, the I-SID is considered remotely discovered. Similarly, if BEB-B receives a mapping from a client for the first time, it creates an I-SID (locally discovered) and also sends an IST message to its peer to create an I-SID (remotely discovered).

Irrespective of whether the I-SID creation on the FA peers is triggered by a local TLV event or by messaging from the IST peer, they can both receive data traffic. Thus in a dual-homed SMLT deployment, any I-SID can be learned irrespective of whether it is discovered locally, discovered remotely or both.



Note

On the IST peers, if an FA TLV is learned on a port or normal MLT (instead of the admin SMLT), only the I-SID is sent to the peer switch.

In the single-homed SMLT, as shown in the following deployment, the FA Server creates either a locally discovered I-SID (if received from a client using FA TLVs) or a remotely discovered I-SID (if synchronized from its IST peer), but not both.

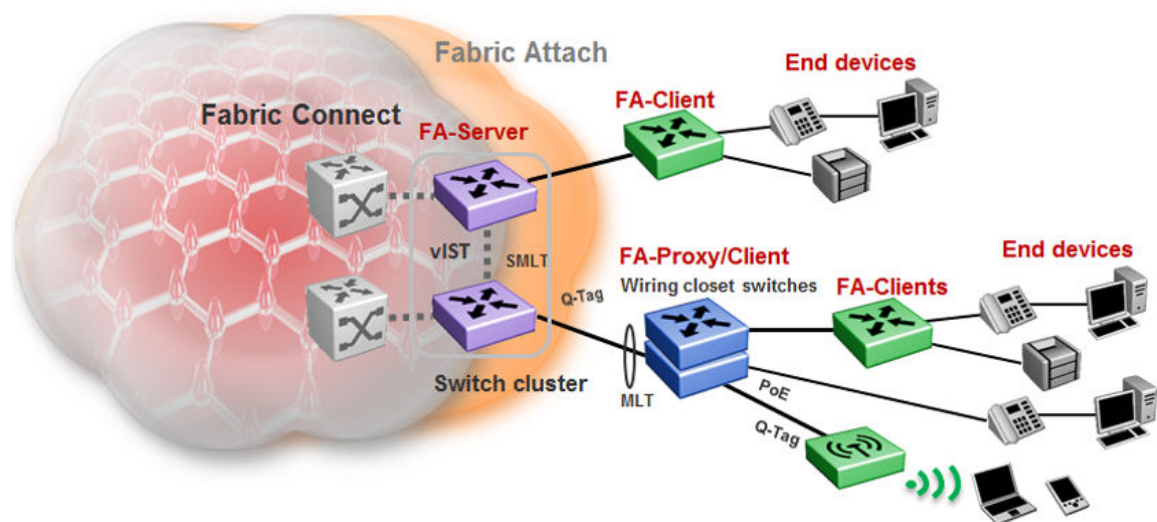


Figure 79: Fabric Attach in a single-homed SMLT deployment

Fabric Attach Considerations

Review the following restrictions, limitations, and behavioral characteristics for Fabric Attach.

- IS-IS and FA must be globally enabled on the FA Server, for FA to operate successfully.
- Static MAC, Static ARP and configuration of a static IGMP group are not supported on FA-enabled ports.
- An FA port cannot be a BROUTER port.
- You cannot enable FA on an existing Transparent Port UNI or a C-VLAN UNI port.

- FA I-SID-to-VLAN assignment mapping requests from a client or proxy device can be accepted or rejected by the FA Server.
- On an FA-enabled port or MLT, you must first disable LACP before you change the LACP key.
- You can only enable VLACP on an FA enabled MLT; VLACP is not supported on FA enabled non-MLT ports.
- On VLACP enabled ports, FA and LLDP signaling run independent of the VLACP state. Therefore, requests and responses are exchanged between the FA Server and client or proxy devices even if VLACP is operationally down. However, forwarding of data traffic is dependent on VLACP being operationally up on the port.

For example, if VLACP is enabled on the FA Server side of the link but not on the proxy or the client side, the FA Server learns the I-SID-to-VLAN assignment mappings and creates the required I-SIDs on its device. However, data traffic is not forwarded on the port until VLACP is operationally up.

- FA uses the virtual MAC to create the FA system ID when the FA is on an SMLT. If you delete the SPBM instance, then this information is no longer available. Therefore, you must delete the FA on SMLT before deleting the SPBM instance.
- You cannot enable FA and Endpoint Tracking simultaneously on the same interface.

Fabric Attach Configuration using the CLI

The following sections provide procedural information you can use to configure Fabric Attach (FA) and Logical Link Discovery Protocol (LLDP) using the Command Line Interface (CLI). For information about LLDP related to FA, see [Link Layer Discovery Protocol configuration using CLI](#) on page 2471.

Configure Fabric Attach Globally

For proper operation, FA must be enabled at both the global level and at the interface level on the FA Server. By default, FA is globally enabled. However, FA is disabled by default at the interface level and must be explicitly enabled on each interface.

Use this procedure to enable Fabric Attach globally on a switch.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable FA:
`fa enable`
3. (Optional) Disable FA:
`no fa enable [-y]`



Caution

Disabling FA flushes all FA element discovery and mappings.

4. View the FA configuration status. Use one of the following commands:

- `show fa`
- `show fa agent`

Example

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#fa enable

Switch:1>show fa

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm

Switch:1>show fa agent

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm
```

Variable Definitions

The following table defines parameters for the **no fa enable** command.

Variable	Value
<code>[-y]</code>	Suppresses the confirmation message.

Configure Fabric Attach Discovery Timeout

Use this procedure to configure the Fabric Attach discovery time-out.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```


2. Configure the FA discovery time-out in seconds:

```
fa discovery-timeout <45-480>
```



Note

The discovery time-out must be greater than or equal to the assignment time-out.

3. (Optional) Configure the default FA discovery time-out:

```
default fa discovery-timeout
```

Example

Configure the FA discovery time-out.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#fa discovery-timeout 50
```

Verify the configuration.

```
Switch:1(config)#show fa

=====
                        Fabric Attach Configuration
=====
                        FA Service : enabled
                        FA Element Type : server
                        FA Assignment Timeout : 45
                        FA Discovery Timeout : 50
                        FA Provision Mode : spbm
```

Variable definitions

The following table defines parameters for the **fa discovery-timeout** command.

Variable	Value
<45-480>	Specifies the Fabric Attach discovery time-out in seconds. The default value is 240 seconds.

Configure Fabric Attach Assignment Timeout

Use this procedure to configure the Fabric Attach assignment time-out.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the FA assignment time-out in seconds:

```
fa assignment-timeout <45-480>
```

**Note**

The assignment time-out must be less than or equal to the discovery time-out.

3. (Optional) Configure the default FA assignment time-out value:

```
default fa assignment-timeout
```

Example

Configure the FA assignment time-out:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#fa assignment-timeout 50
```

Verify the configuration:

```
Switch:1(config)#show fa

=====
                        Fabric Attach Configuration
=====
                        FA Service : enabled
                        FA Element Type : server
                        FA Assignment Timeout : 50
                        FA Discovery Timeout : 240
                        FA Provision Mode : spbm
```

Variable definitions

The following table defines parameters for the **fa assignment-timeout** command.

Variable	Value
<45-480>	Specifies the Fabric Attach assignment time-out in seconds. The default value is 240 seconds.

Enable Fabric Attach on an Interface

Use this procedure to enable Fabric Attach on an interface (port, static MLT or LACP MLT). Enabling FA on an MLT enables FA on all ports of the MLT. If your platform supports channelization, FA can also be enabled on channelized ports.

You can also enable FA Ring Topology Change Notification (TCN) processing on an interface that connects to an ISW-Series Managed Industrial Ethernet Switches (ISW) switch. For more information about FA TCN, see [Fabric Attach Ring with ISW-Series Managed Industrial Ethernet Switches](#) on page 1348.

Before You Begin

Verify that FA is enabled globally on the switch.

About This Task

Enabling FA on a port or MLT is necessary for element discovery.

On the FA Server, FA is enabled globally by default. However, you must explicitly enable FA on the desired port or MLT interface. FA is successfully enabled on an MLT only if all ports of the MLT have FA successfully enabled. Enabling FA automatically configures LLDP on all ports. Tagging is configured and spanning tree is disabled.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable FA on the interface:

```
fa enable
```



Note

You can use the **no** operator to disable FA on the interface. However, disabling FA flushes all FA element discovery and I-SID-to-VLAN mappings associated with the interface.

3. (Optional) Enable FA TCN processing on the interface that connects to ISW switches with ring:

```
fa tcn enable
```

4. View the FA configuration status:

```
show fa interface [disabled-auth] [enabled-auth] [mlt <1-512>] [port  
<{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}>]
```

Examples

Enable FA on a port:

```
Switch:1>en  
Switch:1#conf t  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#interface gigabitethernet 1/2  
Switch:1(config-if)#fa enable
```

Enable FA on an MLT:

```
Switch:1>en  
Switch:1#conf t  
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config)#interface mlt 2
Switch:1(config-mlt)#fa enable
```

Verify that FA is enabled on the interfaces.



Note

When FA is enabled, message authentication is enabled by default. The authentication key is set to the default value and displays encrypted on the output.

```
Switch:1>show fa interface
```

```
=====
                        Fabric Attach Interfaces
=====
```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	TCN	ORIGIN
Port 1/1	enabled	0	0	enabled	****		
Port1/10	enabled	0	0	enabled	****		
Port1/6	enabled	0	0	enabled	****		
Port1/11	enabled	0	0	enabled	****		
Mlt2	enabled	0	0	enabled	****		

```
-----
5 out of 5 Total Num of fabric attach interfaces displayed
-----
```

For example, disable FA on port 1/1.

```
Switch:1(config)#interface gigabitethernet 1/1
Switch:1(config-if)#no fa enable
Switch:1(config-if)#exit
```

Verify that FA is disabled on port 1/1.

```
Switch:1(config)#show fa interface
```

```
=====
                        Fabric Attach Interfaces
=====
```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	TCN	ORIGIN
Port 1/1	disabled	0	0	disabled	****		
Port1/10	enabled	0	0	enabled	****		
Port1/6	enabled	0	0	enabled	****		
Port1/11	enabled	0	0	enabled	****		
Mlt2	enabled	0	0	enabled	****		

```
-----
5 out of 5 Total Num of fabric attach interfaces displayed
-----
```

View the FA interfaces that have authentication enabled:

```
Switch:1(config)#show fa interface enabled-auth
```

```
=====
                        Fabric Attach Interfaces
=====
```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	TCN	ORIGIN
Port1/10	enabled	0	0	enabled	****		
Port1/6	enabled	0	0	enabled	****		
Port1/11	enabled	0	0	enabled	****		
Mlt2	enabled	0	0	enabled	****		
2 out of 2 Total Num of fabric attach interfaces displayed							

Optionally, disable FA message authentication on port 1/6.

```
Switch:1(config)#interface gigabitethernet 1/6
Switch:1(config-if)#no fa message-authentication
Switch:1(config-if)#exit
```

Verify that both FA and FA message authentication are disabled on port 1/6, as indicated by the SERVER STATUS and MSG AUTH STATUS fields respectively.

```
Switch:1(config)#show fa interface
```

Fabric Attach Interfaces							
INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	TCN	ORIGIN
Port 1/1	disabled	0	0	disabled	****		
Port1/10	enabled	0	0	enabled	****		
Port1/6	disabled	0	0	disabled	****		
Port1/11	enabled	0	0	enabled	****		
Mlt2	enabled	0	0	enabled	****		
4 out of 4 Total Num of fabric attach interfaces displayed							

View the FA interfaces that have authentication disabled:

```
Switch:1(config)#show fa interface disabled-auth
```

Fabric Attach Interfaces							
INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	TCN	ORIGIN
Port 1/1	disabled	0	0	disabled	****		
Port1/6	disabled	0	0	disabled	****		
2 out of 2 Total Num of fabric attach interfaces displayed							

Verify the FA interfaces that have TCN enabled.

```
Switch:1>show fa interface
```

Fabric Attach Interfaces							
INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	TCN	ORIGIN

```

Port 1/1    enabled  0      0      enabled  ****    enabled CONFIG
Port1/10    enabled  0      0      enabled  ****
Port1/6      enabled  0      0      enabled  ****
Port1/11    enabled  0      0      enabled  ****
Mlt2        enabled  0      0      enabled  ****    enabled CONFIG

```

```

-----
5 out of 5 Total Num of fabric attach interfaces displayed
-----

```

Verify the FA interfaces that have TCN is disabled when FA is disabled.

```

Switch:1>show fa interface

=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH MSG AUTH  TCN  ORIGIN
                STATUS ISID     CVID     STATUS   KEY
-----
Port 1/1        disabled 0        0        disabled ****    disabled CONFIG
Port1/10        enabled  0        0        enabled  ****
Port1/6          enabled  0        0        enabled  ****
Port1/11        enabled  0        0        enabled  ****
Mlt2            enabled  0        0        enabled  ****    enabled CONFIG

```

```

-----
5 out of 5 Total Num of fabric attach interfaces displayed
-----

```

Variable definitions

The following table defines parameters for the **show fa interface** command.

Variable	Value
disabled-auth	Displays the FA interfaces (port or MLT) that have authentication disabled.
enabled-auth	Displays the FA interfaces (port or MLT) that have authentication enabled.
<1-512>	The valid range for MLT ID. Displays FA configuration on the specified MLT interface.
port {slot/port[/sub-port] [-slot/port[/sub-port]][,...] [slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. Displays FA configuration on the specified port.

Configure FA Message Authentication on an Interface

Use this procedure to configure FA message authentication on an interface (port or MLT).

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure FA message authentication on a port or MLT:

```
[default] [no] fa message-authentication
```



Note

When FA is enabled, message authentication is enabled by default. The authentication key is set to the default value and displays encrypted on the output.

Example

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#
```

Enable message authentication on a port.

```
Switch:1(config)#interface gigabitEthernet 1/2

Switch:1(config-if)#fa message-authentication
Switch:1(config-if)#show fa interface port 1/2
```

```
=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER   MGMT     MGMT     MSG AUTH MSG AUTH  ORIGIN
                STATUS  ISID     CVID     STATUS   KEY
-----
Port1/2         enabled  0         0         enabled  ****

-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
Switch:1(config-if)#exit
Switch:1(config)#
```

Enable message authentication on an MLT.

```
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#fa message-authentication
Switch:1(config-mlt)#show fa interface mlt 10
```

```
=====
                          Fabric Attach Interfaces
=====
```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	ORIGIN
Mlt10	enabled	0	0	enabled	****	

```
-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
Switch:1(config-mlt)#exit
Switch:1(config)#
```

The following example demonstrates disabling message authentication on a port or MLT.

```
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#no fa message-authentication
Switch:1(config-if)#exit
Switch:1(config)
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#no fa message-authentication

Switch:1(config-mlt)#show fa interface
```

```
=====
                          Fabric Attach Interfaces
=====
```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	ORIGIN
Port1/2	enabled	0	0	disabled	****	
Mlt10	enabled	0	0	disabled	****	

```
-----
2 out of 2 Total Num of fabric attach interfaces displayed
-----
```

Configure the FA Authentication Key on an Interface

On the FA Server, you can configure an authentication key on an interface (port, static MLT or LACP MLT), to authenticate a client or proxy device on that interface. The authentication key is stored in encrypted form when you save configuration on the FA Server.

Before You Begin

Ensure that:

- On the FA Server, FA is enabled globally and also on the interface.

- FA message authentication is enabled on the interface.



Note

By default, enabling FA enables message authentication. The authentication key is set to the default value and the system displays the encrypted authentication key on the output.

About This Task

Use this procedure to configure an FA authentication key on a specified port or on all ports of an MLT, on the switch. If you do not configure an authentication key, the default value is used. If you specify a key, the default value is overridden and is stored in encrypted format in a separate file other than the configuration file, when you execute the `save config` command.



Caution

For an FA Client or an FA Proxy device to successfully authenticate and attach to the FA Server, the authentication key must match on both the client and the server. If the authentication key is changed on the FA Server switch, it must correspondingly be changed on the FA Client or Proxy attached to it, for FA to operate properly.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format `slot/port/sub-port`.

2. Configure the FA authentication key:

```
fa authentication-key WORD<0-32>
```

3. (Optional) Configure the default FA authentication key:

```
default fa authentication-key
```

Example

```
Switch:1>en  
Switch:1#conf t  
Enter configuration commands, one per line. End with CNTL/Z.
```

Enable FA and message authentication on a port. Configure the authentication key `phone-network` on the port.

```
Switch:1(config)#interface gigabitEthernet 1/2  
Switch:1(config-if)#fa enable
```

```
Switch:1(config-if)#fa message-authentication
Switch:1(config-mlt)#fa authentication-key phone-network
Switch:1(config-if)#exit
Switch:1(config)#
```

Enable FA and message authentication on an MLT. Configure the authentication key `client-network` on the MLT.

```
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#fa message-authentication
Switch:1(config-mlt)#fa authentication-key client-network
```

Verify configuration of the FA authentication key. The system displays the encrypted authentication key on the output.

```
Switch:1(config-if)#show fa interface

=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER   MGMT    MGMT    MSG AUTH  MSG AUTH  ORIGIN
                STATUS  ISID    CVID    STATUS    KEY
-----
Port1/2        enabled  0        0        enabled   ****
MLT10          enabled  0        0        enabled   ****

-----
2 out of 2 Total Num of fabric attach interfaces displayed
-----
```

Variable Definitions

The following table defines parameters for the **fa authentication-key** command.

Variable	Value
WORD<0-32>	Specifies the authentication key on the port or MLT.

Configure FA Management on a Port or MLT

Use this procedure to configure a management I-SID on a Fabric Attach (FA)-enabled port or MLT.

To configure an FA management I-SID for Auto-sense-enabled ports, see [Configure a Management I-SID for Auto-sense Fabric Attach Proxy Switches](#) on page 48.

Before You Begin

Ensure that the port or MLT is enabled for FA.

About This Task

This command applies to all traffic sent or received on a port or MLT, carrying the VLAN ID specified using the `c-vid` parameter. This parameter is optional.

Depending on whether the `c-vid` parameter is specified or not, the behavior is as follows:

- If you specify the `c-vid` parameter, the FA Server transmits this VLAN ID as the management VLAN in the FA Element TLV. A client or proxy receiving this TLV uses this VLAN-ID for management traffic on the FA Server uplink.
- If you do not specify the `c-vid` parameter, the FA Server transmits a management VLAN with a VLAN ID value of 4095 in the FA Element TLV. A client or proxy receiving this TLV uses untagged traffic for network management on the FA Server uplink.

An FA management I-SID can have a platform VLAN associated with it. For Layer 3 support on the management I-SID, you must create a platform VLAN by port and associate the platform VLAN with the management I-SID. The C-VID can be of the same value or of a different value than that of the platform VLAN.

If the management I-SID matches one of the FA Switched UNI (ELAN) I-SIDs (as displayed by the command `show i-sid elan`), then the platform VLAN is automatically associated with the FA-enabled interface (port or MLT).

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format `slot/port/sub-port`.

2. Configure the FA management I-SID:

```
fa management i-sid <1-16777215> c-vid <1-4094>
```



Important

If you do not specify a C-VID value, the port or MLT is untagged.

3. Verify configuration of FA management on the port or MLT, using the following commands:

- `show i-sid <1-16777215>`
- `show interfaces gigabitEthernet i-sid [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]`
- `show mlt i-sid [<1-512>]`

Examples

Configure FA management on port 1/2:

```
Switch:1>enable  
Switch:1#configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#interface gigabitethernet 1/2

Switch:1(config-if)#fa management i-sid 101 c-vid 101

Switch:1(config-if)#show i-sid 101

```
=====
                        Isid Info
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES  NAME
-----
101       ELAN        3         -         -         CONFIG     EXTRSERVER_101
```

The following example demonstrates the Origin as "auto-sense".

Switch:1(config-if)#show i-sid 500

```
=====
                        Isid Info
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES  NAME
-----
101       ELAN        -         c300:1/45 -         AUTO-SENSE ISID_500
```

Switch:1(config-if)#show interfaces gigabitEthernet i-sid

```
=====
                        PORT Isid Info
=====
PORTNUM   ISID      VLANID    C-VID    ISID      ORIGIN      ISID      MAC
IFINDEX  ID        C-VID    TYPE      NAME      BPDU      SUNI
-----
1/2       193       101       3         101       ELAN        MANAGEMENT EXTRSERVER_101
```

1 out of 1 Total Num of i-sid endpoints displayed

Configure FA management on MLT 10.

Switch:1>enable

Switch:1#configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#interface mlt 10

Switch:1(config-mlt)#fa management i-sid 101

Verify configuration of FA management on the MLT. Because the C-VID is not specified, the MLT displays as untagged.

Switch:1(config-mlt)#show i-sid 101

```
=====
                        Isid Info
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES  NAME
-----
101       ELAN        3         -         u:10        CONFIG     EXTRSERVER_101
```

In the following example, for Layer 3 support, create a platform VLAN 3 and associate it with the management I-SID 101.

Switch:1(config-if)#vlan create 3 type port-mstprstp 0

Switch:1(config-if)#vlan i-sid 3 101

Switch:1(config)#show i-sid

```
=====
                        Isid Info
=====
```

```
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES
-----
15999999  ELAN      4048      -         -           C --- - - - - Onboarding I-SID
16777001  ELAN      N/A       -         -           C --- - - - - FAN-ISID
=====
```

```
c: customer vid    u: untagged-traffic
```

```
All 2 out of 2 Total Num of i-sids displayed
```

```
ORIGIN Legend:
```

```
C: manually configured; D: discovered by FA or EPT
```

```
M: FA management; E: discovered by EAP; A: auto-sense; R: multi-area redist
```

```
l: discover by local switch  r: discover by remote VIST switch
```

```
Switch:1(config-if)#show vlan i-sid
```

```
=====
                                Vlan I-SID
=====
VLAN_ID    I-SID    VLAN NAME    I-SID NAME
-----
1           VLAN-1
2           VLAN-2
3           101      VLAN-3        EXTRSERVER_101
33          VLAN-33
999         VLAN-999
=====
```

Because the management I-SID matches one of the FA Switched UNI (ELAN) I-SIDs, the platform VLAN is automatically associated with the FA-enabled port 1/2.

```
Switch:1(config-if)#show interfaces gigabitEthernet i-sid
```

```
=====
                                PORT Isid Info
=====
ISID      ISID      ISID      ISID      MAC
PORTNUM  IFINDEX  ID        VLANID  C-VID  TYPE    ORIGIN    NAME      BPDU  SUNI
-----
1/2      193      101      3        101    ELAN    MANAGEMENT EXTRSERVER_101
=====
```

```
1 out of 1 Total Num of i-sid endpoints displayed
```

VLAN

Variable Definitions

The following table defines parameters for the **fa management** command.

Variable	Value
<i>i-sid</i> <1-16777215>	Specifies the management I-SID.
<1-4094>	Specifies the customer VLAN ID.
	Important: If you do not specify a C-VID value, the port or MLT is untagged .

Troubleshooting Fabric Attach

The following sections help you troubleshoot problems with Fabric Attach (FA).

Troubleshooting Workflow

Troubleshoot FA in the following sequence:

1. Verify FA configuration:

As a first step, for proper operation, verify that FA is enabled properly at both the global and interface levels. Use the procedures in this section to verify FA configuration.

2. Verify LLDP port-level transmission and reception:

LLDP operates at the interface level. Enabling FA at the port level automatically enables LLDP transmission and reception at the port level. Similarly, enabling FA at the MLT level automatically enables LLDP transmission and reception for all ports in that MLT. Use the procedures in [Link Layer Discovery Protocol](#) on page 2465 to verify LLDP interface (port or MLT) statistics.

3. Verify FA discovery, I-SID-to-VLAN mapping assignments, and Switched UNI I-SID creation:

After you verify LLDP transmission, verify that FA element discovery completed successfully. After a successful FA discovery, FA clients can send I-SID-to-VLAN mapping assignments to the FA Server on an FA-enabled port or MLT. The FA server accepts or rejects these mapping assignments. A prerequisite to successful mapping assignments is that IS-IS and SPBM are properly configured on the FA server. Successful FA mappings result in the creation of ELAN I-SIDs with end-points of type Switched UNI on the FA Server switch.

View Fabric Attach Global Configuration Status

Use this procedure to display the Fabric Attach global configuration status on a switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the FA configuration status using one of the following commands:
 - `show fa`
 - `show fa agent`

Example

Sample output for the **show fa** command:

```
Switch:1>show fa

=====
Fabric Attach Configuration
=====

FA Service : enabled
FA Element Type : server
FA Assignment Timeout : 240
FA Discovery Timeout : 240
FA Provision Mode : spbm
```

Sample output for the **show fa agent** command:

```
Switch:1>show fa agent

=====
                        Fabric Attach Configuration
=====
                        FA Service : enabled
                        FA Element Type : server
                        FA Assignment Timeout : 240
                        FA Discovery Timeout : 240
                        FA Provision Mode : spbm
```

View Fabric Attach Interface Configuration

Use this procedure to view FA interface configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View all FA interfaces (ports and MLTs):

```
show fa interface
```
3. To view FA interface configuration on ports, use one of the following commands:
 - View FA configuration on all ports:

```
show fa interface port
```
 - View FA configuration on a specific port, enter:

```
show fa interface port [{slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]]
```
4. To view FA interface configuration on MLTs, use one of the following commands:
 - View FA configuration on all MLTs:

```
show fa interface mlt
```
 - View FA configuration on a specific MLT:


```
show fa interface mlt [<1-512>]
```

Example

The following example displays sample outputs for the **show fa interface** command.

```
Switch:1>show fa interface

=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH  MSG AUTH  TCN  ORIGIN
                STATUS  ISID    CVID    STATUS    KEY
-----
Port 1/1       enabled  0        0        enabled  ****
Port1/10       enabled  0        0        enabled  ****
Port1/6        enabled  0        0        enabled  ****
Port1/11       enabled  0        0        enabled  ****
Mlt2           enabled  0        0        enabled  ****
-----
5 out of 5 Total Num of fabric attach interfaces displayed
```



The following is a sample output for the **show fa interface** command for the port 1/6.

```
Switch:1>show fa interface port 1/6

=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH MSG AUTH  TCN   ORIGIN
                STATUS ISID   CVID   STATUS   KEY
-----
Port1/6        enabled  0      0      enabled  ****

-----
1 out of 4 Total Num of fabric attach interfaces displayed
-----
```

The following is a sample output for the **show fa interface** command for the MLT 2.

```
Switch:1>show fa interface mlt 2

=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT   MGMT   MSG AUTH MSG AUTH  TCN   ORIGIN
                STATUS ISID   CVID   STATUS   KEY
-----
Mlt2           enabled  0      0      enabled  ****

-----
1 out of 4 Total Num of fabric attach interfaces displayed
-----
```

Variable definitions

The following table defines parameters for the `show fa interface port` command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<div>Identifies the slot and port in one of the following formats:</div> <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) <div>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</div>

The following table defines parameters for the `show fa interface mlt` command.

Variable	Value
<1-512>	The valid range for MLT ID.

Viewing Fabric Attach Discovered Elements

Use this procedure to view Fabric Attach discovered elements.

About This Task

When FA is enabled on an FA Server switch, LLDP PDUs are exchanged between the FA Server and FA Clients or FA Proxies. Standard LLDPs allow neighbors to be learned. With the help of organizational-specific element discovery TLVs, the client or proxy recognizes that it has attached to the FA Server. Only after the discovery handshake is complete, an FA Client or FA Proxy can transmit I-SID-to-VLAN assignments to join the SPB Fabric network through the FA Server.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display FA discovered elements:
show fa elements
3. Display FA discovered elements on a specific port:
show fa elements [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]
[slot/all][all]]

Example

The following example displays the sample output for the **show fa elements** command.

```
Switch:1>show fa elements
```

```
=====
Fabric Attach Discovery Elements
=====
```

PORT	TYPE	MGMT VLAN	STATE	SYSTEM ID	ELEM AUTH	ASGN AUTH
1/1	proxyRing	0	T / S	c8:66:5d:f6:d7:20:00:00:00:09	AP	N
1/5	proxy	710	T / S	50:61:84:ee:8c:00:20:00:00:01	AP	AP
1/6	proxy	710	T / S	50:61:84:ee:8c:00:20:00:00:01	AP	AP

```
=====
Fabric Attach Authentication Detail
=====
```

PORT	ELEM AUTH	OPER STATUS	ASGN AUTH	OPER STATUS
1/1	successAuth		successAuth	
1/5	successAuth		successAuth	
1/6	successAuth		successAuth	

```
State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid
```

```
Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None
```

```
-----
3 out of 3 Total Num of fabric attach discovery elements displayed
```

Variable definitions

The following table defines parameters for the **show fa elements** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]</code> <code>[slot/all][all]}</code>	<ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all)

View Fabric Attach I-SID-to-VLAN Assignments

About This Task

Use this procedure to display the I-SID-to-VLAN assignments advertised by an FA Client or an FA Proxy, to be supported on the FA Server. These assignments can be accepted or rejected by the FA Server. An assignment that is successfully accepted by the FA Server results in the creation of a Switched UNI I-SID on the interface.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display FA I-SID-to-VLAN assignments:
show fa assignment
3. Display FA I-SID-to-VLAN assignments on specific ports:
show fa assignment [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]
[slot/all][all]]}

Example

The following example displays a sample output for the `show fa assignment` command.



Note

The state of I-SID-to-VLAN assignments on a client or proxy device is pending until it is changed by the FA Server to `active` or `reject`.

```
Switch:>en
Switch:1#show fa assignment
=====
Fabric Attach Assignment Map
```

Interface	I-SID	Vlan	State	Origin
1/1	2	2	active	proxy
1/2	3	3	active	proxy
1/2	4	4	active	proxy
1/3	5	5	reject	proxy

4 out of 4 Total Num of fabric attach assignment mappings displayed

Variable definitions

The following table defines parameters for the **show fa assignment** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Troubleshoot FA Server Rejection of I-SID-to-VLAN Assignments Using Trace

Consider an FA solution where the FA Server receives I-SID-to-VLAN assignment requests from a proxy device and some of these assignment requests are rejected by the FA Server. Use this procedure to troubleshoot the cause of the rejection.



Note

When the FA Server rejects an I-SID-to-VLAN assignment request, the error message in the log file lists a generic reason for the failure, such as `rejected due to application error (status 9)`. To troubleshoot further, you must use the **trace** command.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Verify that router IS-IS is enabled. This is required for proper FA operation.
`show isis`



Note

I-SID-to-VLAN assignments are always rejected if router IS-IS is disabled.

3. Verify that FA is enabled on the interface on which I-SID-to-VLAN assignments are expected.

```
show fa interface [disabled-auth] [enabled-auth] [mlt <1-512>] [port
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]]}
```

4. Verify the discovery and authentication status of the proxy device, on the interface.

```
show fa elements [{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]
[slot/all] [all]]}
```

5. Determine the I-SID-to-VLAN assignments received on the interface and which ones are rejected.

```
show fa assignment [{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]
[slot/all] [all]]}
```

6. View the log file to determine the cause of the assignment rejection.

```
show log file module fa
```



Note

When the FA Server rejects an I-SID-to-VLAN assignment request, only a generic reason for the rejection is logged.

7. Configure trace:

- a. Enable keyword search in the trace output:

```
trace grep WORD<0-128>
```

- b. Configure the trace level for FA:

```
trace level <Module_ID>
```



Note

- <Module_ID> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see *Fabric Engine and VOSS Feature Matrix*. To see the list of module IDs, use the **show trace modid-list** command or CLI command completion Help.
- FA uses the trace level 221.

- c. Enable the display of the trace output:

```
trace screen enable
```

Example

The following example simulates a configuration error on the FA Server as a result of which the FA Server rejects I-SID-to-VLAN assignments from the proxy device. When the FA Server rejects an I-SID-to-VLAN assignment, the error message listed in the log file is a generic reason for the rejection, as demonstrated in this example. To troubleshoot further, set up trace.

On the FA Server, assume that the interface MLT 1 consists of ports 1/5 and 1/6. Assume that a proxy device sends I-SID-to-VLAN assignment mapping requests with I-SID 9005 and CVID 400, on this interface.

Simulate a configuration error on the FA Server:

Configure a management I-SID with a C-VID value that is different from that of the C-VID in the I-SID-to-VLAN assignment request from the proxy. So, for example, configure a management I-SID with C-VID 999, which is different from the C-VID advertised by the proxy, which is 400. This causes rejection of I-SID-to-VLAN assignment requests on the interface.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#no fa enable
Switch:1(config-mlt)#fa management i-sid 9005 c-vid 999
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#exit
Switch:1(config)#exit
```

At this stage, the FA Server rejects I-SID-to-VLAN assignments as shown in the following output.

```
Switch:1#show fa assignment

=====
                        Fabric Attach Assignment Map
=====
Interface  I-SID      Vlan      State      Origin
-----
1/5         312        710       active     proxy
1/5         9005       400       reject     proxy
1/6         312        710       active     proxy
1/6         9005       400       reject     proxy
-----
4 out of 4 Total Num of fabric attach assignment mappings displayed
=====
```

Begin troubleshooting on the FA Server:

Verify that IS-IS is enabled.

```
Switch:1>en
Switch:1#show isis

=====
                        ISIS General Info
=====

AdminState : enabled
RouterType : Level 1
System ID : 8404.bcb1.0043
Max LSP Gen Interval : 900
Metric : wide
Overload-on-startup : 20
Overload : false
Csnp Interval : 10
PSNP Interval : 2
Rxmt LSP Interval : 5
spf-delay : 100
Router Name : FAServer
ip source-address : 43.43.43.43
ipv6 source-address : 1:43:0:0:0:0:0:43
ip tunnel source-address : 12.43.43.43
Tunnel vrf : 12
```

```

        ip tunnel mtu :
        Num of Interfaces : 4
        Num of Area Addresses : 1
        inband-mgmt-ip :
            backbone : disabled
        Dynamically Learned Area : 00.0000.0000
        FAN Member : No
        Multi-Area OperState : disabled
        Hello Padding : enabled
Switch:1#

```

Verify that FA is enabled on the interface MLT 1, on which the I-SID-to-VLAN assignments are expected. View the SERVER STATUS field.

```

Switch:1#show fa interface mlt 1

=====
                        Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH MSG AUTH  ORIGIN
                STATUS ISID     CVID     STATUS  KEY
-----
Mlt1            enabled 0        0        enabled ****

-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----

```

Verify the discovery and authentication status of the proxy device on the interface. Note that the proxy is successfully discovered and authenticated on ports 1/5 and 1/6 of MLT 1.

```

Switch:1#show fa elements

=====
                        Fabric Attach Discovery Elements
=====
PORT  TYPE          MGMT      ELEM ASGN
      VLAN STATE  SYSTEM ID AUTH AUTH
-----
1/5   proxy       1    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP
1/6   proxy       1    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP

=====
                        Fabric Attach Authentication Detail
=====
PORT  ELEM OPER      ASGN OPER
      AUTH STATUS AUTH STATUS
-----
1/5   successAuth successAuth
1/6   successAuth successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,  N= None

```

```
-----
2 out of 2 Total Num of fabric attach discovery elements displayed
-----
```

View the log file to determine the cause of the rejection. The log file displays the generic error rejected due to application error (status 9) as follows:

```
Switch:1#show log file module fa
...
CP1 [12/04/15 00:45:51.185:UTC] 0x00374583 00000000 GlobalRouter FA INFO Fabric
Attach Element Discovered on interface 1/5 Element type proxy (3) Id
50:61:84:ee:8c:00:20:00:00:01 CP1 [12/04/15 00:45:51.187:UTC] 0x0037458f 00000000
GlobalRouter FA INFO Fabric Attach Assignment rejected: interface 1/5 i-sid 9005 cvid 400
rejected due to application error (status 9)
...
...
```

To troubleshoot further, use trace.

```
Switch:1#trace grep fa
Switch:1#trace level 221 3
Switch:1#trace screen enable
Screen tracing is on
```

View the trace output. The trace output displays that the error was caused because the FA interface (MLT 1) was configured with a different C-VID for I-SID 9005.

```
Switch:1#0:07:57.801252 1 fa.c :858 [lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUni :FA: faUpdateSwitchedUni port 196 isid 9005 cvid 400
0:07:57.801283 1 fa_swuni.c :2900[lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUniCheck :FA: Call faUpdateSwitchedUniCheckSmlt for mlt 1
0:07:57.801644 1 fa_swuni.c :2421[lcy-ve]
[12898-13062]cbcp-main.x:faSwitchedUniCheckEndpointParms:FA: Failed
rcIsidElanEndPointTblConsistencyCheckCommon for Ifindex 6144 Isid 9005 Cvid 400 error
Switched UNI/Fabric Attach MLT cannot be configured for different c-vid for same I-SID
0:07:57.802074 1 fa.c :858 [lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUni :FA: faUpdateSwitchedUni port 197 isid 9005 cvid 400
0:07:57.802086 1 fa_swuni.c :2900[lcy-ve][12898-13062]cbcp-
main.x:faUpdateSwitchedUniCheck :FA: Call faUpdateSwitchedUniCheckSmlt for mlt 1
0:07:57.802276 1 fa_swuni.c :2421[lcy-ve]
[12898-13062]cbcp-main.x:faSwitchedUniCheckEndpointParms:FA: Failed
rcIsidElanEndPointTblConsistencyCheckCommon for Ifindex 6144 Isid 9005 Cvid 400 error
Switched UNI/Fabric Attach MLT cannot be configured for different c-vid for same I-SID"
```

View Fabric Attach Statistics

If FA discovery fails, use this procedure to display FA statistics to determine if FA discovery TLVs were processed. You can also view the FA assignment statistics to determine the number of FA assignments that were accepted or rejected by the FA Server.

You can view the statistics at either the global level or at the port (interface) level.

Procedure

1. Enter Privileged EXEC mode:
enable

2. View global level FA statistics:

```
show fa statistics [summary]
```

3. View FA statistics at the slot/port level:

```
show fa statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...]  
[slot/all][all]]]
```

**Note**

If a slot is removed from the switch chassis, the FA statistics are not displayed on the slot ports. When the slot is inserted back again, the statistics counters are reset.

4. (Optional) Clear FA statistics:

```
clear fa statistics [summary] [{slot/port[/sub-port]}[-slot/port[/sub-  
port]][,...][slot/all][all]]]
```

Examples

Viewing FA discovery and assignment statistics:

```
Switch:1>en
Switch:1#show fa statistics
```

```
=====
                        Fabric Attach STATISTICS
=====
```

Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	3057	0	1	0
1/2	2000	0	1	0

```
=====
                        Fabric Attach ASSIGNMENTS STATISTICS
=====
```

Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	3149	3	1	3	0	0
1/2	1500	0	1	2	0	0

View a summary of the FA discovery and assignment statistics:

```
Switch:1#show fa statistics summary
```

```
=====
                        Fabric Attach STATISTICS SUMMARY
=====
```

Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	3057	0	1	0
1/2	2000	0	1	0

```
=====
                        Fabric Attach ASSIGNMENTS STATISTICS SUMMARY
=====
```

Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	3149	3	1	3	0	0
1/2	1500	0	1	2	0	0

1/1	3149	3	1	3	0	0
1/2	1500	0	1	2	0	0

Viewing FA statistics on a specific port (port 1/1):

```
Switch:1>en
Switch:1#show fa statistics 1/1
```

```
=====
                        Fabric Attach STATISTICS
=====
```

Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	3057	0	1	0

```
=====
                        Fabric Attach ASSIGNMENTS STATISTICS
=====
```

Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	3149	3	1	3	0	0

Optionally, clear FA statistics and verify that the statistics are cleared.

```
Switch:1#clear fa statistics
Switch:1#show fa statistics
```

```
=====
                        Fabric Attach STATISTICS
=====
```

Port	DiscElem Received	DiscElem Expired	DiscElem Deleted	DiscAuth Failed
1/1	0	0	0	0
1/2	0	0	0	0

```
=====
                        Fabric Attach ASSIGNMENTS STATISTICS
=====
```

Port	Asgn Received	Asgn Accepted	Asgn Rejected	Asgn Expired	Asgn Deleted	AsgnAuth Failed
1/1	0	0	0	0	0	0
1/2	0	0	0	0	0	0

Variable Definitions

The following table defines parameters for the **show fa statistics** command.

Variable	Value
summary	Displays a summary of Fabric Attach element discovery and assignment statistics at the global level.
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display Learned LLDP Neighbors

Use this procedure to verify details of the LLDP neighbors learned.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Verify details of LLDP neighbors learned:
show lldp neighbor
3. Verify details of LLDP neighbors learned on a specific port:
show lldp neighbor port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}

Variable Definitions

The following table defines parameters for the **show lldp neighbor** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <p>Displays LLDP neighbor information on the specified port.</p>

Display Switched UNI (ELAN) I-SID Information

Use this procedure to display information on FA-created Switched UNI (ELAN) I-SIDs.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display all Switched UNI (ELAN) I-SIDs:
show i-sid elan
3. Display ELAN I-SID information on an MLT:
show mlt i-sid [<1-512>]

**Note**

Viewing ELAN I-SID information on an MLT is useful to understand the origin of the I-SID when multiple client or proxy devices connecting to the FA Server using SMLT MLT advertise the same I-SID-to-VLAN mappings. In the event of a link failure on an MLT, the origin of the I-SID helps determine on which MLT, and thereby from which proxy or client device, the mappings were successfully learnt.

4. Display ELAN I-SID information on ports:

```
show interfaces gigabitEthernet i-sid [{slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]]
```

Examples

Display information on all Switched UNI (ELAN) I-SIDs.

The following sample output displays, for example, the I-SID information on one of the peer switches of the FA Server, in a dual-homed SMLT configuration.

```
Switch:1>enable
Switch:1#show i-sid elan

=====
Isid Info
=====
ISID      ISID      VLANID    PORT      MLT        ORIGIN      ISID
ID        TYPE                               INTERFACES INTERFACES
-----
2002      ELAN      N/A       c2002:1/10 -          - - - - -1- - EXTRSERVER_1
4000      ELAN      N/A       -          c4000:1    - - - - --r - EXTRSERVER_12
4001      ELAN      N/A       -          c4001:1    - - - - -1- - EXTRSERVER_101
4030      ELAN      N/A       -          c4030:1    - - - - --r - EXTRSERVER_102
4051      ELAN      N/A       -          c4051:1    - - - - -1- - EXTRSERVER_103
10200     ELAN      N/A       -          c200:1     - - - - --r - EXTRSERVER_2

c: customer vid    u: untagged-traffic

All 6 out of 6 Total Num of Elan i-sids displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
```

l: discover by local switch r: discover by remote VIST switch



Note

The I-SID TYPE field displays once for each I-SID. The I-SID TYPE of an I-SID that is either learned through FA mapping assignments or configured as an FA management I-SID, is always ELAN. If a platform VLAN has the same I-SID value as that of the I-SID in an FA mapping assignment or in an FA management I-SID configuration, then the platform VLAN is associated with the I-SID endpoint and displays in the VLANID column.

Display MLT I-SID information for MLT 1.

In this sample output, the ORIGIN field indicates the origin of the I-SID endpoint.

```
Switch:1>show mlt i-sid
=====
MLT Isid Info
=====
```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU
3	6146	3	N/A	33	ELAN	C --- - --- -	ISID-3	

```
-----
1 out of 1 Total Num of i-sid endpoints displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch   r: discover by remote VIST switch
```

Display I-SID information on the port 1/10:

In this sample output, the ORIGIN field indicates the origin of the I-SID endpoint.

```
Switch:1#show interfaces gigabitEthernet i-sid
=====
PORT Isid Info
=====
```

PORTNUM	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU	MAC SUNI
1/1	192	27	N/A	4000	ELAN	C --- - --- -	ISID-27		FALSE
1/1	192	270	N/A	4001	ELAN	C --- - --- -	ISID-270		FALSE
1/1	192	309	N/A	309	ELAN	C --- - --- -	ISID-309		FALSE
1/1	192	401	N/A	401	ELAN	C --- - --- -	ISID-401		FALSE
1/1	192	1001	N/A	1001	ELAN	C --- - --- -	ISID-1001		FALSE
1/1	192	1111	N/A	1111	ELAN	C --- - --- -	ISID-1111		FALSE
1/1	192	1121	N/A	1121	ELAN	C --- - --- -	ISID-1121		FALSE
1/1	192	1201	N/A	1201	ELAN	C --- - --- -	ISID-1201		FALSE
1/1	192	2001	N/A	2001	ELAN	C --- - --- -	ISID-2001		FALSE
1/2	193	38	N/A	4000	ELAN	C --- - --- -	ISID-38		FALSE
1/2	193	310	N/A	310	ELAN	C --- - --- -	ISID-310		FALSE
1/2	193	380	N/A	4001	ELAN	C --- - --- -	ISID-380		FALSE
1/2	193	402	N/A	402	ELAN	C --- - --- -	ISID-402		FALSE

```
-----
13 out of 152 Total Num of i-sid endpoints displayed

ORIGIN Legend:
```

C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch r: discover by remote VIST switch

Variable Definitions

The following table defines parameters for the **show i-sid** command.

Variable	Value
elan	Displays all ELAN I-SIDs.

The following table defines parameters for the **show mlt i-sid** command.

Variable	Value
<1-512>	The valid range for MLT ID.

The following table defines parameters for the **show interfaces gigabitEthernet i-sid** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Enable or disable FA Zero Touch Client Attachment

Use this procedure to enable or disable the global FA Zero Touch Client Attachment feature on an FA Proxy or Server. By default, FA Zero Touch Client Attachment support is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable an FA Zero Touch client:

```
fa zero-touch-client standard <camera|phone|router|security-device|  
srvr-endpt|switch|video|virtual-switch|wap-type1|wap-type2> i-sid <1-  
15999999>
```

3. Disable an FA Zero Touch client:

```
no fa zero-touch-client standard <camera|phone|router|security-device|
srvr-endpt|switch|video|virtual-switch|wap-type1|wap-type2>
```

Example

```
Switch:1(config)# fa zero-touch-client standard camera i-sid 1003
Switch:1(config)# no fa zero-touch-client standard camera
```

Variable definitions

The following table defines parameters for the **fa zero-touch-client standard** command.

Variable	Value
<i>camera</i>	Specify element type to match camera.
<i>phone</i>	Specify element type to match phone.
<i>router</i>	Specify element type to match router.
<i>security-device</i>	Specify element type to match security-device.
<i>srvr-endpt</i>	Specify element type to match srvr-endpt.
<i>switch</i>	Specify element type to match switch.
<i>video</i>	Specify element type to match video.
<i>virtual-switch</i>	Specify element type to match virtual-switch.
<i>wap-type1</i>	Specify element type to match wap-type1.
<i>wap-type2</i>	Specify element type to match wap-type2.

Display FA Zero Touch Client Attachment

Use this procedure to display the Zero Touch Client Attachment data you have configured on an FA Server.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display Zero Touch Client Attachment data:
show fa zero-touch-client

Example

The following example displays sample output for the **show fa zero-touch-client** command.

```
Switch:1#show fa zero-touch-client

=====
                        Fabric Attach Zero Touch Client
=====
Type      Description      I-SID      VLAN      I-SID Name
```

```

-----
6      wap-type1      11111      123
11     camera        2000      200
-----

```

```

-----
2 out of 2 Total Num of Fabric Attach Zero Touch Client entries displayed
-----

```

Fabric Attach configuration using EDM

The following sections provide procedural information you can use to configure Fabric Attach (FA) and Logical Link Discovery Protocol (LLDP) using EDM. For information about LLDP related to FA, see [Link Layer Discovery Protocol configuration using EDM](#) on page 2494.

Configure Fabric Attach Globally

Use this procedure to configure FA globally or view existing FA global configuration.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**.
3. Select the **Globals** tab.
4. To enable or disable the Fabric Attach service, select **enabled** or **disabled** in the **Service** field.



Caution

Disabling FA flushes all FA element discovery and mappings.

5. View the element type in the **ElementType** field.



Note

The only supported element type is **faServer** (FA Server).

6. To specify the assignment time-out, enter a time-out value in seconds in the **AsgnTimeout** field.
7. View the provision mode in the **ProvisionMode** field.



Note

The supported provision mode is **spbm**.

8. To specify the discovery time-out, enter a time-out value in seconds in the **DiscTimeout** field.
9. To clear the FA statistics, select the **Clear FA Statistics** checkbox.
10. To clear the error counters, select the check boxes **ClearErrorCounters** and/or **ClearGlobalErrorCounters**.
11. Select **Apply**.

Fabric Attach Globals Field Descriptions

Use the data in the following table to use the **Fabric Attach Globals** tab.

Name	Description
Service	Enables or disables Fabric Attach service globally. The default is enable.
ElementType	Specifies the Fabric Attach element type. The supported element type is Fabric Attach Server.
AsgnTimeout	Specifies the Fabric Attach assignment time-out in seconds. The range is 45 to 480 seconds. The default is 240 seconds.
ProvisionMode	Specifies the Fabric Attach provision mode. The supported provision mode is SPB.
DiscTimeout	Specifies the Fabric Attach discovery time-out in seconds. The range is 45 to 480 seconds. The default is 240 seconds.
Clear FA Statistics	Clears Fabric Attach statistics.
ClearGlobalErrorCounters	Clears Fabric Attach global error counters. Disabled by default.

Configure Fabric Attach I-SID-to-VLAN Assignments

Use this procedure to view or configure FA I-SID-to-VLAN assignment information.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**.
3. Select the **Assignment** tab.
4. If you make configuration changes, select **Apply** to save changes.

Assignments Field Descriptions

Use the data in the following table to use the **Assignments** tab.

Name	Description
Ifindex	Specifies the interface identifier of the I-SID-to-VLAN assignment.
Isid	Specifies the I-SID value of the I-SID-to-VLAN assignment.
Vlan	Specifies the VLAN ID component of the I-SID-to-VLAN assignment.
State	Specifies the current state of the I-SID-to-VLAN assignment. It can be one of the following values: <ul style="list-style-type: none"> • Other • Pending • Active • Rejected

Name	Description
Origin	Specifies the origin information of the I-SID-to-VLAN assignment.
Isid name	Specifies the I-SID name.

Configure Fabric Attach Interface-level Settings

Use this procedure to configure FA interface-level settings or view existing interface-level settings.

You can enable Fabric Attach on a port, static MLT or an LACP MLT. Enabling FA on a port not only enables tagging but also disables spanning tree on that port. Enabling FA on an MLT enables FA on all ports of the MLT. When FA is enabled on ports in an MLT or LACP MLT, tagging is enabled and spanning tree is disabled on all those ports.

Before You Begin

Ensure that FA is enabled globally on the switch.

About This Task

Enabling FA on a port or MLT is necessary for element discovery. On the FA Server, FA is enabled globally by default. However, you must explicitly enable FA on a desired port or MLT interface, following which the FA Server can begin transmitting LLDP PDUs that contain the element discovery TLVs. This information is received by FA Client and FA Proxy devices, which in turn also transmit their FA capabilities and settings. After the element handshake completes, the FA Server receives I-SID-to-VLAN assignment mappings from the connected client or proxy devices, on that port or MLT.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**.
3. Select the **Ports** tab.

The FA interface-level settings are displayed.

4. Select the fields on this window to modify existing settings.
5. Select **Apply** to save your changes.
6. To configure FA on a new port or MLT interface:
 - a. Select **Insert**.

The **Insert Ports** dialog box opens.
 - b. To configure FA on a port, enter a port number in the format slot/port[/sub-port], or select **Port** to select from a list of available ports.

- c. To configure FA on an MLT, enter an MLT ID or select **Mlt** to select from a list of configured MLTs.



Note

FA is successfully enabled on the MLT, only if all ports of the MLT have FA successfully enabled. Enabling FA enables LLDP on all ports. Tagging is enabled and spanning tree is disabled.

- d. Select **Insert** to save your changes.
7. To remove (delete) FA on a port or MLT:
 - a. In the content pane, select a port or MLT from the list.
 - b. Select **Delete**.



Caution

Removing FA on an interface flushes all FA element discovery and I-SID-to-VLAN mappings associated with that interface.

Ports Field Descriptions

Use the data in the following table to use the **Ports** tab.

Name	Description
IfIndex	Specifies the interface (port or MLT) on which Fabric Attach is configured.
State	Specifies the current state of the Fabric Attach port. It is either enabled or disabled. This field indicates whether LLDP PDUs (that include FA TLVs) are generated on the port (enabled) or not (disabled).
MsgAuthStatus	Specifies the Fabric Attach message authentication status on the port. It is either enabled or disabled.
MsgAuthKey	Specifies the Fabric Attach message authentication key for the associated port. The maximum length of this key is 32 characters.
MgmtIsid	Specifies the Fabric Attach management I-SID for the associated port. The range is 0 to 16777215. A zero value indicates that the management I-SID is not specified for the interface.
MgmtCvid	Specifies the Fabric Attach management customer VLAN ID (C-VID) for the interface. A zero value indicates that no C-VID is specified for the interface. Using the maximum configuration value for your switch indicates the port is untagged. Platform support determines the C-VID range.

Name	Description
TCNEnable	Specifies whether the associated Fabric Attach interface can receive Topology Change Notification (TCN) BPDUs from ISW-Series Managed Industrial Ethernet Switches.
Origin	Specifies the origin of Fabric Attach port, either manually configured through CLI or EDM, or dynamically configured through Auto-sense.

View Fabric Attach discovered elements

Use this procedure to view discovered Fabric Attach elements.

About This Task

When FA is enabled on an FA Server switch, LLDP PDUs are exchanged between the FA Server and FA Clients or Proxies. Standard LLDPs allow neighbors to be learned. In addition, organizational specific element discovery TLVs allow the Client or Proxy to recognize that it has attached to an FA Server. Only after the discovery handshake is complete, an FA Client or Proxy can transmit I-SID-to-VLAN assignments to join the SPB Fabric through the FA Server.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**.
3. In the content pane, select the **Elements** tab.

Elements Field Descriptions

Use the data in the following table to use the **Elements** tab.

Name	Description
IfIndex	Specifies the interface (port or MLT) at which the Fabric Attach element was discovered.
ElementType	Specifies the element type of the discovered Fabric Attach element, as advertised using LLDP. The supported element type is the Fabric Attach Server.
ElementVlan	Specifies the VLAN ID of the discovered Fabric Attach element, as advertised using LLDP.
ElementId	Specifies the system ID of the discovered Fabric Attach element, as advertised using LLDP.
ElementState	Specifies the state flag data associated with the discovered Fabric Attach element, as advertised using LLDP.
ElementOperAuthStatus	Specifies the authentication status of the discovered Fabric Attach element.
ElementAsgnsOperAuthStatus	Specifies the authentication status of remote assignments.

Name	Description
ElementAuth	Specifies the discovered element authentication status.
AsgnsAuth	Specifies the assignment authentication status.

View FA statistics

Use this procedure to view FA statistics.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**.
3. In the content pane, select the **Stats** tab.

Stats Field Descriptions

Use the data in the following table to use the **Stats** tab.

Name	Description
PortIndex	Specifies the port for which the FA statistics are displayed.
DiscElemReceived	Specifies the number of element discoveries received on the port.
AsgnReceived	Specifies the number of remote assignments received on the port.
AsgnAccepted	Specifies the number of remote assignments accepted on the port.
AsgnRejected	Specifies the number of remote assignments rejected on the port.
AsgnExpired	Specifies the number of remote assignments that have expired, on the port.
AuthFailed	Specifies the number of authentications that have failed on the port.
DiscElemExpired	Specifies the number of discovery elements that have expired on the port.
DiscElemDeleted	Specifies the number of discovery elements that are deleted on the port.
AsgnDeleted	Specifies the number of remote assignments deleted on the port.
AsgnAuthFailed	Specifies the number of remote assignment authentications that failed on the port.

View Global FA Statistics Graphically

Use this procedure to view the global FA statistics graphically.

Procedure

1. In the navigation pane, expand **Configuration > Graph**.

2. Select **Chassis**.
3. Select the **Fabric Attach** tab.
4. To view a graphical representation of the statistics, select a row and select the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
5. Select **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
6. Select **Export**, to export the statistical data to a file.
7. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Fabric Attach Field Descriptions

Use the data in the following table to use the **Fabric Attach** tab.

Name	Description
DiscElemReceived	Specifies the number of discovery elements received globally.
AsgnReceived	Specifies the number of remote I-SID-to-VLAN assignments received globally.
AsgnAccepted	Specifies the number of remote I-SID-to-VLAN assignments accepted globally.
AsgnRejected	Specifies the number of remote I-SID-to-VLAN assignments rejected globally.
AsgnExpired	Specifies the number of remote I-SID-to-VLAN assignments that expired globally.
AuthFailed	Specifies the number of authentications that failed globally.
DiscAuthFailed	Specifies the number of discovery authentications that failed globally.
DiscElemExpired	Specifies the number of discovery elements that expired globally.
DiscElemDeleted	Specifies the number of discovery elements that were deleted globally.
AsgnDeleted	Specifies the number of remote assignments that were deleted globally.

View FA Port Statistics Graphically

Use this procedure to view the FA port statistics graphically.

Before You Begin

Ensure that a switch port is selected in the **Device Physical View** tab.

Procedure

1. In the navigation pane, expand **Graph > Port**.
2. Select the **Fabric Attach** tab.
The FA port statistics are displayed.

3. To view a graphical representation of the port statistics, select a row and select the appropriate icon on the top left-hand-side of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
4. Select **Clear Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
5. Select **Export**, to export the statistical data to a file.
6. To fix a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Fabric Attach Field Descriptions

Use the data in the following table to use the **Fabric Attach** tab.

Name	Description
DiscElemReceived	Specifies the number of discovery elements received on a given port.
AsgnReceived	Specifies the number of remote I-SID-to-VLAN assignments received on a given port.
AsgnAccepted	Specifies the number of remote I-SID-to-VLAN assignments accepted on a given port.
AsgnRejected	Specifies the number of remote I-SID-to-VLAN assignments rejected on a given port.
AsgnExpired	Specifies the number of remote I-SID-to-VLAN assignments that expired on a given port.
AuthFailed	Indicates the number of received TLVs for which authentication was attempted and failed on the identified port.
DiscElemExpired	Specifies the number of discovery elements that expired on a given port.
DiscElemDeleted	Specifies the number of discovery elements that were deleted on a given port.
AsgnDeleted	Specifies the number of remote assignments that were deleted on a given port.
AsgnAuthFailed	Specifies the number of remote assignment authentications that failed on a given port.

Insert a Zero Touch Client

Use this procedure to insert a FA Zero Touch Client.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**.
3. Select the **Zero Touch Client Auto Attach** tab.
4. Select **Insert**.
The **Insert Zero Touch Client** dialog box opens.
5. In the **Type** field, select the ellipsis and select a client.
6. Select **Ok** to select the client or **Refresh** to update the list.

7. In the **Isid** field, enter the I-SID value.
The I-SID value is between 0 and 16777214.
8. Select **Insert**.

Configure FA Zero Touch Client Auto Attach

Use this procedure to configure FA Zero Touch Client auto attach.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Fabric Attach**
3. Select the **Zero Touch Client Auto Attach** tab.
From the Zero Touch Client Auto Attach tab you can configure a number of auto attach settings.
4. Select **Insert**.
5. In the **Type** field, select the ellipsis and select a client.
6. Select **Ok** to select the client or **Refresh** to update the list.
7. In the **Isid** field enter the I-SID value.
8. Select **Insert**.
9. (Optional) To **Delete** a FA Zero Touch client select it from the auto attach table and select **Delete**.

Zero Touch Client Auto Attach Field Descriptions

Use the data in the following table to use the **Zero Touch Client Auto Attach** tab

Field	Description
Type	This column describes the type of client assigned to auto attach. Available FA client types are: <ul style="list-style-type: none">• Wireless AP (Type 1)• Wireless AP (Type 2)• Switch• Router• IP Phone• IP Camera• IP Video• Security Device• Virtual Switch• Server Endpoint
Vlan	Specifies the VLAN ID component of the I-SID-to-VLAN assignment.
Isid	Specifies the I-SID value of the I-SID-to-VLAN assignment.

Fabric Attach configuration examples

This section provides configuration examples to configure Fabric Attach.

Configure a Fabric Attach Solution

The following section describes a simple configuration example to configure Fabric Attach (FA) at the edge of a Fabric Connect network. This is a typical deployment at its simplest level and is powerful because of its use in conjunction with a Fabric Connect core.

Before You Begin

Configure SPBM and IS-IS on the edge and core switches. For more information, see [Configure Minimum SPBM and IS-IS Parameters](#) on page 1221.

About This Task

Configuring FA primarily consists of configuring the FA Server. The FA Server in turn discovers neighboring FA component devices (like the FA Proxies and FA Clients) using FA TLVs within the LLDP PDUs.

In the following deployment, the switch at the edge of the Fabric Connect cloud is configured as the FA Server. On this switch, FA is enabled globally and at the interface (port) level. Another switch, functioning as the FA Proxy connects to the FA enabled port (1/3) on the FA Server. User A is an end user device that needs to send and receive data traffic from User B (another end user device) across the network.

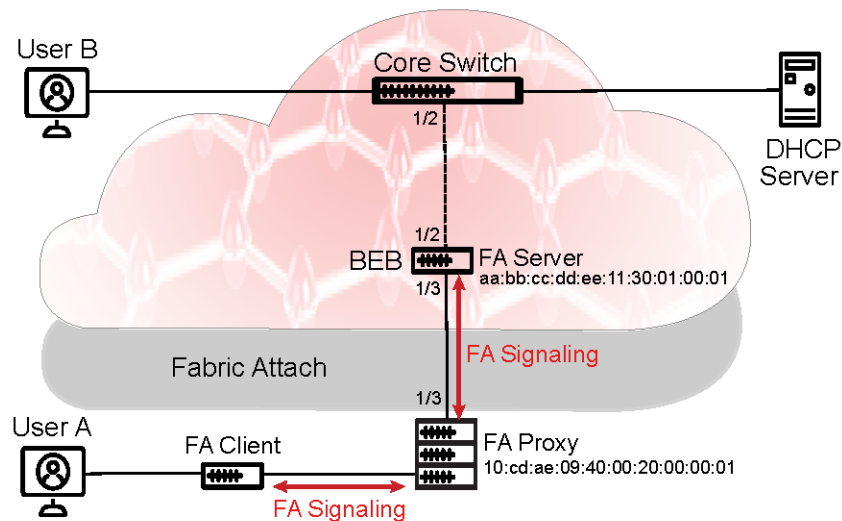


Figure 80: Fabric Attach Solution

Procedure

1. Configure the edge switch (BEB) as the FA Server:

- a. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

- b. Enable FA globally:

```
fa enable
```

- c. Enter port interface configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-  
port]][,...][slot/all][all]}
```

- d. Enable FA on the port:

```
fa enable
```



Note

Enabling FA automatically enables message authentication. Also, the authentication key is set to the default value and the system displays the encrypted authentication key on the output.



Note

Enabling FA on a port not only enables tagging but also disables spanning tree on that port.

2. Verify global configuration of FA using one of the following commands:

- `show fa`
- `show fa agent`

3. Verify interface level configuration of FA:

```
show fa interface
```

4. Verify the discovery of clients attaching to the FA Server:

```
show fa elements
```

5. Display the FA I-SID-to-VLAN assignments:

```
show fa assignment
```

To verify I-SID-to-VLAN assignments on a specific port, enter:

```
show fa assignment {slot/port[/sub-port] [-slot/port[/sub-port]][,...]  
[slot/all][all]}
```

6. Verify creation of Switched UNI (ELAN) I-SIDs:

```
show i-sid elan
```

Example

SPBM and IS-IS configuration on the core and edge switches:

SPBM configuration:

```
Switch:1>en  
Switch:1#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#spbm
Switch:1(config)#spbm ethertype 0x8100
```

IS-IS SPBM configuration:

```
Switch:1(config)#router isis
Switch:1(config)#spbm 1
Switch:1(config-isis)#spbm 1 nick-name 1.00.01
Switch:1(config-isis)#spbm 1 b-vid 41-42 primary 41
Switch:1(config-isis)#spbm 1 ip enable
```

IS-IS router configuration:

```
Switch:1(config-isis)#router isis
Switch:1(config-isis)#sys-name BEB-Switch
Switch:1(config-isis)#ip-source-address 3.3.3.3
Switch:1(config-isis)#is-type ll
Switch:1(config-isis)#system-id 0001.0001.0001
Switch:1(config-isis)#manual-area c0.2000.000.00
Switch:1(config-isis)#exit
```

Interface (port-level) configuration

```
Switch:1(config)#interface GigabitEthernet 1/2
Switch:1(config-if)#no shutdown
Switch:1(config-if)#isis
Switch:1(config-if)#isis spbm 1
Switch:1(config-if)#isis enable
Switch:1(config-isis)#exit
Switch(config)#vlan create 41 type spbm-vlan
Switch(config)#vlan create 42 type spbm-vlan
Switch(config)#router isis enable
Switch(config)#show isis spbm
```

Configuration of the edge switch as the FA Server

Enable FA globally.

```
Switch:1(config)#fa enable
Switch:1(config)#show fa

=====
                          Fabric Attach Configuration
=====
                          FA Service : enabled
                          FA Element Type : server
                          FA Assignment Timeout : 240
                          FA Discovery Timeout : 240
                          FA Provision Mode : spbm
```

Enable FA on the port.

Enabling FA automatically enables message authentication. The authentication key is configured with the default value, which the system displays in encrypted format in the output.

```
Switch:1(config)#int gigabitEthernet 1/3
Switch:1(config-if)#fa enable
Switch:1(config-if)#show fa interface port 1/3
```

```

=====
Fabric Attach Interfaces
=====
INTERFACE      SERVER  MGMT    MGMT    MSG AUTH MSG AUTH  ORIGIN
STATUS        ISID    CVID    STATUS  KEY
-----
Port1/3        enabled 0        0        enabled ****

1 out of 1 Total Num of fabric attach interfaces displayed

Switch:1(config-if)#exit
Switch:1(config)#exit

```

Verify that the FA Proxy is discovered by the FA Server.

```

Switch:1(config)#show fa elements

=====
Fabric Attach Discovery Elements
=====
PORT  TYPE          MGMT          ELEM ASGN
VLAN  STATE  SYSTEM ID    AUTH AUTH
-----
1/3   proxy          2    T / S  10:cd:ae:09:40:00:20:00:00:01  AP  AP

Fabric Attach Authentication Detail
=====
ELEM OPER          ASGN OPER
PORT  AUTH STATUS      AUTH STATUS
-----
1/3   successAuth      successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,   N= None

2 out of 2 Total Num of fabric attach discovery elements displayed

```

Verify the FA I-SID-to-VLAN assignment. An active state indicates that the FA (ELAN) I-SID is successfully created with endpoint of type Switched UNI. By default, this I-SID is created for Layer 2.

```

Switch:1#show fa assignment

=====
Fabric Attach Assignment Map
=====
Interface  I-SID    Vlan    State    Origin
-----
1/3        44        2        active    proxy

1 out of 1 Total Num of fabric attach assignment mappings displayed

```

For Layer 3 support, you must configure a platform VLAN. The platform VLAN can have the same value as that of the C-VID or it can have a different value.

In this example, the platform VLAN has the same value as the C-VID.

```
Switch:1(config)#vlan create 2 type port-mstprstp 0
Switch:1(config)#vlan i-sid 2 44
Switch:1#show i-sid elan
```

```
=====
                                Isid Info
=====
```

ISID	ISID	PORT	MLT	ORIGIN
ID	TYPE	VLANID	INTERFACES	INTERFACES
NAME				
44	ELAN	2	c2:1/3	DISC_LOCAL ISID-44

```
c: customer vid      u: untagged-traffic

All 1 out of 1 Total Num of Elan i-sids displayed
```

Verify neighbor discovery on the FA Proxy switch:

Note that the edge switch (BEB) is discovered as the FA Server by the FA Proxy.

```
Switch:1(config)#show fa agent
```

```
Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Zero Touch Status: Enabled
Fabric Attach Auto Provision Setting: Proxy
Fabric Attach Provision Mode: SPBM
Fabric Attach Client Proxy Status: Enabled
Fabric Attach Standalone Proxy Status: Disabled
Fabric Attach Agent Timeout: 50 seconds
Fabric Attach Extended Logging Status: Enabled
Fabric Attach Primary Server Id: aa:bb:cc:dd:ee:11:30:01:00:01 (SPBM)
Fabric Attach Primary Server Descr:BEB-Switch (6.0.0.0_GA)
```

```
Switch:1(config)#show fa elements
```

Unit/ Port	Element Type	Element Subtype	Element VLAN	Auth	System ID
1/3	Server	Server (Auth)	0	AP	aa:bb:cc:dd:ee:11:30:01:00:01

```
Switch:1(config)#show fa i-sid
```

I-SID	VLAN	Source	Status
44	2	Proxy	Active

Configure Fabric Attach in an SMLT

This example describes FA configuration and behavior in a dual-homed SMLT deployment.

The following figure shows a simple FA solution in a dual-homed SMLT deployment. In this deployment, a pair of BEB switches (BEB A and BEB B) operating as IST peers are

configured as the FA Server. An access switch or a wiring closet switch, configured as an FA Proxy, connects to the FA Server.

The FA Proxy advertises I-SID-to-VLAN assignment mappings to the FA Server. Both BEB switches receive the mapping information using LLDP PDUs containing assignment TLVs. The switch that learns the mapping first considers the I-SID to be discovered locally and creates the I-SID on its device. The mapping information is then shared with its IST peer switch. When the peer switch receives the mapping across IST in a new SMLT message, it too creates the I-SID on its device. This I-SID however, is considered to be discovered remotely because it is learnt from synchronization with the peer switch. The mappings can also be learned on the FA Server from both LLDP PDUs and from IST synchronization.




Figure 81: FA configuration in dual-homed SMLT

Before You Begin

Ensure that the proxy device (for example, an access switch) is properly configured for FA. See the corresponding product documentation for information on how to configure FA on the switch.

Procedure

1. Configure SMLT and vIST on switches BEB A and BEB B.



Caution

For the IST peer switches acting as the FA Server to transmit the same FA System ID (based on the virtual MAC), SMLT configuration on both the switches must be the same.

For detailed information on configuring SMLT and vIST, see [MultiLink Trunking and Split MultiLink Trunking](#) on page 2657.

Configure BEB A and BEB B as the FA Server:

Perform the following configuration on each switch.

2. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

3. Enable FA globally:

```
fa enable
```

4. Enter MLT interface configuration mode:

```
interface mlt <1-512>
```

5. Enable FA on the MLT:

```
fa enable
```

**Note**

Enabling FA automatically enables message authentication. Also, the authentication key is set to the default value and the system displays the encrypted authentication key on the output.

6. (Optional) Configure an FA authentication key with a value different from that of the default value:

```
fa authentication-key [WORD<0-32>]
```

**Caution**

When you configure the FA authentication key, you must configure the same value on both BEB switches in the SMLT.

Verify global and MLT-level FA configuration on BEB A and BEB B:

7. Verify global configuration of FA using one of the following commands:

- `show fa`
- `show fa agent`

8. Verify MLT-level (interface-level) configuration of FA:

```
show fa interface
```

Verify FA discovery on BEB A and BEB B:

9. Verify discovery of the FA Proxy.

```
show fa elements
```

View FA I-SID-to-VLAN assignments on BEB A and BEB B:

10. View the FA I-SID-to-VLAN assignments:

```
show fa assignment
```

To view FA I-SID-to-VLAN assignments on specific ports, enter:

```
show fa assignment {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]  
[slot/all] [all]}
```

Verify creation of Switched UNI I-SIDs on BEB A and BEB B:

11. Verify creation of Switched UNI (ELAN) I-SIDs:

- View ELAN I-SID information using `show i-sid elan`.
- View ELAN I-SID information on a specific MLT using `show mlt i-sid [<1-512>]`.



Note

Viewing ELAN I-SID information on an MLT is very useful to understand the origin of the I-SID, when multiple client or proxy devices connecting to the FA Server using SMLT MLT advertise the same I-SID-to-VLAN mappings. In the event of a link failure on an MLT, the origin of the I-SID helps determine on which MLT, and thereby from which proxy or client device, the mappings were successfully learnt.

Examples

SMLT configuration on BEB A and BEB B

Perform the same configuration on both BEB A and BEB B:

```
Switch:1>enable
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface mlt 1
Switch:1(config)#smlt
```

vIST configuration on BEB A and BEB B

On BEB A:

```
Switch:1(config)#vlan create 2261 type port-mstprstp 0
Switch:1(config)#vlan i-sid 2261 1502261
Switch:1(config)#interface vlan 2261
Switch:1(config)#ip address 192.0.2.0 255.255.255.0 2
```

Configure BEB B (IP address 192.0.2.1) as the IST peer.

```
Switch:1(config)#virtual-ist peer-ip 192.0.2.1 vlan 2261
```

On BEB B:

```
Switch:1(config)#vlan create 2261 type port-mstprstp 0
Switch:1(config)#vlan i-sid 2261 1502261
Switch:1(config)#interface vlan 2261
Switch:1(config)#ip address 192.0.2.1 255.255.255.0 2
```

Configure BEB A (IP address 192.0.2.1) as the IST peer.

```
Switch:1(config)#virtual-ist peer-ip 192.0.2.1 vlan 2261
```

FA configuration on BEB A

Enable FA globally and on the MLT:

```
Switch:1(config)#fa enable
Switch:1>show fa
```

```
=====
```

```

Fabric Attach Configuration
=====
FA Service : enabled
FA Element Type : server
FA Assignment Timeout : 240
FA Discovery Timeout : 240
FA Provision Mode : spbm

```

Optionally, configure an FA authentication key with the value `dual-homed-smlt`. Ensure that you configure the same value on both switches BEB A and BEB B.

```

Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#fa authentication-key dual-homed-smlt

```

Enable FA on the MLT:

```

Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#exit
Switch:1(config)#show fa interface

```

```

=====
Fabric Attach Interfaces
=====

```

INTERFACE	SERVER STATUS	MGMT ISID	MGMT CVID	MSG AUTH STATUS	MSG AUTH KEY	ORIGIN
Mlt1	enabled	0	0	enabled	****	

```

-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----

```

Verify discovery of the FA Proxy:

```

Switch:1(config)#show fa elements

```

```

=====
Fabric Attach Discovery Elements
=====

```

PORT	TYPE	MGMT VLAN	STATE	SYSTEM ID	ELEM AUTH	ASGN AUTH
1/21	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP
1/22	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP
1/23	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP

```

=====
Fabric Attach Authentication Detail
=====

```

PORT	ELEM AUTH	OPER STATUS	ASGN AUTH	OPER STATUS
1/21	successAuth		successAuth	
1/22	successAuth		successAuth	
1/23	successAuth		successAuth	

State Legend: (Tagging/AutoConfig)
T= Tagged, U= Untagged, D= Disabled, S= Spbm, V= Vlan, I= Invalid

Auth Legend:
AP= Authentication Pass, AF= Authentication Fail,
NA= Not Authenticated, N= None


```
-----
3 out of 3 Total Num of fabric attach discovery elements displayed
```

The FA Proxy advertises I-SID-to-VLAN assignment mappings to BEB A, on MLT ports 1/21 to 1/23. View the FA I-SID-to-VLAN assignments on BEB-A:

All ports in the MLT receive the FA assignment mappings, as shown in the following output.

```
Switch:1(config)#show fa assignment

=====
Fabric Attach Assignment Map
=====
Interface  I-SID      Vlan      State      Origin
-----
1/21       2          2         active     proxy
1/21       3          3         active     proxy
1/21       4          4         active     proxy
1/22       2          2         active     proxy
1/22       3          3         active     proxy
1/22       4          4         active     proxy
1/23       2          2         active     proxy
1/23       3          3         active     proxy
1/23       4          4         active     proxy
```

FA configuration on BEB B

Enable FA globally and on the MLT:

```
Switch:1(config)#fa enable
Switch:1(config)#show fa

=====
Fabric Attach Configuration
=====
FA Service : enabled
FA Element Type : server
FA Assignment Timeout : 240
FA Discovery Timeout : 240
FA Provision Mode : spbm
```

Configure the FA authentication key dual-homed-smlt. Ensure that you configure the same value as on BEB A.

```
Switch:1(config)#interface mlt 1
Switch:1(config-mlt)#fa authentication-key dual-homed-smlt
```

Enable FA on the MLT:

```
Switch:1(config-mlt)#fa enable
Switch:1(config-mlt)#exit
Switch:1(config)#show fa interface

=====
Fabric Attach Interfaces
=====
INTERFACE  SERVER  MGMT  MGMT  MSG AUTH  MSG AUTH  ORIGIN
          STATUS ISID  CVID  STATUS   KEY
-----
```

```
Mlt1          enabled 0          0          enabled ****
```

```
-----
1 out of 1 Total Num of fabric attach interfaces displayed
-----
```

Verify discovery of FA Proxy:

```
Switch:1(config)#show fa elements
```

```
=====
                        Fabric Attach Discovery Elements
=====
```

PORT	TYPE	MGMT		SYSTEM ID	ELEM ASGN	
		VLAN	STATE		AUTH	AUTH
1/17	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP
1/18	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP
1/19	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP
1/20	proxy	2	T / S	10:cd:ae:09:40:00:20:00:00:01	AP	AP

```
=====
                        Fabric Attach Authentication Detail
=====
```

PORT	ELEM OPER		ASGN OPER	
	AUTH	STATUS	AUTH	STATUS
1/17	successAuth		successAuth	
1/18	successAuth		successAuth	
1/19	successAuth		successAuth	
1/20	successAuth		successAuth	

```
-----
State Legend: (Tagging/AutoConfig)
T= Tagged,      U= Untagged,      D= Disabled,      S= Spbm,      V= Vlan,      I= Invalid

Auth Legend:
AP= Authentication Pass,  AF= Authentication Fail,
NA= Not Authenticated,   N= None

-----
4 out of 4 Total Num of fabric attach discovery elements displayed
```

The FA Proxy device advertises I-SID-to-VLAN assignment mapping requests to BEB B on MLT ports 1/17 to 1/20.

View FA I-SID-to-VLAN assignments on BEB-B:

```
Switch:1(config)#show fa assignment 1/17
```

```
=====
                        Fabric Attach Assignment Map
=====
```

Interface	I-SID	Vlan	State	Origin
1/17	2	2	active	proxy
1/17	3	3	active	proxy
1/17	4	4	active	proxy

Verify the creation of FA Switched UNI (ELAN) I-SIDs on BEB A and BEB B.

On BEB A:

```
Switch:1(config)#show i-sid elan
=====
Isid Info
=====
```

ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN	ISID NAME
2	ELAN	N/A	-	c2:1	- --- - -lr -	ISID-2
3	ELAN	N/A	-	c3:1	- --- - -lr -	ISID-3
4	ELAN	N/A	-	c4:1	- --- - -lr -	ISID-4

```

=====
ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:1(config-if)#shlw mlt i-sid

```

View the I-SID information for MLT1 on BEB A.

```
Switch:1(config)#show mlt i-sid 1
=====
MLT Isid Info
=====
```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU
1	6144	2	N/A	2	ELAN	- --- - -lr -	ISID-2	
1	6144	3	N/A	3	ELAN	- --- - -lr -	ISID-3	
1	6144	4	N/A	4	ELAN	- --- - -lr -	ISID-4	

```

=====
3 out of 3 Total Num of i-sid endpoints displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:1(config-if)#shlw mlt i-sid

```

On BEB B:

```
Switch:1(config)#show i-sid elan
=====
Isid Info
=====
```

ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN	ISID NAME
2	ELAN	N/A	-	c2:1	- --- - -lr -	ISID-2
3	ELAN	N/A	-	c3:1	- --- - -lr -	ISID-3
4	ELAN	N/A	-	c4:1	- --- - -lr -	ISID-4

```

=====
ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:1(config-if)#shlw mlt i-sid

```

View the I-SID information for MLT1 on BEB B.

```
Switch:1(config)#show mlt i-sid 1
```

```
=====
```

MLT Isid Info									
MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN		ISID NAME	BPDU
1	6144	2	N/A	2	ELAN	- --- -	-lr -	ISID-2	
1	6144	3	N/A	3	ELAN	- --- -	-lr -	ISID-3	
1	6144	4	N/A	4	ELAN	- --- -	-lr -	ISID-4	

```
=====
```

3 out of 3 Total Num of i-sid endpoints displayed

ORIGIN Legend:
 C: manually configured; D: discovered by FA or EPT
 M: FA management; E: discovered by EAP; A: auto-sense
 l: discover by local switch r: discover by remote VIST switch

```
8k_fanout:1(config-if)#shlw mlt i-sid
```

The following section describes the behavior if, for example, a link failure occurs between the FA Proxy and BEB B, as shown in the following figure.




Figure 82: FA behavior in dual-homed SMLT during a link failure

View the I-SID-to-VLAN assignments on BEB A:

```
Switch:1(config)#show fa assignment 1/21
```

```
=====
```

Fabric Attach Assignment Map				
Interface	I-SID	Vlan	State	Origin
1/21	2	2	active	proxy
1/21	3	3	active	proxy
1/21	4	4	active	proxy

```
=====
```

View the Switched UNI (ELAN) I-SIDs created on BEB A.

```
Switch:1(config)#show i-sid elan
```

```
=====
```

Isid Info									
ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN	ISID NAME			
2	ELAN	N/A	-	c2:1	- --- -	-1- - ISID-2			
3	ELAN	N/A	-	c3:1	- --- -	-1- - ISID-3			
4	ELAN	N/A	-	c4:1	- --- -	-1- - ISID-4			

```
=====
```

```

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:1(config-if)#shlw mlt i-sid

```

View the I-SID information for MLT 1 on BEB A.

```

Switch:1(config)#show mlt i-sid 1

=====
MLT Isid Info
=====

```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU
1	6144	2	N/A	2	ELAN	- --- - -l-	ISID-2	
1	6144	3	N/A	3	ELAN	- --- - -l-	ISID-3	
1	6144	4	N/A	4	ELAN	- --- - -l-	ISID-4	

```

-----
3 out of 3 Total Num of i-sid endpoints displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:1(config-if)#shlw mlt i-sid

```

View the Switched UNI (ELAN) I-SIDs created on BEB B.

```

BEB-B:1(config-mlt)#show i-sid elan

=====
Isid Info
=====

```

ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN	ISID NAME
2	ELAN	N/A	-	c2:1	- --- - --r	ISID-2
3	ELAN	N/A	-	c3:1	- --- - --r	ISID-3
4	ELAN	N/A	-	c4:1	- --- - --r	ISID-4

```

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:1(config-if)#shlw mlt i-sid

```

View the I-SID information for MLT 1 on BEB B.

```

Switch:1(config)#show mlt i-sid 1

=====
MLT Isid Info
=====

```

MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU
1	6144	2	N/A	2	ELAN	- --- - --r	ISID-2	
1	6144	3	N/A	3	ELAN	- --- - --r	ISID-3	
1	6144	4	N/A	4	ELAN	- --- - --r	ISID-4	

```

-----
3 out of 3 Total Num of i-sid endpoints displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch  r: discover by remote VIST switch
8k_fanout:l(config-if)#shlw mlt i-sid

```

Endpoint Tracking

Table 138: Endpoint Tracking product support

Feature	Product	Release introduced
Endpoint Tracking	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.7
	VSP 7400 Series	VOSS 8.1.1

Endpoint Tracking Overview

Endpoint Tracking provides dynamic assignment of virtual machines (VMs) to IP subnets as they attach to a Shortest Path Bridging (SPB) cloud. Deployment scenarios include VMs connecting to DvR Leaf nodes, or regular SPBM deployments.

ExtremeCloud IQ Site Engine is integral to the Endpoint Tracking solution. ExtremeCloud IQ Site Engine delivers automation; there is no need to manually configure server VLANs on data center access switches. Additionally, ExtremeCloud IQ Site Engine provides the ability to see what VM MACs exist, and where they are located.

ExtremeCloud IQ Site Engine's ExtremeConnect module integrates with third-party virtualization software (such as VMware or Microsoft HyperV) and communicates with the ExtremeControl module to automatically extract all of the VM MACs (including VLAN assignment for each MAC) and then automatically create all of the necessary authentication profiles, rules and mappings.

When the switch detects a new VM on a port, it sends a RADIUS request to ExtremeCloud IQ Site Engine. ExtremeConnect checks with VCenter for the Port Group, VLAN ID, and I-SID information that corresponds with the VM, communicates with the ExtremeControl module for the RADIUS authentication, and sends the RADIUS

response back to the switch with the VLAN:ISID binding information. Based on the binding, the switch then automatically creates a dynamic Switched UNI (S-UNI). Dynamic S-UNIs are not saved into the configuration file.

Typical Endpoint Tracking Implementation Example

The following example shows a typical implementation of Endpoint Tracking and the dynamic I-SID assignment process, as provisioned in ExtremeCloud IQ Site Engine.




Figure 83: Endpoint Tracking Example

The sequence within and among the four example VLANs in this configuration is as follows:

1. The RADIUS server authenticates VM1, and the switch automatically creates a Switched UNI with VLAN 10 and I-SID 10 binding, (using the outbound attributes received from the RADIUS server). Subsequently, the server authenticates VM2, which uses the same Switched UNI.
2. Similarly, on the other side of the SPB cloud, the RADIUS server authenticates VM5 and the switch automatically creates a Switched UNI with VLAN 30 and I-SID 10 binding, (using the outbound attributes received from the RADIUS server). Subsequently, the server authenticates VM6, which uses the same Switched UNI.
3. The same sequence occurs for VMs 3 and 4, and PCs 7 and 8, with the first authentication in each VLAN providing the outbound RADIUS attributes needed for the creation of a Switched UNI for that VLAN.
4. The final result is that VMs 1, 2, 5, and 6 can access each other on I-SID 10, and VMs 3, 4, 7, and 8 can access each other on I-SID 20.

Static S-UNIs and Visibility Mode

Endpoint Tracking can also be used in cases where static S-UNIs are configured on Endpoint Tracking-enabled ports. In this case, the MACs are allowed by default on the static S-UNI. However, by default, the MACs learned on a static S-UNI are not learned at the Endpoint Tracking level. Endpoint Tracking Visibility Mode allows tracking of MACs that are learned on static S-UNIs. This implies that a binding is created for these MACs,

but these bindings do not create dynamic S-UNIs, they are used for tracking purposes only.

Interface Support

Endpoint Tracking is supported on Ethernet ports, MLTs, and SMLTs.

If the switch is a Virtual IST (vIST) peer, the dynamic Switched UNI is synchronized to its vIST peer as follows:

- If the MAC is learned on an SMLT UNI interface, all Switched UNI information is synchronized to the vIST peer.
- If the MAC is learned on a non-SMLT UNI interface, only the I-SID is synchronized to the vIST peer.

VM Moves and VLAN:ISID Bindings

When a VM moves to a new switch within a network (with no change to the VLAN segment), the new switch triggers a new RADIUS authentication, which points that VM MAC to the new switch, and new bindings are applied on the new switch. The old switch detects that the VM MAC is moved and automatically deletes the old binding, if the old binding has not already aged out.

However, if a VM remains attached to the same (previously authenticated) switch, but the VLAN segment is changed, you must push a reauthentication request from ExtremeCloud IQ Site Engine to force the required binding updates. For more information about managing binding updates using RADIUS Change-of-Authorization (CoA) functionality, see [ExtremeCloud IQ Site Engine Integration](#) on page 1410.

Operational Considerations

Consider the following when implementing Endpoint Tracking:

- A RADIUS server used for Endpoint Tracking provides authorization only; no accounting processes are supported. Although accounting is enabled by default for all RADIUS servers, it is not currently supported for use with Endpoint Tracking, even if left enabled.
- Fabric Attach is not supported on ports or MLT/SMLTs that have Endpoint Tracking enabled.

ExtremeCloud IQ Site Engine Integration

Endpoint Tracking integrates with ExtremeCloud IQ Site Engine ExtremeConnect and ExtremeControl modules. The ExtremeConnect module offers API integration with third party products, such as VMware or Microsoft HyperV, from which VM endpoint information is extracted and automatically converted into usable policies for use in the ExtremeControl module, which acts as a RADIUS server for authorizing Endpoint Tracking MACs.

The following diagram illustrates an example of ExtremeCloud IQ Site Engine interaction with a switch for Endpoint Tracking:




Figure 84: ExtremeCloud IQ Site Engine Endpoint Tracking Interaction Example

RADIUS Server Attributes

The RADIUS attributes to configure in either standard or custom ExtremeCloud IQ Site Engine RADIUS profiles for Endpoint Tracking depend on your deployment and traffic type:

- For tagged traffic, if the RADIUS server provides both the VLAN ID and I-SID value, use only the **FA-VLAN-ISID** attribute.
- For tagged traffic, if the RADIUS server provides only the VLAN ID (and you are therefore using an I-SID offset value), use only the **Tunnel-Private-Group-ID** attribute.
- For untagged traffic, if the RADIUS server provides both the VLAN ID and I-SID value, use the **FA-VLAN-ISID** and **Egress-VLANID** or **Egress-VLAN-name** attributes.
- For untagged traffic, if the RADIUS server provides only the VLAN ID (and you are therefore using an I-SID offset value), use the **Tunnel-Private-Group-ID** and **Egress-VLANID** or **Egress-VLAN-name** attributes.
- Use the **Session-Timeout** attribute to override the default timeout period of 24 hours, which is amount of time, in seconds, between a MAC address authentication and the deletion of that MAC address from the Endpoint Tracking binding table.

All other RADIUS attributes are ignored.

Managing Binding Updates using RADIUS Change-of-Authorization

Endpoint Tracking uses RADIUS RFC 5176 Change-of-Authorization (CoA) functionality to enable forced VLAN:ISID binding updates.

For example, when a VLAN segment is changed on a VM that resides on a previously authenticated switch, that VM requires a new VLAN:ISID binding to reflect the new VLAN segment. Because the switch has previously been authenticated, you must force a new authentication request to update the binding information.

Using ExtremeControl, you can manually push a reauthentication request for the VM MAC. This action sends a disconnect-request from the RADIUS server to the switch, which deletes the old binding. When the switch detects the VM again, a new RADIUS

authentication request is sent from the switch to the RADIUS server, resulting in updated binding information upon successful authentication.

For more information about RADIUS Dynamic Session Change Support (RFC 5176), see [RFC 5176 — Dynamic Session Change](#) on page 3054.

Deployment Examples

Endpoint Tracking deployment scenarios include Distributed Virtual Routing (DvR) deployments, or regular SPBM deployments.

The following example illustrates a DvR deployment:

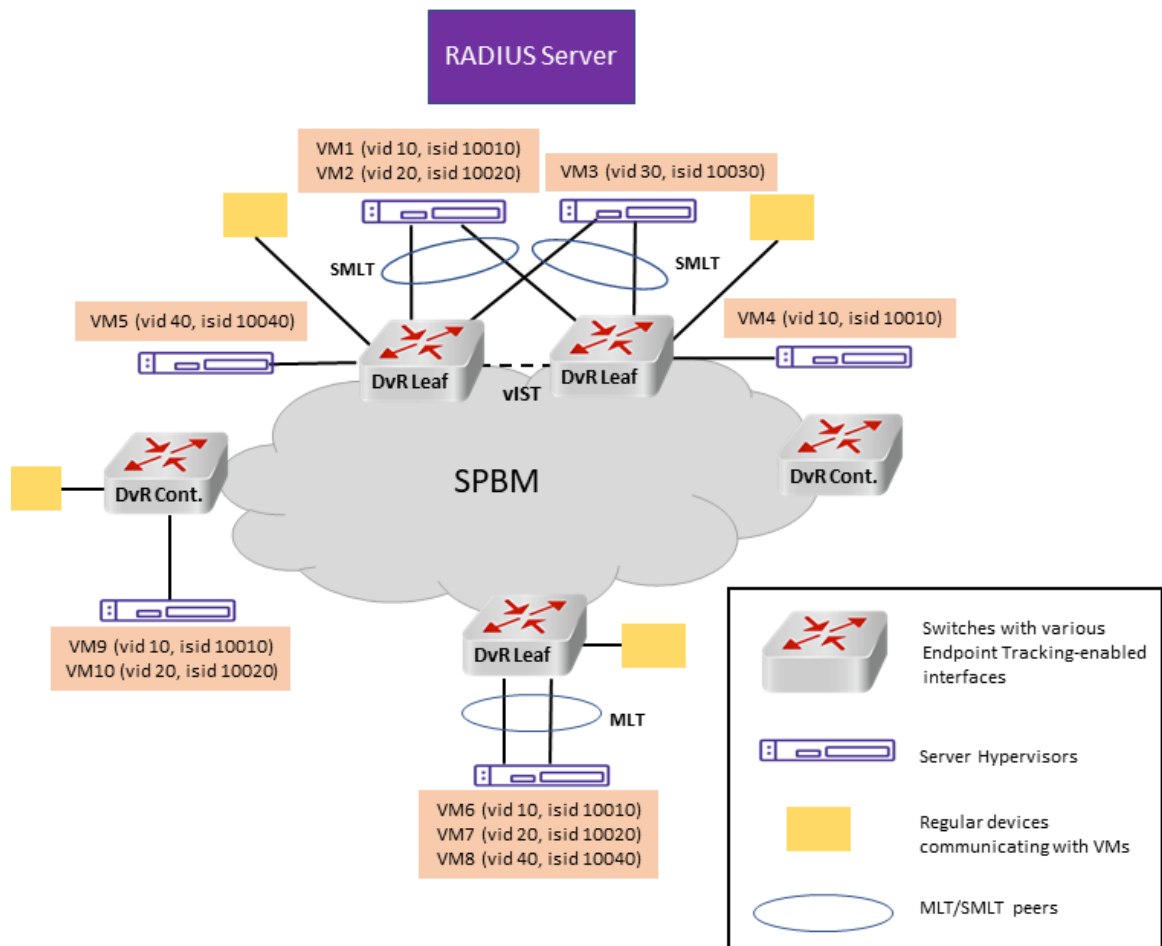


Figure 85: DvR Topology Example

The following example illustrates a regular SPBM deployment:




Figure 86: SPBM Cloud Example

Configuration Fundamentals

ExtremeCloud IQ Site Engine Configuration

To configure Endpoint Tracking, you must perform the following:

- Using the ExtremeConnect component, configure and manage your third-party virtualization platform.
- Using the ExtremeControl component, configure and manage the RADIUS server used for Endpoint Tracking authentication.

For information about configuring ExtremeCloud IQ Site Engine, see the ExtremeCloud IQ Site Engine documentation at <https://www.extremenetworks.com/support/documentation/>.

Switch Configuration

To configure Endpoint Tracking, you must perform the following:

- Add, configure, and enable the RADIUS server host as configured in the ExtremeCloud IQ Site Engine to function as the switch authentication server

for Endpoint Tracking. Ensure that you select **endpoint-tracking** for the **used-by** variable. Add, configure and enable the RADIUS dynamic-server client. For information about adding a RADIUS server host and a RADIUS dynamic-server client to the switch, see [RADIUS](#) on page 3045.

- Optionally configure a global I-SID offset value.

When you provision the Endpoint Tracking RADIUS server in ExtremeCloud IQ Site Engine, you choose which outbound attributes the RADIUS server includes in each authentication response. If you always include an I-SID value in those outbound attributes, you do not need to configure an I-SID offset value on the switch.

For MACs that do not receive an I-SID attribute from the RADIUS server, use Auto-ISID-Offset functionality. The configured I-SID offset value is used to calculate an I-SID value for a switched UNI when no I-SID value is provided by the RADIUS server in the outbound attributes. In that case, the I-SID value is calculated as follows:

$$\text{I-SID} = \text{VLAN ID} + \text{configured I-SID offset value.}$$

- After optionally configuring a global I-SID offset value, enable Endpoint Tracking globally on the switch.
- Create and enable Endpoint Tracking on each interface. Ensure that you have deleted any existing VLAN bindings on the interfaces, as the Endpoint Tracking bindings are dynamic.

CLI commands provide the functionality to separate the creation, deletion, enabling, and disabling of Endpoint Tracking on interfaces. For example, if you want to flush all VLAN:ISID bindings on a port, you can disable (but not delete) Endpoint Tracking on that port, keeping the port distinct from other ports where Endpoint Tracking is not yet created.

Configure Endpoint Tracking Using CLI

The following sections provide procedural information to configure Endpoint Tracking using CLI.

Configure Endpoint Tracking Interfaces

Create and enable Endpoint Tracking on ports and MLT/SMLT interfaces. Creating, deleting, enabling, and disabling Endpoint Tracking on interfaces can be accomplished as separate steps using this procedure.

Before You Begin

- In ExtremeCloud IQ Site Engine, configure your third-party virtualization platform, and the RADIUS server used for Endpoint Tracking authentication. For information about configuring ExtremeCloud IQ Site Engine, see the ExtremeCloud IQ Site Engine documentation at <https://www.extremenetworks.com/support/documentation/>.
- On the switch, add and configure the RADIUS server as configured in ExtremeCloud IQ Site Engine.

- To support E-Tree Private VLAN with Endpoint Tracking, you must create a Private VLAN with I-SID and configure access ports as PVLAN isolated or promiscuous.

About This Task

Configure ports and MLT/SMIT interfaces to function as Switched UNI interfaces, and then create and enable Endpoint Tracking on those interfaces.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable Flex UNI on the interface:

```
flex-uni enable
```

3. Create and enable Endpoint Tracking:

- Create Endpoint Tracking on the interface:

```
endpoint-tracking
```

- Create and enable Endpoint Tracking on the interface:

```
endpoint-tracking enable
```

What to Do Next

Configure Endpoint Tracking globally on the switch.

Configure Endpoint Tracking Globally

Configure Endpoint Tracking globally on the switch.

Before You Begin

- In ExtremeCloud IQ Site Engine, configure your third-party virtualization platform, and the RADIUS server used for Endpoint Tracking authentication. For information about configuring ExtremeCloud IQ Site Engine, see the ExtremeCloud IQ Site Engine documentation at <https://www.extremenetworks.com/support/documentation/>.
- On the switch, add and configure the RADIUS server as configured in ExtremeCloud IQ Site Engine.
- Create and enable Endpoint Tracking on interfaces.

About This Task

Optionally, if the RADIUS outbound attributes do not include an I-SID value, configure an I-SID offset value, and globally enable I-SID offset for Endpoint Tracking. The I-SID offset value is used to calculate an I-SID value for a switched UNI if no I-SID value is provided by the RADIUS server. In that case, the I-SID value is calculated as follows:

$$\text{I-SID} = \text{VLAN ID} + \text{configured I-SID offset value}.$$

After optionally configuring an I-SID offset value, enable Endpoint Tracking globally on the switch.



Note

If you have previously enabled Endpoint Tracking globally and want to change the currently configured I-SID offset value, you must disable Endpoint Tracking globally, change the I-SID value, and then re-enable Endpoint Tracking globally.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. (Optional) Configure an I-SID offset value, and enable I-SID offset globally on the switch:

```
endpoint-tracking auto-isid-offset <0-15995903>
```

```
endpoint-tracking auto-isid-offset enable
```
3. Enable Endpoint Tracking globally on the switch:

```
endpoint-tracking enable
```

Variable Definitions

The following table defines parameters for the **endpoint-tracking auto-isid-offset** command.

Variable	Value
<code><0-15995903></code>	The I-SID offset value. The default is 15990000.
<code>enable</code>	Enables or disables I-SID offset value globally on the switch. The default is disabled.

Configure Endpoint Tracking Visibility Mode

Configure Endpoint Tracking visibility mode on the switch.

Before You Begin

- In ExtremeCloud IQ Site Engine, configure your third-party virtualization platform, and the RADIUS server used for Endpoint Tracking authentication. For information about configuring ExtremeCloud IQ Site Engine, see the

ExtremeCloud IQ Site Engine documentation at <https://www.extremenetworks.com/support/documentation/>.

- On the switch, add and configure the RADIUS server as configured in ExtremeCloud IQ Site Engine.
- Create and enable Endpoint Tracking on interfaces.
- Enable Endpoint Tracking globally on the switch.

About This Task

Enable visibility mode to allow MAC learning on static S-UNIs for Endpoint Tracking.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable Endpoint Tracking visibility mode on the switch:

```
endpoint-tracking visibility-mode
```

Display Endpoint Tracking Configuration Information

Perform this procedure to display configuration information for Endpoint Tracking.

About This Task

Perform this procedure to display global, interface and binding information for Endpoint Tracking.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the global status of Endpoint Tracking on the switch, and the configured I-SID offset value, if applicable:

```
show endpoint-tracking
```
3. Display the status of all interfaces that have Endpoint Tracking created:

```
show endpoint-tracking interfaces [gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]] | [mlt <1-512>]]
```
4. Display a summary of the VLAN:ISID binding information for all ports, or MLT/SMLT interfaces:

```
show endpoint-tracking bindings summary
```
5. Display the VLAN:ISID binding information for the switch, for ports, or for MLT/SMLT interfaces:

```
show endpoint-tracking bindings [gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]] | [mlt <1-512>]]
```

Example

The following example displays all of the Endpoint Tracking configuration information for a switch.

```
Switch:1>show endpoint-tracking
```

```

Endpoint Tracking Configuration
=====

endpoint tracking status : ENABLED
auto-isid-offset value : 15990000
auto-isid-offset enabled : ENABLED
visibility-mode status : ENABLED

Switch:1>show endpoint-tracking interfaces
=====

Endpoint Tracking Interfaces
=====

PORT
NUM      INDEX      STATUS
-----
1/1       192      Enabled
1/10      201      Enabled
MLT-2     6145     Enabled
MLT-5     6148     Disabled

4 out of 4 Total Num of Endpoint Tracking interfaces displayed

Switch:1>show endpoint-tracking bindings summary
=====

Endpoint Tracking Bindings
=====

PORT/MLT  INDEX  TOTAL  ACCEPTED  REJECTED  PENDING  TIMEOUT  SERVER-UNREACHABLE
-----
1/10      201    5       5         0         0         0         0

Switch:1>show endpoint-tracking bindings
=====

Endpoint Tracking Bindings
=====

PORT/MLT  INDEX  MAC              STATUS  VLAN ID  ISID    SOURCE      TIMEOUT      TIME REMAINING
-----
1/10      201    00:00:00:00:1b:01 accept  27       15990027 autoconfig  0 day(s), 00:01:40 0 day(s), 00:00:00
1/10      201    00:00:00:00:1b:02 accept  27       15990027 autoconfig  0 day(s), 00:01:40 0 day(s), 00:00:00
1/10      201    00:00:00:00:1b:03 accept  27       15990027 autoconfig  0 day(s), 00:01:40 0 day(s), 00:00:00
1/10      201    00:00:00:00:1b:04 accept  27       15990027 autoconfig  0 day(s), 00:01:40 0 day(s), 00:00:00
1/10      201    00:00:00:00:1b:05 accept  27       15990027 autoconfig  0 day(s), 00:01:40 0 day(s), 00:00:00

5 out of 5 Total Num of Endpoint Tracking bindings displayed.

```


Variable Definitions

The following table defines parameters for the **show endpoint-tracking bindings** command.

Variable	Value
<code>gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>MLT <1-512></code>	Specifies the MLT ID.
<code>summary</code>	Provides a summary of the total number and status of bindings for all interfaces.

Configure Endpoint Tracking Using EDM

The following sections provide procedural information to configure Endpoint Tracking using Enterprise Device Manager (EDM).

Configure Endpoint Tracking Interfaces

Before You Begin

- In ExtremeCloud IQ Site Engine, configure your third-party virtualization platform, and the RADIUS server used for Endpoint Tracking authentication. For information about configuring ExtremeCloud IQ Site Engine, see the ExtremeCloud IQ Site Engine documentation at <https://www.extremenetworks.com/support/documentation/>.
- On the switch, add and configure the RADIUS server as configured in ExtremeCloud IQ Site Engine.

About This Task

Enable Endpoint Tracking on ports or MLT/SMLT interfaces.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Endpoint Tracking**.
3. Select the **Interface** tab.
4. Select **Insert**.
5. Select **Port** or **Mlt**, select the slot and port number or MLT ID, and select **OK**.
6. Select **InterfaceEnable**.
7. Select **Insert**.
8. Select **Apply**.

What to Do Next

Next, configure Endpoint Tracking globally on the switch.

Interface Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
InterfaceIndex	Specifies the interface index of the selected port or MLT.
InterfaceEnable	Enables Endpoint Tracking on the selected port or MLT.

Configure Endpoint Tracking Globally

Configure Endpoint Tracking globally on the switch.

Before You Begin

- In ExtremeCloud IQ Site Engine, configure your third-party virtualization platform, and the RADIUS server used for Endpoint Tracking authentication. For information about configuring ExtremeCloud IQ Site Engine, see the ExtremeCloud IQ Site Engine documentation at <https://www.extremenetworks.com/support/documentation/>.
- On the switch, add and configure the RADIUS server as configured in ExtremeCloud IQ Site Engine.
- Create and enable Endpoint Tracking on interfaces.

About This Task

Optionally, if the RADIUS outbound attributes do not include an I-SID value, configure an I-SID offset value, and globally enable I-SID offset for Endpoint Tracking. The I-SID offset value is used to calculate an I-SID value for a switched UNI if no I-SID value is provided by the RADIUS server. In that case, the I-SID value is calculated as follows: I-SID = VLAN ID + configured I-SID offset value.

After optionally configuring an I-SID offset value, enable Endpoint Tracking globally on the switch.



Note

If you have previously enabled Endpoint Tracking globally and want to change the currently configured I-SID offset value, you must disable Endpoint Tracking globally, change the I-SID value, and then re-enable Endpoint Tracking globally.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Endpoint Tracking**.
3. Select the **Globals** tab.

4. (Optional) Configure an I-SID offset value, and enable I-SID offset globally on the switch:

Enter a value into the **AutolsidOffset** field and select **AutolsidOffsetEnable**.

5. Select **GlobalEnable**.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AutolsidOffset	The I-SID offset value. The default is 15990000.
AutolsidOffsetEnable	Enables or disables I-SID offset value globally on the switch. The default is disabled.
GlobalEnable	Enables or disables Endpoint Tracking globally on the switch. The default is disabled.
VisibilityEnable	Enables or disables visibility mode for Endpoint Tracking. The default is disabled.

Display Binding Information

Display Endpoint Tracking binding information.

About This Task

Display all VLAN:ISID binding information on the switch for Endpoint Tracking.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Endpoint Tracking**.
3. Select the **Binding** tab.

Binding Field Descriptions

Use the data in the following table to use the **Binding** tab.

Name	Description
IfIndex	Specifies the interface index of the selected port or MLT.
MacAddress	Specifies the MAC address that corresponds to the VLAN:ISID binding.

Name	Description
Status	Specifies the Endpoint Tracking data binding status as follows: <ul style="list-style-type: none"> • pending: indicates that a request has been sent to the RADIUS server • accept: indicates that the RADIUS server has successfully returned the request • reject: indicates that the RADIUS server has rejected the request • timeout: indicates that the RADIUS server request has timed out. The entry is deleted if it remains in this state for 15 minutes. • serverNotConfigured: indicates that the RADIUS server is not configured for Endpoint Tracking. The entry is deleted if it remains in this state for 15 minutes.
VlanId	Specifies the VLAN ID.
Isid	Specifies the I-SID value, either provided by the RADIUS server, or calculated using the VLAN ID plus the configured I-SID offset value.
IsidSource	Specifies whether the I-SID value is provided by the RADIUS server (radius), or calculated using the VLAN ID plus the configured endpoint-tracking offset value (autoconfig).
Timeout	Specifies the timeout period that is applied to the MAC in the bindings table when the MAC is aged out. If the MAC is in timeout state (there is no response from the RADIUS server), the timeout triggers immediately with a 15 minute period. Otherwise, the default timeout is one day, and triggers the moment the MAC ages out from the VLAN/I-SID bridge forwarding database (FDB) table. The default timeout of one day can be overridden by the RADIUS server if the Session-Timeout attribute is configured and returned.
TimeRemaining	Specifies the time remaining until the Endpoint Tracking data binding entry expires.

Multi-area SPB

Table 139: Multi-area SPB product support

Feature	Product	Release introduced
Multi-area SPB Boundary Node	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.4

Table 139: Multi-area SPB product support (continued)

Feature	Product	Release introduced
Static data I-SID redistribution for Multi-area SPB Boundary Node	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.8
Virtual NNI links for Multi-area SPB Boundary Nodes	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.0.3
	7720 Series	Fabric Engine 9.0.3
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 9.0.3

The Multi-area SPB feature enables building a scalable Shortest Path Bridging (SPB) network, consisting of multiple Intermediate System-to-Intermediate System (IS-IS) areas that connect to each other through the boundary nodes. The Multi-area SPB feature provides flexible network design through which you can group the nodes in the areas on a need basis.

Each area in the Multi-area SPB network contains the interior nodes and the boundary nodes.




Figure 87: Boundary nodes in the Multi-area SPB network



Tip

As a best practice, ensure you have a high-speed direct physical connection between boundary nodes.

Each boundary node runs two Intermediate-System-to-Intermediate-System (IS-IS) instances, representing the home area and the remote area. Each IS-IS instance runs the Shortest Path First (SPF) computation and uses its own Link State Database (LSDB). The boundary nodes support Backbone Edge Bridge (BEB) functionality in addition to the Backbone Control Bridge (BCB) configuration and the UNIs belong only to the home area. You can configure IS-IS interfaces on each instance.



Note

The boundary nodes do not support Auto-sense for remote interfaces. You must manually configure all remote NNI on the boundary nodes. Each time the port receives a Fabric Connect TLV, the port is configured as NNI in the home area. You must disable Auto-sense on the IS-IS remote area ports.

The multiple areas in the Multi-area SPB network connect in various patterns, for example, the hub and spoke topology and the daisy chain topology.

Designated and Non-designated Boundary Nodes

Among the boundary nodes in the Multi-area SPB network, the system elects the designated boundary node on basis of the unique chassis MAC value of the node. The nodes in the SPB area advertise their chassis MAC value using the link state packet (LSP) and the system elects the node with the smallest chassis MAC value as the designated boundary node.

The functions of the boundary nodes in the Multi-area SPB network are:

- Designated boundary node:
 - Creates the area virtual node LSPs on each IS-IS instance.
 - Configures the virtual node's adjacency with the other boundary node on basis of the reachability of the node.
 - Advertises the services on the area virtual node on basis of the redistribution policies that you configure.
- Non-designated boundary node:
 - Connects to the area virtual node and the system uses it to forward the traffic on basis of the IS-IS computation.
 - Takes over as the designated boundary node if the designated boundary node is no longer available in any area.

For more information about the area virtual node, see [Area Virtual Node](#) on page 1425.

Minimum Requirements

The minimum requirements for the Multi-area SPB to function are:

- Two boundary nodes to provide redundancy.
- Multi-area SPB redistribution policies between areas, like Layer 2 redistribution, Layer 3 unicast redistribution, multicast redistribution, Distributed Virtual Routing (DvR) backbone redistribution and so on. For more information, see [Multi-area SPB Redistribution and Policies](#) on page 1427.

Area Virtual Node

The designated boundary node in the Multi-area SPB network creates the area virtual node. The area virtual node sums up an entire network topology in a single virtual node representing all the services from the represented area.




Figure 88: Representation of area virtual node

When considering the shortest path between two nodes in different areas in the Multi-area SPB network, the area virtual node logically connects the shortest path within one area to the shortest path within another area. The area virtual nodes help the Multi-area SPB network to visualize any external area as a single IS-IS entity.

The boundary nodes in the Multi-area SPB network exposes one virtual node in each of two connected areas, to represent a collection of nodes from the other area. On the basis of the policies that you configure on the boundary nodes, you can control the services that the system exports from one area to the other area as the system includes the configuration in the Link State Packets (LSP) of the corresponding area virtual node.

You can configure the area virtual node parameters on both home and remote IS-IS instances. For more information about configuring the area virtual node, see [Configure IS-IS Multi-area SPB Virtual Node](#) on page 1438.

Layer 2 Ping and Layer 2 Traceroute Support

Multi-area SPB supports Connectivity Fault Management (CFM) on boundary nodes in both home and remote area. For the remote area, the boundary nodes respond to the Layer 2 ping and Layer 2 traceroute messages that contain the remote area system ID. You must enable CFM on the boundary nodes for the functionality to work.

Multi-area SPB supports the following Layer 2 ping, Layer 2 traceroute, and Layer 2 tracertree interactions:

- from boundary node to remote area Backbone Edge Bridge (BEB)

- from remote area BEB to boundary node

**Note**

Multi-area SPB does not support CFM between the home and remote area.

The following list identifies CFM restrictions for virtual node interactions:

- The boundary node cannot use Layer 2 ping or Layer 2 traceroute to the virtual node.
- The virtual node does not reply to Layer 2 tracertree or Layer 2 tracemroute.
- The virtual node does not initiate CFM messages.

For more information about CFM, see [Connectivity Fault Management](#) on page 3799.

Remote IS-IS Instance

The Multi-area SPB feature introduces the remote Intermediate-System-to-Intermediate-System (IS-IS) instance. You can configure the manual area, system ID, nickname, and area virtual node parameters on the remote IS-IS instance. When enabling Remote IS-IS Instance, make sure that the physical node nickname, virtual node nickname and system ID are different. To configure the remote IS-IS instance on the switch, you must:

- Configure manual area and nickname manually on the switch.
- Enable IS-IS and remote IS-IS instances on the switch, globally.
- Configure Shortest Path Bridging MAC (SPBM) instance globally.
- Configure Backbone-VLANs (B-VLANs) globally on both home and remote instances.

For more information, see [Enable Remote IS-IS Router Configuration Globally](#) on page 1445.

Multi-area SPB Redistribution and Policies

The Multi-area SPB redistribution and policies govern the process of importing routes in the Multi-area SPB network across the boundary nodes, that is from the home area to the remote area or vice versa. The traffic flow between the areas within the Multi-area SPB network is the same as the packet flow between the source and destination in a flat SPB network. If more than one boundary node connects to same two areas then the redistribution configuration on the boundary nodes must be consistent, that is the system must permit or deny the same routes to the cross area boundaries. With Multi-area SPB you can configure layer 2 redistribution, layer 3 unicast redistribution, multicast over SPB redistribution, and Distributed Virtual Routing backbone redistribution on the switch.

**Note**

5520 Series does not support Distributed Virtual Routing Backbone Redistribution

Layer 2 Redistribution

In the Multi-area SPB network the layer 2 redistribution identifies the services that the system transports from the home area to the remote area and the I-SIDs that the system extends across the boundary nodes in an area. The destination area permits or denies the services on basis of the Layer 2 redistribution policies that you configure.

You can configure the Layer 2 I-SID redistribution on the switch globally. The Layer 2 redistribution policies are bidirectional, which means only one redistribution filter applies to both the home area and the remote area.

**Note**

Any changes in the Layer 2 I-SIDs in the home area will trigger updates to the redistributed I-SIDs in the remote area or vice versa.

For more information about configuring Multi-area SPB Layer 2 redistribution, see:

- [Configure Multi-area SPB Layer 2 I-SID Redistribution](#) on page 1446
- [Configure Multi-area SPB Layer 2 I-SID List](#) on page 1446

Layer 3 Unicast Redistribution

The layer 3 VSN functionality within the Multi-area SPB network is same as the layer 3 VSN functionality in the flat SPB network. In the Multi-area SPB network, for the layer 3 services to pass the area boundary, you must explicitly configure them on the boundary nodes. Inter-VRF routes on the boundary node will not be redistributed across area boundaries. You can configure redistribution unidirectionally on the boundary nodes. You can also apply different policies from home to remote and remote to home.

The configuration guidelines for the layer 3 Multi-area SPB redistribution are:

- To redistribute inter-area routes in GRT, configure IPSC on the boundary nodes.
- To redistribute inter-area routes in a specific VRF, configure IPSC and layer 3 VSN for that VRF.

**Note**

The system does not redistribute inter-vrf routes on the boundary nodes across the area boundaries.

For information about configuring layer 3 unicast redistribution, see:

- [Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution](#) on page 1447
- [Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution on a VRF Instance](#) on page 1451
- [Configure IS-IS Multi-area SPB IPv6 Unicast Redistribution](#) on page 1449
- [Configure IS-IS Multi-area SPB IPv6 Unicast Redistribution on a VRF Instance](#) on page 1452

Multicast over SPB Redistribution

Multicast over Shortest Path Bridging (SPB) redistribution manages filtering of the IP multicast streams that the system transports from one area to another. The IP

multicast services pass through to the other area on the boundary node depending on the policy that you configure specific to the sender IP address, the multicast group IP address and the Virtual Router Forwarding (VRF) ID.

Multicast over SPB policy is a unidirectional solution, the system filters the IP multicast streams from the area that they originate to the other area or vice versa, independently. Each IP multicast stream that the system advertises in one area will also be advertised on the corresponding area virtual node after it passes the multicast over SPB policy.

Following are the configuration guidelines for the multicast over SPB redistribution:

- Enable multicast on the boundary nodes.
- For layer 3 VSNs, configure Multicast Virtual Private Network (MVPN) on the boundary nodes.
- For layer 2 multicast, the scope I-SID must pass the layer 2 VSN policy.
- Configure the same policies as on the designated boundary node on the peer boundary nodes.

For information about configuring Multi-area SPB multicast redistribution, see:

- [Configure Multi-area SPB IPv4 Routed Multicast Redistribution](#) on page 1457
- [Configure IS-IS Multi-area SPB IPv4 Routed Multicast Redistribution on a VRF Instance](#) on page 1458

Distributed Virtual Routing Backbone Redistribution

**Note**

5520 Series does not support Distributed Virtual Routing Backbone Redistribution

The Distributed Virtual Routing (DvR) functionality within the Multi-area SPB network is same as the DvR functionality in the flat SPB network. For more information, see [Distributed Virtual Routing Fundamentals](#) on page 919. To configure Multi-area SPB DvR backbone redistribution, the boundary node must be a functional DvR controller in a domain in the home area or a non-DvR Backbone Edge Bridge.

**Note**

DvR domains can only be a part of one area, you must not configure the same DvR domain in multiple areas.

When you enable the Multi-area SPB DvR backbone redistribution the system extends the DvR backbone across the area boundaries. The boundary node in the Multi-area SPB network performs the DvR controller functions to receive and propagate the hosts in the DvR domain. Since the interior nodes in one area of the Multi-area SPB network are not visible to the interior nodes in other areas, the system sends the area virtual nodes representing the specific area as a gateway in the DvR route TLV. When the system sends the DvR traffic to another DvR domain in a specific destination area, the system sends the traffic to the destination area's boundary nodes and the boundary nodes perform the route lookup and forward the traffic to the correct DvR domain in the destination area.

For information about configuring Multi-area SPB DvR backbone redistribution, see [Configure Multi-area SPB DvR Backbone Redistribution](#) on page 1454.

Multi-area SPB Considerations and Restrictions

The following list identifies the restrictions and considerations that apply to the Multi-area SPB feature:

- Two boundary nodes can be either in a vIST configuration (paired with each other) or in a non-vIST configuration. Three or more boundary nodes can only exist in a non-vIST configuration. Any other combination of boundary nodes is not supported.



Note

Only two boundary nodes can be in a vIST configuration.

- Up to four nodes can function as boundary nodes between any given pair of areas.



Note

5520 Series and 5720 Series support a maximum of two nodes that can function as boundary nodes.

- You must not connect the same Protocol Independent Multicast (PIM) domain to the SPB-PIM Gateway nodes that are in different Intermediate-System-to-Intermediate-System (IS-IS) areas, to avoid the inter-area redistribution of the same multicast information.
- You can enable the Dynamic Nickname server on the boundary nodes in the home area, but the boundary nodes cannot be clients in any of the two areas. The boundary nodes do not support the Dynamic Nickname server in the remote area.
- You must manually configure the backbone VLANs (B-VLAN) on the boundary nodes, so the system does not learn the dynamic values that it receives through the Link Layer Discovery Protocol (LLDP). However, the system sends the manually configured B-VLANs on the BN through LLDP, so that other neighbors can learn them (both in home and remote areas).
- Each time the port receives a Fabric Connect TLV, the port is configured as NNI in the home area. You must disable Auto-sense on the IS-IS remote area ports. For more information, see [Auto-sense Port States](#) on page 28.
- If the system forms an adjacency between two boundary nodes that are part of the home and remote area, the hello packets in the home area use the home manual area and the hello packets in the remote area use the remote manual area.
- If the system forms a home and a remote adjacency on the same port then the Multi-area SPB feature uses different Backbone VLAN IDs (B-VIDs) for each adjacency, the home adjacency uses the primary B-VID and the remote adjacency uses the secondary B-VID.
- If the system forms an IS-IS adjacency in both the home and remote areas on a boundary node of the same port then the remote adjacency stays up only with another boundary node that also has IS-IS configured on both the home and remote areas of the same port.

- If a boundary node connects to a Backbone Edge Bridge (BEB) in the remote area and if you configure IS-IS in the home area on the same interface, then the remote adjacency goes down.
- On the boundary node, to install a route from a remote area in the routing table manager (RTM), the route must pass the accept policy and the Multi-area SPB redistribution policy that you configure on the specific Virtual Router Forwarding (VRF) instance.
- On the boundary node, to install an inter-VRF route from a remote area in the routing table manager (RTM), the inter-VRF route must pass the accept policy and the Multi-area SPB redistribution policy that you configure on both the source and destination VRF instances.
- Nickname and system ID for the physical node and virtual node must be different.
- When enabling Remote IS-IS Instance, make sure that the physical node nickname, virtual node nickname and system ID are different.
- You cannot establish an SSH connection to the boundary node from an IS-IS remote area.

Multi-area Deployment Guidelines

Use the following guidelines to design a Multi-area network:

- In order to achieve optimal convergence performance, use the following order of preference for inter-connecting Boundary nodes:
 1. Direct links between all Boundary Nodes with adjacencies configured for both home and remote area.
 2. Reliable and redundant high bandwidth forwarding path between any two Boundary nodes in both areas.
 3. Reliable and redundant high bandwidth forwarding path between any two Boundary nodes in home area and Virtual NNI Links in remote area.
 4. Reliable and redundant high bandwidth forwarding path between any two Boundary nodes in remote area and Virtual NNI Links in home area.
- Manually configure the Virtual NNI Links IS-IS metric to reflect the most desired traffic forwarding pattern. As a general rule, physical NNIs are preferred over Virtual NNI Links.
- Manually configure the IS-IS metric of any relevant NNI in order to avoid using Virtual NNI Links to forward unicast and multicast traffic, when possible.
- In a non-vIST configuration with three or more boundary nodes, use the following guidelines to ensure there is a robust Fabric path between them in both areas (home and remote):
 - If a direct connection exists between the boundary nodes, configure both areas NNIs onto the link or links.
 - If there is no direct connection between the boundary nodes, make sure a robust Fabric connectivity exists in both areas.
 - If a WAN with low bandwidth exists on one side, you can create a virtual NNI link and establish a virtual Fabric adjacency over the area with the robust connection. With this, you can interconnect your WAN area over the LAN area with high

performance connectivity and avoid the path between the boundary nodes in the WAN area.

For more information, see [Virtual NNI Links for Multi-Area Boundary Nodes](#) on page 1432.

Duplicate System IDs Between IS-IS Areas

Duplicate system IDs between adjacent IS-IS areas are not supported.

If a duplicate system ID exists between adjacent areas, the boundary node logs an event like the following:

```
1 2021-09-21T02:08:44.734-04:00 GB12 CP1 - 0x001dc705 - 00000000
GlobalRouter ISIS WARNING Duplicate System ID 0022.2222.2415 detected on
Boundary Node between HostName:B15 Chassis MAC:b0:ad:aa:41:c0:00 (HOME
area) and HostName:B14 Chassis MAC:b0:ad:aa:41:b8:00 (REMOTE area)
```

To recover from this event and ensure the information on the boundary node is correctly updated, you must perform the following tasks:

1. Change the system ID on both of the conflicting nodes.
2. Disable and re-enable IS-IS on both of the conflicting nodes.

Delete errors can display on the boundary node but they have no functional impact.

Duplicate Nicknames Between IS-IS Areas

Duplicate nicknames between adjacent IS-IS areas are not supported.

If a duplicate nickname exists between adjacent areas, the boundary node logs an event like the following:

```
GlobalRouter ISIS WARNING Duplicate Nickname f.f0.3c detected on
Boundary Node between SysId 4088.2fe9.8884 (REMOTE area) and SysId
f4ce.48a1.f084 (HOME area)
```

To recover from this event and ensure the information on the boundary node is correctly updated, you must perform the following tasks:

1. Change the nickname on both of the conflicting nodes.
2. Disable and re-enable IS-IS on both of the conflicting nodes.

Delete errors can display on the boundary node but they have no functional impact.

Virtual NNI Links for Multi-Area Boundary Nodes

Boundary nodes in the Multi-area SPB network require a robust Fabric path between them in both areas (home and remote). If a robust connection for one of the areas is not possible, you can create a virtual NNI link and establish a virtual Fabric adjacency over the area with the robust connection. For instance, if an adjacency exists between

two boundary nodes in the home area, you can use virtual NNI functionality to establish a Fabric adjacency in the remote area, and vice versa.

To create virtual NNI link functionality between boundary nodes, you must complete the following tasks:

- Configure the loopback IP address as the source IP address using the Multi-area virtual link flag.
- Configure a logical IS-IS interface using the Multi-area virtual link flag.
- Configure IS-IS on the logical interface.

IP Shortcuts automatically redistributes the IP address in the specific area without the need for an IS-IS redistribution policy. Boundary nodes receive this IP address in the corresponding area as an IS-IS redistributed route.

Virtual NNI Links on Boundary Nodes Considerations and Restrictions

The following list identifies considerations and restrictions that apply to virtual NNI links on boundary nodes:

- Function on boundary nodes only.
- Support only a CLIP (loopback) interface configured on the GRT.
- Support only one CLIP (loopback) interface.
- Do not support Bidirectional Forwarding Detection (BFD). You cannot enable BFD on a logical interface for a virtual NNI link.
- Cannot coexist with VXLAN Gateway. You cannot enable VXLAN Gateway on a loopback interface for a virtual NNI link and you cannot create a loopback interface for a virtual NNI link when VXLAN Gateway is enabled.

Multi-area SPB Configuration using CLI

Perform the procedures in this section to configure Multi-area SPB parameters like virtual node,router remote Intermediate-System-to-Intermediate-System (IS-IS), layer 2 unicast and multicast redistribution, and IPv4 and IPv6 route redistribution on the switch using the Command Line Interface (CLI).



Note


The following procedures do not apply to 4220 Series, 5320 Series, or 5420 Series.

Multi-area SPB Configuration Flow

Perform the following procedures in sequence to configure the Multi-area SPB feature on the switch.

Procedure

1. [Configure Remote IS-IS Multi-area SPB Nickname for an SPBM Instance](#) on page 1434.
2. [Display the Remote IS-IS SPBM Nickname Configuration](#) on page 1466.

3. [Configure Remote IS-IS Manual Area](#) on page 1435.
4. [Configure Remote IS-IS System ID](#) on page 1436.
5. [Configure an IS-IS Area Name](#) on page 1437
6. [Display IS-IS Remote Area Configuration](#) on page 1467.
7. [Display IS-IS Area Configuration](#) on page 1467.
8.  **Note**
As a best practice, use the derived values for the virtual node to avoid misconfiguration between peer boundary nodes. However, you can use this procedure to override the derived values.

(Optional) [Configure IS-IS Multi-area SPB Virtual Node](#) on page 1438.

9. [Display IS-IS Area Virtual Node Configuration](#) on page 1468.
10. [Configure Remote IS-IS Hello Authentication on a Port](#) on page 1439.
11. [Configure Remote IS-IS Layer 1 Designated Router Priority on a Port](#) on page 1440.
12. [Configure Remote IS-IS Hello Interval on a Port](#) on page 1441.
13. [Configure Remote IS-IS Hello Multiplier on a Port](#) on page 1442.
14. [Configure Remote IS-IS SPBM Parameters on a Port](#) on page 1443.
15. [Enable Remote IS-IS on a Port](#) on page 1444.
16. [Enable Remote IS-IS Router Configuration Globally](#) on page 1445.
17. [Display IS-IS Remote Area Configuration](#) on page 1467.
18. [Configure Multi-area SPB Layer 2 I-SID List](#) on page 1446.
19. [Display IS-IS Multi-area SPB Layer 2 I-SID List Information](#) on page 1468.
20. [Configure Multi-area SPB Layer 2 I-SID Redistribution](#) on page 1446.
21. [Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution](#) on page 1447.
22. [Configure IS-IS Multi-area SPB IPv6 Unicast Redistribution](#) on page 1449.
23. [Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution on a VRF Instance](#) on page 1451.
24. [Configure IS-IS Multi-area SPB IPv6 Unicast Redistribution on a VRF Instance](#) on page 1452.
25. [Configure Multi-area SPB DvR Backbone Redistribution](#) on page 1454.
26. [Display IS-IS Multi-area SPB DvR Configuration](#) on page 1471.
27. [Configure Multi-area SPB Layer 2 Multicast Snooping Redistribution](#) on page 1455.
28. [Configure Multi-area SPB IPv4 Routed Multicast Redistribution](#) on page 1457.
29. [Configure IS-IS Multi-area SPB IPv4 Routed Multicast Redistribution on a VRF Instance](#) on page 1458.
30. [Configure Multi-area SPB IPv4 Static I-SID Routed Multicast Redistribution](#) on page 1461.

Configure Remote IS-IS Multi-area SPB Nickname for an SPBM Instance

About This Task

Perform this procedure to configure remote IS-IS Multi-area SPB nickname for a specific SPBM instance on the switch.

Procedure

- 1. Enter IS-IS Router Remote Configuration mode:
`enable`

`configure terminal`

`router isis remote`
- 2. Configure IS-IS Multi-area SPB nickname:
`spbm <1-100> nick-name <x.xx.xx>`

Example

Configuring remote IS-IS Multi-area SPB nickname for SPBM instance:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis remote
Switch:1(config-remote)#spbm 4 nick-name 1.11.16
```

What to Do Next

To verify the configuration, see [Display the Remote IS-IS SPBM Nickname Configuration](#) on page 1466.

Variable Definitions

The following table defines parameters for the **spbm** command.

Variable	Value
<1-100>	Specifies the SPBM instance on the remote IS-IS interface.
nick-name<x.xx.xx>	Specifies a nickname for the remote IS-IS SPBM Multi-area SPB instance. The value is 2.5 bytes in the format <x.xx.xx>.

Configure Remote IS-IS Manual Area

About This Task

Perform this procedure to configure the remote Intermediate System-to-Intermediate System (IS-IS) manual area.

Procedure

- 1. Enter IS-IS Router Remote Configuration mode:
`enable`

`configure terminal`

`router isis remote`
- 2. Configure the manual-area:
`manual-area xx.xxxx.xxxx...xxx`

Example

Configuring remote IS-IS manual area:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis remote
Switch:1(config-isis-remote)#manual-area c1.1000.000.00
```

What to Do Next

To verify the configuration, see [Display IS-IS Remote Area Configuration](#) on page 1467 and [Display IS-IS Area Configuration](#) on page 1467.

Variable Definitions

The following table defines parameters for the **manual-area** command.

Variable	Value
<xx.xxxx.xxxx...xxx x>	Specifies the remote IS-IS manual-area in the range of 1 to 13 bytes.

Configure Remote IS-IS System ID

About This Task

Perform this procedure to configure the remote Intermediate System-to-Intermediate System (IS-IS) system ID on the switch.

Procedure

- 1. Enter IS-IS Router Remote Configuration mode:

enable

configure terminal

router isis remote
- 2. Configure the remote IS-IS system ID:

system-id xxxx.xxxx.xxxx

Example

Configuring the remote IS-IS system ID:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis remote
Switch:1(config-isis-remote)#system-id 0017.c6ed.44df
```

What to Do Next

To verify the configuration, see [Display IS-IS Remote Area Configuration](#) on page 1467 and [Display IS-IS Area Configuration](#) on page 1467.

Variable Definitions

The following table defines parameters for the **system-id** command.

Variable	Value
<code>xxxx.xxxx.xxxx</code>	Specifies the remote IS-IS system ID for the switch.

Configure an IS-IS Area Name

Before You Begin

- You must disable IS-IS or remote IS-IS.

About This Task

Perform this procedure to configure the area name for home and remote areas.

Procedure

- Enter either IS-IS Router Configuration or IS-IS Router Remote Configuration mode:
`enable`
`configure terminal`
`router isis` OR `router isis remote`
- Configure the area name:
`area-name WORD<0-255>`
- Enable IS-IS or remote IS-IS:
`router isis enable` OR `router isis remote enable`

Example

Configure the home and remote area names:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router isis
Switch:1(config-isis)#area-name 30
Switch:1(config)#router isis remote
Switch:1(config-isis-remote)# area-name 50
```

Variable Definitions

The following table defines parameters for the **area-name** command.

Variable	Value
<code>WORD<0-255></code>	Specifies the area name. The default value is area-manual-area, where manual-area represents the IS-IS or remote IS-IS manual-area value that you configure.

*Configure IS-IS Multi-area SPB Virtual Node***About This Task****Note**

As a best practice, use the derived values for the virtual node to avoid misconfiguration between peer boundary nodes. However, you can use this procedure to override the derived values.

The following list describes how the switch derives the virtual node parameters:

- The switch derives the virtual node nickname from the last 5 nibbles of the manual area address. For example, area ID 49.0001 generates a nickname of 9.00.01.
- The switch derives the system ID and system name from the nickname. For example, a nickname of 9.00.01 generates system ID 92:00:01:ff:ff:f0.

Perform this procedure to configure the IS-IS Multi-area SPB virtual node parameters like nickname, system ID, and system name in the home or remote area.

Procedure

1. Enter either IS-IS Router Configuration or IS-IS Router Remote Configuration mode:

```
enable

configure terminal

router isis OR router isis remote
```
2. Configure the nickname for the IS-IS Multi-area SPB virtual node:

```
area-vnode nick-name <x.xx.xx>
```
3. Configure the system name for the IS-IS Multi-area SPB virtual node:

```
area-vnode sys-name WORD<0-255>
```
4. Configure the system ID for the IS-IS Multi-area SPB virtual node:

```
area-vnode system-id xxxx.xxxx.xxxx
```

Examples

Configuring the IS-IS Multi-area SPB virtual node in the home area:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#area-vnode nick-name 0.82.40
Switch:1(config-isis)#area-vnode sys-name SwitchC
Switch:1(config-isis)#area-vnode system-id 1222.3300.0000
```

Configuring the IS-IS Multi-area SPB virtual node in the remote area:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis remote
Switch:1(config-isis)#area-vnode nick-name 0.82.60
Switch:1(config-isis)#area-vnode sys-name SwitchA
Switch:1(config-isis)#area-vnode system-id 1222.3311.2222
```

What to Do Next

To verify the configuration, see [Display IS-IS Area Virtual Node Configuration](#) on page 1468.

Variable Definitions

The following table defines parameters for the **area-vnode** command.

Variable	Value
<i>nick-name</i> x.xx.xx	Specifies a nickname for the IS-IS SPBM Multi-area SPB instance. The value is 2.5 bytes in the format x.xx.xx.
<i>sys-name</i> WORD<0-255>	Specifies the name of the system.
<i>system-id</i> xxxx.xxxx.xxxx	Specifies the IS-IS Multi-area SPB virtual node system ID.

Configure Remote IS-IS Hello Authentication on a Port

About This Task

Perform this procedure to configure authentication type used for remote Intermediate-System-to-Intermediate-System (IS-IS) hello packets on the interface. The type can be one of the following: none, simple, hmac-md5, or hmac-sha-256.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the hello authentication type:

```
isis remote hello-auth type {none | simple | hmac-md5 | hmac-sha-256}  
[key WORD<1-16>] [key-id <1-255>]
```

Example

Configuring the simple authentication type for remote IS-IS hello packets on port 1/2:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/2  
Switch:1(config-if)#isis remote hello-auth type simple key Test key-id 125
```

Variable Definitions

The following table defines parameters for the **isis remote hello-auth** command.

Variable	Value
<i>key WORD<1-16></i>	Specifies the authentication key (password) that the receiving router uses to verify the packet.
<i>key-id <1-255></i>	Specifies the optional key ID.
<i>type {none simple hmac-md5 hmac-sha-256}</i>	<p>Specifies the authentication type used for remote IS-IS hello packets on the interface. The type can be one of the following:</p> <ul style="list-style-type: none"> Simple - Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. You can also specify a key value. hmac-md5 - MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. You can also specify a key value and key-id. hmac-sha-256 - With SHA-256 authentication, the switch adds an HMAC-SHA256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. If the digests match, the packet is accepted. If the digests do not match, the receiving switch discards the packet. You can also specify a key value and key-id. <p>The default type is none.</p>

Configure Remote IS-IS Layer 1 Designated Router Priority on a Port

About This Task

Perform this procedure to configure the Level 1 remote Intermediate-System-to-Intermediate-System (IS-IS) designated router priority to the specified value.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the Level 1 remote IS-IS designated router priority:

```
isis remote ll-dr-priority <0-127>
```

Example

Configuring the remote IS-IS designated router priority for port 1/4:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#isis remote ll-dr-priority 1
```

Variable Definitions

The following table defines parameters for the **isis remote ll-dr-priority** command.

Variable	Value
<0-127>	Specifies the level 1 remote Intermediate-System-to-Intermediate-System (IS-IS) designated router priority value. The default value is 64.

Configure Remote IS-IS Hello Interval on a Port

About This Task

Perform this procedure to configure the remote hello interval to change how often hello packets are sent out from an interface level.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the remote hello interval:

```
isis remote ll-hello-interval <1-600>
```

Example

Configuring the remote hello interval on port 1/3:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/3
Switch:1(config-if)#isis remote ll-hello-interval 1
```

Variable Definitions

The following table defines parameters for the **isis remote l1-hello-interval** command.

Variable	Value
<1-600>	Configures the Level 1 remote hello interval. The default value is 9 seconds.

Configure Remote IS-IS Hello Multiplier on a Port

About This Task

Perform this procedure to configure the remote hello multiplier to specify how many hellos the switch must miss before it considers the adjacency with a neighboring switch down.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the level 1 hello multiplier:

```
isis remote l1-hello-multiplier <1-600>
```

Example

Configuring the remote hello multiplier on port 1/2:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#isis remote l1-hello-multiplier 1
```

Variable Definitions

The following table defines parameters for the **isis remote l1-hello-multiplier** command.

Variable	Value
<1-600>	Configures the Level 1 hello multiplier. The default value is 3 seconds.

Configure Remote IS-IS SPBM Parameters on a Port

About This Task

Perform this procedure to configure the remote Intermediate-System-to-Intermediate-System (IS-IS) Shortest Path Bridging MAC (SPBM) parameters on a specific port.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the remote IS-IS SPBM interface type:

```
isis remote spbm <1-100> interface-type {broadcast | pt-pt}
```

3. Configure the cost of the remote SPBM instance:

```
isis remote spbm <1-100> ll-metric {<1-16777215> | auto}
```

Example

Configuring the remote IS-IS SPBM interface type and layer 1 metric on port 1/2:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/2  
Switch:1(config-if)#isis remote spbm 2 interface-type broadcast  
Switch:1(config-if)#isis remote spbm 2 ll-metric 2
```

Variable Definitions

The following table defines parameters for the **isis remote spbm** command.

Variable	Value
<1-100>	Specifies the SPBM instance.
<i>interface-type</i> {broadcast pt-pt}	Specifies the IS-IS SPBM interface type as broadcast or point-to-point (pt-pt).
<i>ll-metric</i> {<1-16777215> auto}	Specifies the SPBM instance Level 1 metric on the IS-IS interface located on a specific port. You can configure a static value, or use auto. When you use the auto setting, the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed.

Enable Remote IS-IS on a Port

Before You Begin

- Configure manual-area in both home and remote area.
- Configure the Shortest Path Bridging MAC (SPBM) instance and backbone VLANs (B-VLAN) manually.
- Configure the nickname in both home and remote area manually.
- Disable Auto-sense on the port.
- Make sure that the physical node nickname, virtual node nickname and system ID are different.

About This Task

Perform this procedure to enable remote Intermediate-System-to-Intermediate-System (IS-IS) interface on the selected port(s).



Note

When the switch receives a Fabric Connect TLV through LLDP on a port that has remote IS-IS enabled on it, the port transitions to the Auto-sense NNI pending state. For more information, see [Auto-sense Port States](#) on page 28.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable remote IS-IS interface:

```
isis remote enable
```

Example

Enabling remote IS-IS interface on port 1/2:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/2  
Switch:1(config-if)#isis remote enable
```

Variable Definitions

The following table defines parameters for the **isis remote** command.

Variable	Value
<i>enable</i>	Enables a remote IS-IS interface on the selected ports.

Enable Remote IS-IS Router Configuration Globally

Before You Begin

- Configure manual-area in both home and remote area.
- Configure the Shortest Path Bridging MAC (SPBM) instance and backbone VLANs (B-VLAN) manually.
- Configure the nickname in both home and remote area manually.

About This Task

Perform this procedure to enable remote IS-IS router configuration on the switch, globally.



Note

You cannot enable TTL handling for bridged traffic with routed SPB IP Multicast if **router isis remote** is administratively enabled. You cannot enable **router isis remote** if TTL handling is already configured.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Enable IS-IS router remote configuration:

```
router isis remote [enable]
```

Example

Enabling IS-IS router remote configuration globally:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis remote enable
```

What to Do Next

To verify the configuration, see [Display IS-IS Remote Area Configuration](#) on page 1467.

Variable Definitions

The following table defines parameters for the **router isis remote** command.

Variable	Value
<i>enable</i>	Enables router IS-IS remote configuration globally on the switch.

Configure Multi-area SPB Layer 2 I-SID List

About This Task

Perform this procedure to configure the Multi-area SPB layer 2 I-SID list.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

2. Configure the Multi-area SPB layer 2 I-SID list:

```
multi-area 12 isid-list WORD<1-32> WORD<1-1024>
```

Example

Configuring the Multi-area SPB layer 2 I-SID list that includes specific I-SID values:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area 12 isid-list testlist 1,3,5,100-200
```

What to Do Next

To verify the configuration, see [Display IS-IS Multi-area SPB Layer 2 I-SID List Information](#) on page 1468.

Variable Definitions

The following table defines parameters for the **multi-area 12 isid-list** command.

Variable	Value
<i>WORD<1-32></i>	Specifies the name of the I-SID list.
<i>WORD<1-1024></i>	Specifies the list of I-SID values.

Configure Multi-area SPB Layer 2 I-SID Redistribution

About This Task

Perform this procedure to configure Multi-area SPB Layer 2 I-SID redistribution.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```
2. Configure the Multi-area SPB Layer2 I-SID redistribution:

```
multi-area l2 redistribute i-sid {deny-all | permit-all} [except-isid-
list WORD<1-32>]
```
3. Verify the Multi-area SPB Layer2 I-SID redistribution configuration:

```
show isis multi-area l2 redistribute i-sid
```
4. Apply the Layer 2 I-SID redistribution:

```
isis multi-area l2 apply redistribute i-sid
```

Example

Configuring Multi-area SPB Layer 2 I-SID redistribution and exclude the I-SID values specified in the I-SID list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area l2 redistribute i-sid permit-all except-isid-list
testlist
Switch:1(config-isis)#show isis multi-area l2 redistribute i-sid
=====
MULTI AREA L2 ISID REDIST POLICY
=====
Permit          Except List Name
-----
permit-all      testlist
-----
Switch:1(config-isis)#isis multi-area l2 apply redistribute i-sid
```

Variable Definitions

The following table defines parameters for the **multi-area l2 redistribute i-sid** command.

Variable	Value
<i>deny-all</i>	Denies the Multi-area SPB IPv6 unicast redistribution configuration for the specified I-SID list name.
<i>except-isid-list</i> <i>WORD<1-32></i>	Specifies the name of the I-SID list.
<i>permit-all</i>	Permits the Multi-area SPB IPv6 unicast redistribution configuration for the specified I-SID list name.

Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution

Before You Begin

Configure IP Shortcuts on the boundary nodes.

About This Task

Perform this procedure to configure Multi-area SPB IPv4 unicast redistribution on the switch.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```

2. Create the Multi-area SPB IPv4 unicast redistribution instance:

```
multi-area ip redistribute unicast [home-to-remote | remote-to-home]
```



Note

If you do not specify the direction for redistribution then the system applies the redistribution in both directions by default.

3. (Optional) Configure a route policy to govern the redistribution:

```
multi-area ip redistribute unicast [home-to-remote | remote-to-home]
route-map WORD<1-64>
```

4. (Optional) Disable the transmission of route tags in Multi-area SPB:

```
multi-area ip redistribute unicast delete-tags
```

5. Enable the Multi-area SPB IPv4 unicast redistribution instance:

```
multi-area ip redistribute unicast [home-to-remote | remote-to-home]
enable
```

6. Verify the Multi-area SPB IPv4 unicast redistribution configuration:

```
show isis multi-area ip redistribute unicast [home-to-remote | remote-to-home]
```

7. Apply the Multi-area SPB IPv4 unicast redistribution:

```
isis multi-area ip apply redistribute unicast [home-to-remote |
remote-to-home]
```

Example

Configuring Multi-area SPB IPv4 unicast redistribution for home to remote direction:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area ip redistribute unicast home-to-remote
Switch:1(config-isis)#multi-area ip redistribute unicast home-to-remote route-map test
Switch:1(config-isis)#multi-area ip redistribute unicast home-to-remote enable
Switch:1(config-isis)#show isis multi-area ip redistribute home-to-remote
=====
                ISIS Multiarea Redistribute List for ip unicast - GlobalRouter
=====
DIRECTION          ENABLE      RPOLICY
-----
home-to-remote      TRUE        test
```

```
Switch:1(config-isis)#isis multi-area ip apply redistribute unicast home-to-remote
```

Variable Definitions

The following table defines parameters for the **multi-area ip redistribute unicast** and the **isis multi-area ip apply redistribute unicast** commands.

Variable	Value
<i>delete-tags</i>	Disables the transmission of route tags in Multi-area SPB. Use this parameter if platforms do not correctly calculate the TLV length in TLV 135. By default, route tags are transported in Multi-area SPB.
<i>enable</i>	Enables Multi-area SPB IPv4 unicast redistribution.
<i>home-to-remote</i>	Specifies the IPv4 unicast redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the IPv4 unicast redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD<1-64></i>	Specifies the name of the route policy.
<i>vrf WORD<1-16></i>	Specifies the unicast redistribution configuration for specific VRF instance.

Configure IS-IS Multi-area SPB IPv6 Unicast Redistribution

Before You Begin

Configure IP Shortcuts on the boundary nodes.

About This Task

Perform this procedure to configure Multi-area SPB IPv6 unicast redistribution on the switch.

Procedure

1. Enter IS-IS Router Configuration mode:


```
enable
configure terminal
router isis
```
2. Create the Multi-area SPB IPv6 unicast redistribution:


```
multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home]
```



Note

If you do not specify the direction for redistribution then the system applies the redistribution in both directions by default.

3. (Optional) Configure a route policy to govern the redistribution:

```
multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home]
route-map WORD<1-64>
```

4. Enable the Multi-area SPB IPv6 unicast redistribution instance:

```
multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home]
enable
```

5. Verify the Multi-area SPB IPv6 unicast redistribution configuration:

```
show isis multi-area ipv6 redistribute unicast [home-to-remote |
remote-to-home]
```

6. Apply the Multi-area SPB IPv6 unicast redistribution:

```
isis multi-area ipv6 apply redistribute unicast [home-to-remote |
remote-to-home]
```

Example

Configuring Multi-area SPB IPv6 unicast redistribution for home to remote direction:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area ipv6 redistribute unicast home-to-remote
Switch:1(config-isis)#multi-area ipv6 redistribute unicast home-to-remote route-map test
Switch:1(config-isis)#multi-area ipv6 redistribute unicast home-to-remote enable
Switch:1(config-isis)#show isis multi-area ipv6 redistribute home-to-remote
=====
                ISIS Multiarea Redistribute List for ipv6 unicast - GlobalRouter
=====
DIRECTION          ENABLE      RPOLICY
-----
home-to-remote      TRUE        test
-----
Switch:1(config-isis)#isis multi-area ipv6 apply redistribute unicast home-to-remote
```

Variable Definitions

The following table defines parameters for the **multi-area ipv6 redistribute unicast** and the **isis multi-area ipv6 apply redistribute unicast** commands.

Variable	Value
<i>enable</i>	Enables Multi-area SPB IPv6 unicast redistribution on the switch.
<i>home-to-remote</i>	Specifies the IPv6 unicast redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the IPv6 unicast redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD<1-64></i>	Specifies the name of the route policy.
<i>vrf WORD<1-16></i>	Specifies the IPv6 unicast redistribution configuration for a specific VRF instance.

*Configure IS-IS Multi-area SPB IPv4 Unicast Redistribution on a VRF Instance***Before You Begin**

Configure IP Shortcuts and layer 3 VSN on the Virtual Router Forwarding (VRF) instance.

About This Task

Perform this procedure to configure IS-IS Multi-area SPB IPv4 unicast redistribution on a specific VRF instance on the switch.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```

**Note**

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the IS-IS Multi-area SPB IPv4 unicast redistribution instance on the VRF:

```
isis multi-area ip redistribute unicast [home-to-remote | remote-to-home]
```

**Note**

If you do not specify the direction for redistribution then the system applies the redistribution in both directions by default.

3. (Optional) Configure a route policy to govern the redistribution:

```
isis multi-area ip redistribute unicast [home-to-remote | remote-to-home] route-map WORD<1-64>
```

4. (Optional) Disable the transmission of route tags in Multi-area SPB:

```
isis multi-area ip redistribute unicast delete-tags
```

5. Enable the Multi-area SPB IPv4 unicast redistribution instance on the VRF:

```
isis multi-area ip redistribute unicast [home-to-remote | remote-to-home] enable
```

6. Verify the Multi-area SPB IPv4 unicast redistribution configuration on the VRF:

```
show isis multi-area ip redistribute unicast [home-to-remote | remote-to-home] vrf WORD<1-16>
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the Multi-area SPB IPv4 unicast redistribution on the VRF:

```
isis multi-area ip apply redistribute unicast [home-to-remote | remote-to-home] vrf WORD<1-16>
```

Example

Configuring IS-IS Multi-area SPB IPv4 unicast redistribution for home to remote direction:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf VRF1
Switch:1(router-vrf)#isis multi-area ip redistribute unicast home-to-remote
Switch:1(router-vrf)#isis multi-area ip redistribute unicast home-to-remote route-map test
Switch:1(router-vrf)#isis multi-area ip redistribute unicast home-to-remote enable
Switch:1(router-vrf)#show isis multi-area ip redistribute unicast home-to-remote vrf VRF1
=====
                ISIS Multiarea Redistribute List for ip unicast - VRF VRF1
=====
DIRECTION          ENABLE      RPOLICY
-----
home-to-remote      TRUE        test
-----
Switch:1(router-vrf)#exit
Switch:1(config)#isis multi-area ip apply redistribute unicast home-to-remote vrf VRF1
```

Variable Definitions

The following table defines parameters for the **isis multi-area ip redistribute unicast** and the **isis multi-area ip apply redistribute unicast** commands.

Variable	Value
<i>delete-tags</i>	Disables the transmission of route tags in Multi-area SPB. Use this parameter if platforms do not correctly calculate the TLV length in TLV 135. By default, route tags are transported in Multi-area SPB.
<i>enable</i>	Enables IPv4 unicast redistribution on the VRF instance.
<i>home-to-remote</i>	Specifies the IPv4 unicast redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the IPv4 unicast redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD<1-64></i>	Specifies the name of the route policy.
<i>vrf WORD<1-16></i>	Specifies the IPv4 unicast redistribution configuration for specific VRF instance.

Configure IS-IS Multi-area SPB IPv6 Unicast Redistribution on a VRF Instance

Before You Begin

Configure IP Shortcuts and layer 3 VSN on the Virtual Router Forwarding (VRF) instance.

About This Task

Perform this procedure to configure IS-IS Multi-area SPB IPv6 unicast redistribution on a specific Virtual Router Forwarding (VRF) instance on the switch.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable

configure terminal

router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the IS-IS Multi-area SPB IPv6 unicast redistribution instance on the VRF:

```
isis multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home]
```



Note

If you do not specify the direction for redistribution then the system applies the redistribution in both directions by default.

3. (Optional) Configure a route policy to govern the redistribution:

```
isis multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home] route-map WORD<1-64>
```

4. Enable the Multi-area SPB IPv6 unicast redistribution instance on the VRF:

```
isis multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home] enable
```

5. Verify the Multi-area SPB IPv6 unicast redistribution configuration on the VRF:

```
show isis multi-area ipv6 redistribute unicast [home-to-remote | remote-to-home] vrf WORD<1-16>
```

6. Exit to Global Configuration mode:

```
exit
```

7. Apply the Multi-area SPB IPv6 unicast redistribution on the VRF:

```
isis multi-area ipv6 apply redistribute unicast [home-to-remote | remote-to-home] vrf WORD<1-16>
```

Example

Configuring IS-IS Multi-area SPB IPv6 unicast redistribution for home to remote direction:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf VRF1
Switch:1(router-vrf)#isis multi-area ipv6 redistribute unicast home-to-remote
Switch:1(router-vrf)#isis multi-area ipv6 redistribute unicast home-to-remote route-map test
Switch:1(router-vrf)#isis multi-area ipv6 redistribute unicast home-to-remote enable
Switch:1(router-vrf)#show isis multi-area ipv6 redistribute unicast home-to-remote vrf VRF1
=====
ISIS Multiarea Redistribute List for ipv6 unicast - VRF VRF1
```

```

=====
DIRECTION          ENABLE      RPOLICY
-----
home-to-remote     TRUE       test
-----
Switch:1(router-vrf)#exit
Switch:1(config)#isis multi-area ipv6 apply redistribute unicast home-to-remote vrf VRF1

```

Variable Definitions

The following table defines parameters for the **isis multi-area ipv6 redistribute unicast** and the **isis multi-area ipv6 apply redistribute unicast** commands.

Variable	Value
<i>enable</i>	Enables IPv6 unicast redistribution on the VRF instance.
<i>home-to-remote</i>	Specifies the IPv6 unicast redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the IPv6 unicast redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD<1-64></i>	Specifies the name of the route policy.
<i>vrf WORD<1-16></i>	Specifies the IPv6 unicast redistribution configuration for a specific VRF instance.

Configure Multi-area SPB DvR Backbone Redistribution



Note

This procedure does not apply to 5520 Series.

About This Task

Perform this procedure to configure Multi-area SPB Distributed Virtual Routing (DvR) backbone redistribution on the switch.

Procedure

1. Enter IS-IS Router Configuration mode:


```
enable
configure terminal
router isis
```
2. Configure IS-IS Multi-area SPB DvR backbone redistribution:


```
multi-area dvr redistribute backbone
```

Example

Configuring Multi-area SPB DvR backbone redistribution:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area dvr redistribute backbone

```

What to Do Next

To verify the configuration, see [Display IS-IS Multi-area SPB DvR Configuration](#) on page 1471.

Configure Multi-area SPB Layer 2 Multicast Snooping Redistribution

About This Task

Perform this procedure to configure Multi-area SPB layer 2 multicast snooping redistribution for specific I-SID values or a specific I-SID list.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

2. Create the Multi-area SPB layer 2 snooping multicast redistribution for a specific I-SID or I-SID list:

```
multi-area l2 redistribute snoop-multicast {home-to-remote | remote-to-home} {i-sid <1-16777215> | isid-list WORD<1-32>}
```

3. (Optional) Configure a route policy to govern the redistribution:

```
multi-area l2 redistribute snoop-multicast {home-to-remote | remote-to-home} {i-sid <1-16777215> | isid-list WORD<1-32>} route-map WORD<1-64>
```

4. Enable the Multi-area SPB layer 2 snooping multicast redistribution instance:

```
multi-area l2 redistribute snoop-multicast {home-to-remote | remote-to-home} {i-sid <1-16777215> | isid-list WORD<1-32>} enable
```

5. Verify the Multi-area SPB layer 2 snooping multicast redistribution configuration:

```
show isis multi-area l2 redistribute snoop-multicast
```

6. Apply the Multi-area SPB layer 2 snooping multicast redistribution:

```
isis multi-area l2 apply redistribute snoop-multicast [home-to-remote | remote-to-home]
```

Examples

Configuring Multi-area SPB layer 2 home to remote multicast snooping redistribution:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
```

Configuring Multi-area SPB layer 2 home to remote multicast snooping redistribution for I-SID 1100:

```
Switch:1(config-isis)#multi-area l2 redistribute snoop-multicast home-to-remote i-sid 1100
Switch:1(config-isis)#multi-area l2 redistribute snoop-multicast home-to-remote i-sid 1100
route-map test
Switch:1(config-isis)#multi-area l2 redistribute snoop-multicast home-to-remote i-sid
```

```
1100
enable
```

Configuring Multi-area SPB layer 2 home to remote multicast snooping redistribution for I-SID list List1:

```
Switch:1(config-isis)#multi-area l2 redistribute snoop-multicast home-to-remote isid-list List1
Switch:1(config-isis)#multi-area l2 redistribute snoop-multicast home-to-remote isid-list List1 route-map test
Switch:1(config-isis)#multi-area l2 redistribute snoop-multicast home-to-remote isid-list List1 enable
```

Verifying the Multi-area SPB layer 2 home to remote multicast snooping redistribution configuration:

```
Switch:1(config-isis)#show isis multi-area l2 redistribute snoop-multicast
=====
MULTI-AREA L2 SNOOP MULTICAST REDISTRIBUTE BY ISID
=====
ISID          DIRECTION          ENABLE    RPOLICY
-----
1100          home-to-remote      TRUE      test
=====
MULTI-AREA L2 SNOOP MULTICAST REDISTRIBUTE BY ISID LIST
=====
ISID LIST      DIRECTION          ENABLE    RPOLICY
-----
List1          home-to-remote      TRUE      test
-----
```

Applying the Multi-area SPB layer 2 home to remote multicast snooping redistribution:

```
Switch:1(config-isis)#isis multi-area l2 apply redistribute snoop-multicast home-to-remote
```

Variable Definitions

The following table defines parameters for the **multi-area l2 redistribute snoop-multicast** and the **isis multi-area l2 apply redistribute snoop-multicast** commands.

Variable	Value
<i>enable</i>	Enables Multi-area SPB layer 2 multicast snooping redistribution.
<i>i-sid</i> <1-16777215>	Specifies the I-SID value.
<i>isid-list</i> <i>WORD</i> <1-32>	Specifies the name of the I-SID list.
<i>home-to-remote</i>	Specifies the layer 2 multicast snooping redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the layer 2 multicast snooping redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD</i> <1-64>	Specifies the name of the route policy.

Configure Multi-area SPB IPv4 Routed Multicast Redistribution

Before You Begin

- Enable multicast on the boundary nodes.
- Configure Multiprotocol Label Switching Virtual Private Network (MVPN) on the boundary nodes, for layer 3 VSN.
- Configure the same policies as the boundary node on the peer boundary nodes.

About This Task

Perform this procedure to configure Multi-area SPB IPv4 routed multicast redistribution on the switch.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

2. Create the Multi-area SPB IPv4 routed multicast redistribution instance:

```
multi-area ip redistribute routed-multicast [home-to-remote | remote-to-home]
```



Note

If you do not specify the direction for redistribution then the system applies the redistribution in both directions by default.

3. (Optional) Configure a route policy to govern the redistribution:

```
multi-area ip redistribute routed-multicast [home-to-remote | remote-to-home] route-map WORD<1-64>
```

4. Enable the Multi-area SPB IPv4 routed multicast redistribution instance:

```
multi-area ip redistribute routed-multicast [home-to-remote | remote-to-home] enable
```

5. Verify the Multi-area SPB IPv4 routed multicast redistribution configuration:

```
show isis multi-area ip redistribute routed-multicast [home-to-remote | remote-to-home]
```

6. Apply the Multi-area SPB IPv4 routed multicast redistribution:

```
isis multi-area ip apply redistribute routed-multicast [home-to-remote | remote-to-home]
```

Example

Configuring Multi-area SPB IPv4 routed multicast redistribution for home to remote direction:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area ip redistribute routed-multicast home-to-remote
```

```

Switch:1(config-isis)#multi-area ip redistribute routed-multicast home-to-remote route-
map test
Switch:1(config-isis)#multi-area ip redistribute routed-multicast home-to-remote enable
Switch:1(config-isis)#show isis multi-area ip redistribute routed-multicast home-to-remote
=====
      ISIS Multiarea Redistribute List for routed multicast - GlobalRouter
=====
DIRECTION                ENABLE      RPOLICY
-----
home-to-remote           TRUE        test
-----
Switch:1(config-isis)#isis multi-area ip apply redistribute routed-multicast home-to-
remote

```

Variable Definitions

The following table defines parameters for the **multi-area ip redistribute routed-multicast** and the **isis multi-area ip apply redistribute routed-multicast** command.

Variable	Value
<i>enable</i>	Enables Multi-area SPB IPv4 routed multicast redistribution.
<i>home-to-remote</i>	Specifies the IPv4 routed multicast redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the IPv4 routed multicast redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD<1-64></i>	Specifies the name of the route policy.
<i>vrf WORD<1-16></i>	Specifies the multicast routing redistribution configuration for specific VRF instance.

Configure IS-IS Multi-area SPB IPv4 Routed Multicast Redistribution on a VRF Instance

Before You Begin

- Enable multicast on the boundary nodes.
- Configure Multiprotocol Label Switching Virtual Private Network (MVPN) on the boundary nodes, for layer 3 VSN.
- Configure the same policies as the boundary node on the peer boundary nodes.

About This Task

Perform this procedure to configure IS-IS Multi-area SPB IPv4 routed multicast redistribution on a specific Virtual Router Forwarding (VRF) instance on the switch.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the IS-IS Multi-area SPB IPv4 routed multicast redistribution instance on the VRF:

```
isis multi-area ip redistribute routed-multicast [home-to-remote |  
remote-to-home]
```



Note

If you do not specify the direction for redistribution then the system applies the redistribution in both directions by default.

3. (Optional) Configure a route policy to govern the redistribution:

```
isis multi-area ip redistribute routed-multicast [home-to-remote |  
remote-to-home] route-map WORD<1-64>
```

4. Enable the Multi-area SPB IPv4 routed multicast redistribution instance on the VRF:

```
isis multi-area ip redistribute routed-multicast [home-to-remote |  
remote-to-home] enable
```

5. Verify the Multi-area SPB IPv4 routed multicast redistribution configuration on the VRF:

```
show isis multi-area ip redistribute routed-multicast [home-to-remote  
| remote-to-home] [vrf WORD<1-16> | vrfids WORD<0-512>]
```

6. Exit to Global Configuration mode:

```
exit
```

7. Apply the Multi-area SPB IPv4 routed multicast redistribution on the VRF:

```
isis multi-area ip apply redistribute routed-multicast [home-to-remote  
| [remote-to-home] vrf WORD<1-16>
```

Example

Configuring IS-IS Multi-area SPB IPv4 routed multicast redistribution for home to remote direction on VRF VRF1:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf VRF1
Switch:1(router-vrf)#isis multi-area ip redistribute routed-multicast home-to-remote
Switch:1(router-vrf)#isis multi-area ip redistribute routed-multicast home-to-remote
route-map test
Switch:1(router-vrf)#isis multi-area ip redistribute routed-multicast home-to-remote
enable
Switch:1(router-vrf)#show isis multi-area ip redistribute routed-multicast home-to-remote
```

```

vrf VRF1
=====
                ISIS Multiarea Redistribute List for routed multicast - VRF VRF1
=====
DIRECTION          ENABLE      RPOLICY
-----
home-to-remote      TRUE        test
-----
Switch:1(router-vrf)#exit
Switch:1(config)#isis multi-area ip apply redistribute routed-multicast home-to-remote
vrf VRF1

```

Variable Definitions

The following table defines parameters for the **isis multi-area ip redistribute routed-multicast** and the **isis multi-area ip apply redistribute routed-multicast** commands.

Variable	Value
<i>enable</i>	Enables IPv4 routed multicast redistribution on the VRF instance.
<i>home-to-remote</i>	Specifies the IPv4 routed multicast redistribution configuration for home to remote direction.
<i>remote-to-home</i>	Specifies the IPv4 routed multicast redistribution configuration for remote to home direction.
<i>route-map</i> <i>WORD<1-64></i>	Specifies the name of the route policy.
<i>vrf WORD<1-16></i>	Specifies the multicast routing redistribution configuration for specific VRF instance.

Forward UNI Traffic Without Home-Area IS-IS Adjacency

About This Task

Use this feature to forward traffic from UNI ports on the Multi-area SPB designated boundary node into the remote-area, and receive traffic from the remote area that is destined for the UNIs ports, without requiring an established IS-IS adjacency in the home area.



Important

To avoid packet loops back into the home area, this feature does not operate on non-designated boundary nodes.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```
2. Configure the multi-area flags functionality:

```
multi-area flags home-always-up
```

Configure Multi-area SPB IPv4 Static I-SID Routed Multicast Redistribution

Before You Begin

- Enable multicast on the boundary nodes.
- Configure Multicast Virtual Private Network (MVPN) on the boundary nodes, for layer 3 VSN.

About This Task

Perform this procedure to configure Multi-area SPB IPv4 static I-SID routed multicast redistribution on the switch.



Note

- Routed multicast redistribution is available for IPv4 only.
- Multi-area SPB boundary nodes can redistribute static data I-SIDs, and the static data I-SID multicast redistribution policy takes precedence over the Layer 2 multicast redistribution.
- Multi-area SPB boundary nodes with remote IS-IS enabled do not support IP SPB Multicast Policy .

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

2. Configure the Multi-area SPB IPv4 static I-SID list to redistribute on the switch:

```
multi-area ip redistribute static-isid-routed-multicast [home-to-remote | remote-to-home] isid-list WORD<1-32>
```

The system displays the warning message while configuring the Multi-area SPB IPv4 static I-SID routed multicast redistribution:

```
WARNING: L3 static i-sid routed multicast will not be inter-area redistributed until  
apply command is issued after enable command.
```

3. Enable the Multi-area SPB IPv4 static I-SID routed multicast redistribution instance:

```
multi-area ip redistribute static-isid-routed-multicast [home-to-remote | remote-to-home] enable
```

4. Verify the Multi-area SPB IPv4 static I-SID routed multicast redistribution configuration:

```
show isis multi-area ip redistribute static-isid-routed-multicast  
[home-to-remote | remote-to-home]
```

5. Apply the Multi-area SPB IPv4 static I-SID routed multicast redistribution:

```
isis multi-area ip apply redistribute static-isid-routed-multicast
```

Example

Configuring Multi-area SPB IPv4 static I-SID list to redistribute for remote to home direction:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#multi-area ip redistribute static-isis-routed-multicast remote-to-home isid-list
list2
WARNING: L3 static i-sid routed multicast will not be inter-area redistributed until apply command is
issued after enable command.
Switch:1(config-isis)#multi-area ip redistribute static-isis-routed-multicast [home-to-remote | remote-
to-home] enable
Switch:1(config-isis)#show isis multi-area ip redistribute static-isis-routed-multicast remote-to-home
=====
MULTI-AREA IP REDISTRIBUTE STATIC ISID ROUTED MULTICAST - ALL VRF
=====
DIRECTION          ENABLE      ISID LIST
-----
remote-to-home     TRUE       list2

Switch:1(config-isis)#isis multi-area ip apply redistribute static-isis-routed-multicast
```

Variable Definitions

The following table defines parameters for the **multi-area ip redistribute static-isis-routed-multicast** command.

Variable	Value
<i>enable</i>	Enables Multi-area SPB IPv4 static I-SID routed multicast redistribution.
<i>home-to-remote</i>	Specifies the IPv4 static I-SID routed multicast redistribution configuration for home to remote direction.
<i>isid-list</i> WORD<1-32>	Specifies the name of the I-SID list.
<i>remote-to-home</i>	Specifies the IPv4 static I-SID routed multicast redistribution configuration for remote to home direction.

Create a Virtual NNI Link Between Multi-Area Boundary Nodes

Before You Begin

- Create basic SPBM and IS-IS infrastructure.
- Configure a CLIP interface.

About This Task

Perform the following procedure to create a virtual NNI link between Multi-area boundary nodes.

Procedure

1. Enter Loopback Interface Configuration mode

```
enable
```



```
configure terminal
```



```
interface Loopback <1-256>
```
2. Configure the loopback interface IP address to use as the source IP address for the Multi-area virtual NNI link:

```
ip address {<A.B.C.D/x> | <A.B.C.D> <A.B.C.D>} multi-area-virtual-link  
<home | remote> [name WORD<0-64>]
```
3. Exit the Loopback Configuration mode:

```
exit
```
4. Create a IS-IS logical interface for the Multi-area virtual NNI link:

```
logical-intf isis <1-255> dest-ip <A.B.C.D> multi-area-virtual-link  
[name WORD<1-64>]
```
5. Configure IS-IS on the logical interface:



Note

If you configure the loopback interface for the Multi-area virtual NNI link in the home area, you must configure IS-IS on the logical interface in the remote area, and vice versa.

- a. Create an IS-IS interface:

```
isis [remote]
```
 - b. Enable the SPBM instance on the IS-IS interface:

```
isis [remote] spbm <1-100>
```
 - c. Enable the IS-IS interface:

```
isis [remote] enable
```
6. Verify the configuration using the following commands:
 - `show interfaces loopback`
 - `show isis logical-interface`

Example

Configure a Multi-area virtual link NNI for the home area. To do this, create the loopback source to be redistributed in the remote area and then enable IS-IS on a virtual-link logical interface in the home area.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 2.2.2.2/32 multi-area-virtual-link remote
exit
Switch:1(config)#logical-intf isis 2 dest-ip 4.4.4.4 multi-area-virtual-link
Switch:1(config-isis-2-4.4.4.4)#isis
Switch:1(config-isis-2-4.4.4.4)#isis spbm 1
Switch:1(config-isis-2-4.4.4.4)#isis enable
exit
```

Configure a Multi-area virtual NNI link for the remote area. To do this, create the loopback source to be redistributed in the home area and then enable IS-IS on a virtual-link logical interface in the remote area.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 2.2.2.2/32 multi-area-virtual-link home
exit
Switch:1(config)#logical-intf isis 2 dest-ip 4.4.4.4 multi-area-virtual-link
Switch:1(config-isis-2-4.4.4.4)#isis remote
Switch:1(config-isis-2-4.4.4.4)#isis remote spbm 1
Switch:1(config-isis-2-4.4.4.4)#isis remote enable
exit
```

The following example displays loopback interface information:

```
Switch:1#show interfaces loopback
```

Circuitless IP Interface - GlobalRouter												
INTF ID	IP_ADDRESS	NET_MASK	OSPF STATUS	PIM STATUS	AREA_ID	IF INDX	IP NAME	MULTI-AREA VIRTUAL LINK				
1	2.2.2.2	255.255.255.255	disable	disable	0.0.0.0	1344	Virtual-link Remote					
2	2.3.4.5	255.255.255.255	disable	disable	0.0.0.0	1345						
15	192.0.2.2	255.255.255.0	disable	disable	0.0.0.0	1358	EXTR					
Loopback IPv6 Interface												
IF INDX	VRF NAME	Descr	VLAN ADDRESS	PHYSICAL ADDRESS	ADMIN STATE	OPER STATE	TYPE	MTU	HOP LMT	REACHABLE TIME	RETRANSMIT TIME	IPSEC STATE
1234	GREEN	CLIPv6-11	00:00:00:00:00:0b		enable	up	ETHER	1500	64	30000	1000	disable
Loopback IPv6 Address												
IPV6 ADDRESS/PREFIX	LENGTH	LOOPBACK-ID	TYPE	ORIGIN	STATUS	VALID LIFETIME	PREF LIFETIME	NAME				
2001:DB8:2000::1/128		C-11	UNICAST	MANUAL	PREFERRED	INF	INF	EXTRSER200				

The following example displays logical interface information:

```
Switch:1#show isis logical-interface
```

ISIS Logical Interfaces												
IFIDX	NAME	ENCAP TYPE	L2_INFO PORT/MLT	VIDS (PRIMARY)	TUNNEL DEST-IP	L3_TUNNEL PORT/MLT	NEXT_HOP_INFO VLAN	BFD STATUS	TUNNEL SRC-IP	ORIGIN	ISIS MTU	SDWAN OPER STATE
1	SD-WAN-1	IP	--	--	192.0.2.3	Port1/44	4047	sd-wan	disabled	192.0.2.1	ZTF	1400 UP
2	SPBoIP_T1	IP	--	--	192.0.2.15	Port1/25	500	vrf23	disabled	192.0.2.16	CONFIG	1000 N/A
3	SPBoIP_T2	IP	--	--	192.0.2.224	MLT10	2	vrf24	disabled	192.0.2.22	CONFIG	1600 N/A
4	Virtual-link-4	IP	--	--	4.4.4.4	PortRX-NNI	4051	GlobalRouter	disabled	2.2.2.2	CONFIG	1600 N/A
5	Virtual-link-5	IP	--	--	5.5.5.5	Null	0	GlobalRouter	disabled	2.2.2.2	CONFIG	1600 N/A
5 out of 5 Total Num of Logical ISIS interfaces												

Variable Definitions

The following table defines parameters for the **ip address** command.

Variable	Value
<A.B.C.D/X> <A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask in the format A.B.C.D/X or A.B.C.D A.B.C.D.
<i>multi-area-virtual-link</i>	Specifies that the source IP address is used for a Multi-area virtual NNI link.
<home remote>	Specifies the transport area for the Multi-area virtual NNI link. Note: If you configure the loopback interface for the virtual NNI link in the home area, you must configure the IS-IS logical interface for the virtual NNI link in the remote area, and vice versa.
<i>name WORD<0-64></i>	Specifies a name associated with the IP address. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 5720 Series, 7520 Series, and 7720 Series, the system autogenerates the name and displays the name as <i>Virtual-link</i> .

The following tables define parameters for the **logical-intf isis** command.

Variable	Value
<1-255>	Specifies the index number that uniquely identifies this logical interface.
<i>dest-ip <A.B.C.D></i>	Specifies the destination IP address of the Multi-area virtual NNI link. Note: This IP address is the loopback source IP address.
<i>multi-area-virtual-link</i>	Specifies that the logical interface is used for a Multi-area virtual NNI link.
<i>name WORD<1-64></i>	Specifies a name associated with the logical interface. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 5720 Series, 7520 Series, and 7720 Series, the system autogenerates the name and displays the name as <i>Virtual-link-x</i> , where x is the logical interface ID.

The following tables define parameters for the **isis** command.

Variable	Value
<i>remote</i>	Specifies the remote area as the IS-IS interface. If the loopback interface for the virtual NNI link is configured in the remote area, omit this parameter.
<i>enable</i>	Enables the IS-IS interface for the virtual NNI link. The default is disabled.
<i>spbm <1-100></i>	Specifies the SPBM instance on the IS-IS interface.

Multi-area SPB Verification using CLI

Perform the procedures in this section to verify the Multi-area SPB configuration on the switch using the Command Line Interface (CLI).



Note

The following procedures do not apply to 4220 Series, 5320 Series, or 5420 Series.

Display the Remote IS-IS SPBM Nickname Configuration

About This Task

Perform this procedure to display the remote Intermediate-System-to-Intermediate-System (IS-IS) Shortest Path Bridging MAC (SPBM) nickname configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the remote IS-IS SPBM nickname configuration:

```
show isis spbm nick-name remote
```

Example

Displaying the remote IS-IS SPBM nickname configuration:

```
Switch:1>show isis spbm nick-name remote
=====
ISIS SPBM NICK-NAME
=====
```

LSP ID	LIFETIME	NICK-NAME	VIRTUAL-BMAC	HOST-NAME	AREA	AREA NAME
0001.0004.0005.00-00	1057	0.01.45	00:55:11:55:22:55	Switch-A	REMOTE	area-9.00.01
209e.f77e.2086.00-00	432	0.84.78	00:00:00:00:00:00	Switch-B	REMOTE	area-9.00.01
5005.5005.5005.00-00	1185	0.01.55	00:55:11:55:22:55	Switch-C	REMOTE	area-9.00.01
840c.8404.8404.00-00	528	8.84.04	00:00:00:00:00:00	Switch-D	REMOTE	area-9.00.01
9200.02ff.fff0.00-00	432	9.00.02	00:00:00:00:00:00	Switch-E	REMOTE	area-9.00.01
f46e.959f.8c86.00-00	363	0.01.40	00:00:00:00:00:00	Switch-F	REMOTE	area-9.00.01

```
Remote: 6 out of 6 Total Num of Entries
=====
```


*Display IS-IS Remote Area Configuration***About This Task**

Perform this procedure to display the IS-IS remote area configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IS-IS remote configuration:

```
show isis remote
```

Example

Displaying IS-IS remote area information:

```
Switch:1>show isis remote
=====
                        ISIS Remote Area Info
=====
                        AdminState : enabled
                        System ID : 209e.f77e.2086
                        Num of Interfaces : 4
                        Num of Area Addresses : 1
                        Dynamically Learned Area : 49.0002
                        Multi-Area OperState : designated
                        Designated BN Hostname: BN4435
                        Designated BN Home System ID : 0010.7400.4435
                        Designated BN Remote System ID : 0020.7400.4435
```

*Display IS-IS Area Configuration***About This Task**

Perform this procedure to display the IS-IS area configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IS-IS area configuration:

```
show isis area
```

Example

Displaying IS-IS area configuration:

```
Switch:1>show isis area
=====
                        ISIS Area Address
=====
ADDRESS                ORIGIN  AREA   AREA-NAME
-----
49.0030                 config  HOME   30
49.0050                 config  REMOTE 50
=====
```

*Display IS-IS Area Virtual Node Configuration***About This Task**

Perform this procedure to display the IS-IS area virtual node configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IS-IS area virtual node configuration:

```
show isis area-vnode
```

Example

Displaying IS-IS area virtual node configuration:

```
Switch:1>show isis area-vnode
=====
ISIS SPBM Multi-Area VNode Info
=====
VNODE          VNODE          VNODE          REPRESENTED    REPRESENTED    VNODE
SYSTEM-ID      NICK-NAME      HOST-NAME      AREA-ADDRESS   AREA           USED-IN-AREA
-----
9200.30ff.fff0  9.00.30        vn-30          49.0030        HOME           REMOTE
9200.50ff.fff0  9.00.50        vn-50          49.0050        REMOTE        HOME
```

*Display IS-IS Multi-area SPB Layer 2 I-SID List Information***About This Task**

Perform this procedure to display the IS-IS Multi-area SPB layer 2 I-SID list information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the I-SID list information on the switch:

```
show isis multi-area l2 isid-list WORD<1-32>
```

Example

Displaying the IS-IS Multi-area SPB layer 2 I-SID list information:

```
Switch:1>show isis multi-area l2 isid-list test
=====
MULTI AREA L2 ISID LIST
=====
List Name          I-SID / I-SID RANGE
-----
test                50510-50513
All 1 out of 1 Total Num of Isid Lists displayed
```

*Display IS-IS Multi-area SPB I-SID Redistribution Configuration***About This Task**

Perform this procedure to display the IS-IS Multi-area SPB layer 2 I-SID redistribution configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IS-IS Multi-area SPB layer 2 I-SID redistribution configuration:

```
show isis multi-area l2 redistribute i-sid
```

Example

Displaying the IS-IS Multi-area SPB layer 2 I-SID redistribution configuration:

```
Switch:1>show isis multi-area l2 redistribute i-sid
=====
MULTI AREA L2 ISID REDIST POLICY
=====
Permit          Except List Name
-----
permit-all      TestList1
```

Display IS-IS Multi-area SPB IPv4 Redistribution Configuration

About This Task

Perform this procedure to display the IS-IS Multi-area SPB IPv4 redistribution configuration for home to remote, remote to home, unicast, routed multicast, specific Virtual Router Forwarding (VRF) instance, and VRF IDs on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IS-IS Multi-area SPB IPv4 home to remote redistribution configuration:


```
show isis multi-area ip redistribute [home-to-remote [vrf WORD<1-16> |
vrfids WORD<0-512>]]
```
3. Display the IS-IS Multi-area SPB IPv4 remote to home redistribution configuration:


```
show isis multi-area ip redistribute [remote-to-home [vrf WORD<1-16> |
vrfids WORD<0-512>]]
```
4. Display the IS-IS Multi-area SPB IPv4 routed multicast redistribution configuration:


```
show isis multi-area ip redistribute routed-multicast [home-to-remote
[vrf WORD<1-16> | vrfids WORD<0-512>] | remote-to-home [vrf WORD<1-16>
| vrfids WORD<0-512>] | vrf WORD<1-16> | vrfids WORD<0-512>]]
```
5. Display the IS-IS Multi-area SPB IPv4 unicast redistribution configuration:


```
show isis multi-area ip redistribute unicast [home-to-remote [vrf
WORD<1-16> | vrfids WORD<0-512>] | remote-to-home [vrf WORD<1-16> |
vrfids WORD<0-512>] | vrf WORD<1-16> | vrfids WORD<0-512>]]
```
6. Display the IS-IS Multi-area SPB IPv4 redistribution configuration for VRF instance:


```
show isis multi-area ip redistribute vrf WORD<1-16>
```
7. Display the IS-IS Multi-area SPB IPv4 redistribution configuration for VRF ID:


```
show isis multi-area ip redistribute vrfids WORD<0-512>
```

Example

Displaying the IS-IS Multi-area SPB IPv4 unicast redistribution configuration on the switch:

```
Switch:1>show isis multi-area ip redistribute unicast
=====
                ISIS Multiarea Redistribute List for ip unicast - GlobalRouter
=====
DIRECTION          ENABLE      RPOLICY    DELETE-TAGS
-----
home-to-remote      TRUE                FALSE
remote-to-home      TRUE                FALSE
```

Display IS-IS Multi-area SPB IPv6 Redistribution Configuration

About This Task

Perform this procedure to display the IS-IS Multi-area SPB IPv6 redistribution configuration for home to remote, remote to home, unicast, specific Virtual Router Forwarding (VRF) instance, and VRF IDs on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IS-IS Multi-area SPB IPv6 home to remote redistribution configuration:

```
show isis multi-area ipv6 redistribute [home-to-remote [vrf WORD<1-16>
| vrfids WORD<0-512>]]
```
3. Display the IS-IS Multi-area SPB IPv6 remote to home redistribution configuration:

```
show isis multi-area ipv6 redistribute [remote-to-home [vrf WORD<1-16>
| vrfids WORD<0-512>]]
```
4. Display the IS-IS Multi-area SPB IPv6 unicast redistribution configuration:

```
show isis multi-area ipv6 redistribute unicast [home-to-remote [vrf
WORD<1-16> | vrfids WORD<0-512>] | remote-to-home [vrf WORD<1-16> |
vrfids WORD<0-512>] | vrf WORD<1-16> | vrfids WORD<0-512>]]
```
5. Display the IS-IS Multi-area SPB IPv6 redistribution configuration for VRF instance:

```
show isis multi-area ipv6 redistribute [vrf WORD<1-16>]
```
6. Display the IS-IS Multi-area SPB IPv6 redistribution configuration for VRF ID:

```
show isis multi-area ipv6 redistribute [vrfids WORD<0-512>]
```

Example

Displaying the IS-IS Multi-area SPB IPv6 unicast redistribution configuration:

```
Switch:1>show isis multi-area ipv6 redistribute unicast
=====
                ISIS Multiarea Redistribute List for ipv6 unicast - GlobalRouter
=====
DIRECTION          ENABLE      RPOLICY
-----
home-to-remote      TRUE        Test1
remote-to-home      TRUE        Test2
```

Display IS-IS Multi-area SPB DvR Configuration**Note**

This procedure does not apply to 5520 Series.

About This Task

Perform this procedure to display the IS-IS Multi-area SPB Distributed Virtual Routing (DvR) redistribution configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IS-IS Multi-area SPB DvR redistribution configuration:

```
show isis multi-area dvr redistribute
```

Example

Displaying the IS-IS Multi-area SPB DvR redistribution configuration on the switch:

```
Switch:1>show isis multi-area dvr redistribute
```

```
=====
MULTI-AREA DVR BACKBONE REDISTRIBUTE
=====
```

```
Admin State
```

```
Enabled
=====
```

Display IS-IS Multi-area SPB Layer 2 Multicast Snooping Redistribution Configuration**Note**

The following procedure only applies to 5520 Series and 5720 Series.

About This Task

Perform this procedure to display the IS-IS Multi-area SPB layer 2 multicast snooping redistribution configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the multicast snooping redistribution configuration:

```
show isis multi-area l2 redistribute snoop-multicast
```

Example

Displaying the IS-IS Multi-area SPB layer 2 multicast snooping redistribution configuration on the switch:

```
Switch:1>show isis multi-area l2 redistribute snoop-multicast
```

```
=====
MULTI-AREA L2 SNOOP MULTICAST REDISTRIBUTE BY ISID
=====
```

```
ISID          DIRECTION          ENABLE          RPOLICY
=====
```

```

1100      home-to-remote      TRUE      test
=====
MULTI-AREA L2 SNOOP MULTICAST REDISTRIBUTE BY ISID LIST
=====
ISID LIST      DIRECTION      ENABLE      RPOLICY
-----
List1          home-to-remote      TRUE      test
-----

```

Multi-area SPB Configuration using EDM

Perform the procedures in this section to configure Multi-area SPB parameters like area virtual node, remote Intermediate-System-to-Intermediate-System (IS-IS), remote Shortest Path Bridging MAC (SPBM), layer 2 unicast and multicast redistribution, and IPv4 and IPv6 route redistribution on the switch using the Enterprise Device Manager (EDM).



Note

The following procedures apply to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.

Configure Multi-area SPB Layer 2 I-SID List

About This Task

Perform this procedure to configure layer 2 I-SID list.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 ISID List** tab.
4. Select **Insert**.
5. Configure the layer 2 I-SID list parameters.
6. Select **Insert**.

L2 I-SID List Field Descriptions

Use data in the following table to use the **L2 I-SID List** tab.

Name	Description
Name	Specifies the name of the layer 2 I-SID list.
RangeStart	Specifies the I-SID value as the starting range for the layer 2 I-SID list.
RangeEnd	Specifies the I-SID value as the ending range for the layer 2 I-SID list.

Configure Layer 2 Multi-area SPB I-SID Redistribution

About This Task

Perform this procedure to configure Multi-area SPB Layer 2 redistribution.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 ISID Redist** tab.
4. Select **Insert**.
5. Configure the layer 2 I-SID redistribution parameters.
6. Select **Insert**.

L2 I-SID Redistribute Field Descriptions

Use the data in the following table to use the **L2 ISID Redistribute** tab.

Name	Description
Type	Specifies the layer 2 Redistribution type.
Permit	Configures the Multi-area SPB layer 2 I-SID redistribution policy as permit or deny.
ExceptListName	Configures the name of layer 2 I-SID list. The system does not apply the Multi-area SPB layer 2 I-SID redistribution policy to I-SID values or the range of I-SID values in the layer 2 I-SID list.

Configure Multi-area SPB Layer 2 Multicast I-SID Redistribution for Home Area

About This Task

Perform this procedure to configure the Multi-area SPB layer 2 multicast redistribution for specific I-SID value(s) in the home area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 MC Redist Home ISID** tab.
4. Select **Insert**.
5. In the **I-sid** field, type the I-SID value.
6. Select **Enable**.
7. (Optional) In the **PolicyName** field, type the name of the redistribution policy.
8. Select **Insert**.

L2 MC Redist Home ISID Field Descriptions

Use the data in the following table to use the **L2 MC Redist Home ISID** tab.

Name	Description
I-sid	Specifies the I-SID value.
Enable	Enables layer 2 multicast redistribution for specific I-SID value in the home area.
PolicyName	Specifies the name of the redistribution policy.

*Configure Multi-area SPB L2 Multicast I-SID List Redistribution for Home Area***About This Task**

Perform this procedure to configure the Multi-area SPB layer 2 multicast I-SID list redistribution for the home area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 MC Redist Home ISID List** tab.
4. Select **Insert**.
5. In the **ISID List Name** field, type the name of the I-SID list.
6. Select **Enable**.
7. (Optional) In the **PolicyName** field, type the policy name.
8. Select **Insert**.

L2 MC Redist Home ISID List Field Descriptions

Use the data in the following table to use the **L2 MC Redist Home ISID List** tab.

Name	Description
ISID List Name	Specifies the name of the I-SID list.
Enable	Enables Multi-area SPB layer 2 multicast redistribution for the specific I-SID list in the home area.
PolicyName	Specifies the name of the redistribution policy.

*Apply Multi-area SPB Layer 2 Multicast Redistribution in Home Area***About This Task**

Perform this procedure to apply the Multi-area SPB layer 2 multicast redistribution that you configure in the home area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 MC Apply Home** tab.
4. In the **McastApply** field, select **apply**.

L2 MC Apply Home Field Descriptions

Use data in the following table to use the **L2 MC Apply Home** tab.

Name	Description
McastApply	Applies or disables the Multi-area SPB layer 2 multicast redistribution in the home area. <ul style="list-style-type: none"> noAction apply

Configure Multi-area SPB Layer 2 Multicast I-SID Redistribution for Remote Area

About This Task

Perform this procedure to configure the Multi-area SPB layer 2 multicast redistribution for specific I-SID value(s) in the remote area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 MC Redist Remote ISID** tab.
4. Select **Insert**.
5. In the **I-sid** field, type the I-SID value.
6. Select **Enable**.
7. (Optional) In the **PolicyName** field, type the name of the redistribution policy.
8. Select **Insert**.

L2 MC Redist Remote ISID Field Descriptions

Use the data in the following table to use the **L2 MC Redist Remote ISID** tab.

Name	Description
I-sid	Specifies the I-SID value.
Enable	Enables layer 2 multicast redistribution for specific I-SID value in the remote area.
PolicyName	Specifies the name of the redistribution policy.

Configure Multi-area SPB L2 Multicast I-SID List Redistribution for Remote Area

About This Task

Perform this procedure to configure the Multi-area SPB layer 2 multicast I-SID list redistribution for the remote area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L2 MC Redist Remote ISID List** tab.
4. Select **Insert**.
5. In the **ISID List Name** field, type the name of the I-SID list.
6. Select **Enable**.

7. (Optional) In the **PolicyName** field, type the policy name.
8. Select **Insert**.

L2 MC Redist Remote ISID List Field Descriptions

Use the data in the following table to use the **L2 MC Redist Remote ISID List** tab.

6

Name	Description
ISID List Name	Specifies the name of the I-SID list.
Enable	Enables Multi-area SPB layer 2 multicast redistribution for the specific I-SID list in the remote area.
PolicyName	Specifies the name of the redistribution policy.

Apply Multi-area SPB Layer 2 Multicast Redistribution in Remote Area

About This Task

Perform this procedure to apply the Multi-area SPB layer 2 multicast redistribution that you configure in the remote area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**
3. Select the **L2 MC Apply Remote** tab.
4. In the **McastApply** field, select **apply**.

L2 MC Apply Remote Field Descriptions

Use data in the following table to use the **L2 MC Apply Remote** tab.

Name	Description
McastApply	Applies or disables the Multi-area SPB layer 2 multicast redistribution in the remote area. <ul style="list-style-type: none">• noAction• apply

Configure Multi-area SPB Layer 3 Redistribution in Home Area

About This Task

Perform this procedure to configure Multi-area SPB Layer 3 redistribution in the home area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L3 Redist Home** tab.

4. Select **Insert**.
5. In **VrfId**, type the Virtual Router Forwarding ID.
6. In **Type**, select the Layer 3 redistribution type.
7. (Optional) In **RoutePolicy**, type the name of the redistribution policy.
8. (Optional) Select **DeleteTags** to disable the transmission of route tags in Multi-area SPB.
9. Select **Enable**.
10. In **Apply**, select **apply**.
11. Select **Insert**.

L3 Redist Home Field Descriptions

Use the data in the following table to use the **L3 Redist Home** tab.

Name	Description
VrfId	Specifies the Virtual Router Forwarding ID.
Type	Specifies the layer 3 redistribution type. The options are: <ul style="list-style-type: none"> • ip-ucast - IPv4 unicast • ip-mcast - IPv4 multicast • ipv6-ucast - IPv6 unicast
RoutePolicy	Associates a route policy to a specific layer 3 redistribution entry in the home area.
Enable	Enables layer 3 redistribution in the home area. The default value is disabled.
Apply	Applies the Multi-area SPB layer 3 redistribution filters in the home area. The default value is noAction.
DeleteTags	Disables the transmission of route tags in Multi-area SPB. Select this parameter if platforms do not correctly calculate the TLV length in TLV 135. By default, route tags are transported in Multi-area SPB.

Configure Multi-area SPB Layer 3 Redistribution in Remote Area

About This Task

Perform this procedure to configure Multi-area SPB Layer 3 redistribution in the remote area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L3 Redist Remote** tab.
4. Select **Insert**.
5. In **VrfId**, type the Virtual Router Forwarding ID.
6. In **Type**, select the Layer 3 redistribution type.
7. (Optional) In **RoutePolicy**, type the name of the redistribution policy.
8. (Optional) Select **DeleteTags** to disable the transmission of route tags in Multi-area SPB.

9. Select **Enable**.
10. In **Apply**, select **apply**.
11. Select **Insert**.

L3 Redist Remote Field Descriptions

Use the data in the following table to use the **L3 Redist Remote** tab.

Name	Description
Vrflid	Specifies the Virtual Router Forwarding ID.
Type	Specifies the layer 3 redistribution type. The options are: <ul style="list-style-type: none"> • ip-ucast - IPv4 unicast • ip-mcast - IPv4 multicast • ipv6-ucast - IPv6 unicast
RoutePolicy	Associates a route policy to a specific layer 3 redistribution entry in the remote area.
Enable	Enables layer 3 redistribution in the remote area. The default value is disabled.
Apply	Applies the Multi-area SPB layer 3 redistribution filters in the remote area. The default value is noAction.
DeleteTags	Disables the transmission of route tags in Multi-area SPB. Select this parameter if platforms do not correctly calculate the TLV length in TLV 135. By default, route tags are transported in Multi-area SPB.

Configure Multi-area SPB Layer 3 Static I-SID Routed Multicast Redistribution Policy in Home Area

About This Task

Perform this procedure to configure Multi-area SPB Layer 3 static I-SID routed multicast redistribution policy in the home area.



Note

Routed multicast redistribution is available for IPv4 only. With the MAC learning disabled, the static data I-SID multicast redistribution policy takes precedence over the Layer 2 multicast redistribution.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L3 MC Redist Home Static ISID List** tab.
4. Select **Insert**.
5. In the **Type** field, select **staticisid**.
6. Select **Enable**.
7. In the **IsidListName** field, type the name of the I-SID list.
8. In the **Apply** field, select **apply**.
9. Select **Insert**.

L3 MC Redist Home Static ISID List Field Descriptions

Use the data in the following table to use the **L3 MC Redist Home Static ISID List** tab.

Name	Descriptions
Type	Specifies the layer 3 static I-SID routed multicast redistribution type. The default type is staticisid.
Enable	Enables layer 3 static I-SID routed multicast redistribution in the home area. The default value is disabled.
IsidListName	Specifies the name of the I-SID list. The default I-SID list name is htr.
Apply	Applies the Multi-area SPB layer 3 static I-SID routed multicast redistribution filters in the home area. The default value is noAction.

Configure Multi-area SPB Layer 3 Static I-SID Routed Multicast Redistribution Policy in Remote Area

About This Task

Perform this procedure to configure Multi-area SPB Layer 3 static I-SID routed multicast redistribution policy in the remote area.



Note

Routed multicast redistribution is available for IPv4 only. With the MAC learning disabled, the static data I-SID multicast redistribution policy takes precedence over the Layer 2 multicast redistribution.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**.
3. Select the **L3 MC Redist Remote Static ISID List** tab.
4. Select **Insert**.
5. In the **Type** field, select **staticisid**.
6. Select **Enable**.
7. In the **IsidListName** field, type the name of the I-SID list.
8. In the **Apply** field, select **apply**.
9. Select **Insert**.

L3 MC Redist Remote Static ISID List Field Descriptions

Use the data in the following table to use the **L3 MC Redist Remote Static ISID List** tab.

Name	Descriptions
Type	Specifies the layer 3 static I-SID routed multicast redistribution type. The default type is staticisid.
Enable	Enables layer 3 static I-SID routed multicast redistribution in the remote area. The default value is disabled.
IsidListName	Specifies the name of the I-SID list. The default I-SID list name is htr.
Apply	Applies the Multi-area SPB layer 3 static I-SID routed multicast redistribution filters in the remote area. The default value is noAction.

Configure Multi-area SPB DvR Redistribution



Note

This procedure does not apply to 5520 Series.

About This Task

Perform this procedure to configure the Multi-area SPB Distributed Virtual Routing (DvR) backbone redistribution on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **MultiArea**
3. Select the **DVR Redistribute** tab.
4. Select **DvrBbRedistribute**.
5. Select **Apply**.

DVR Redistribute Field Descriptions

Use data in the following table to use the **DVR Redistribute** tab.

Name	Description
DvrBbRedistribute	Configures the Multi-area SPB DvR redistribution on the switch. The default is disabled.

Display Remote IS-IS and SPBM Summary

About This Task

Perform this procedure to view the remote Intermediate-System-to-Intermediate-System (IS-IS) and SPBM summary.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.

2. Select **IS-IS Remote**.
3. Select the **Protocol Summary** tab.

Protocol Summary Field Descriptions

Use data in the following table to use the **Protocol Summary** tab.

Name	Description
AdminState	Indicates the global status of remote IS-IS on the switch.
SystemId	Indicates the remote IS-IS system ID of the switch.
NodeNickName	Indicates the nickname for the remote SPBM instance globally.
Circuit Index	Indicates the identifier of the IS-IS circuit. This value is for SNMP indexing purposes only and does not have any relation to any protocol value.
IfIndex	Indicates the interface to which the circuit corresponds.
AdminState	Indicates the administrative state of the circuit. <ul style="list-style-type: none"> • on • off
OperState	Indicates the operational state of the circuit. <ul style="list-style-type: none"> • up • down
Circuit Index	Indicates the unique identifier of the remote IS-IS circuit. This value is for SNMP Indexing purposes only and does not have any relation to any protocol value.
AdjIndex	Indicates value identifying the IS-IS adjacency from all other such adjacencies on the circuit. The system dynamically assigns this value when it forms an adjacency.
Interface	Indicates the remote IS-IS interface.
AdjState	Indicates the state of the adjacency: <ul style="list-style-type: none"> • down • initializing • up • failed

Name	Description
AdjNeighSysID	Indicates the system ID of the neighboring IS-IS.
AdjHostName	Indicates the host name listed in the LSP, or the IS-IS system name. The IS-IS system name uses global system name, if the host name is not configured.

Configure Remote IS-IS Global Parameters

About This Task

Perform this procedure to configure remote IS-IS global parameters. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.



Note

You cannot enable TTL handling for bridged traffic with routed SPB IP Multicast if remote IS-IS is already configured. Likewise, you cannot configure remote IS-IS if TTL handling is already configured.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Globals** tab.
4. Configure the remote IS-IS global parameters.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminState	Specifies the global status of remote IS-IS on the switch. The default is off.
SystemId	<p>Specifies the IS-IS system ID for the switch. Valid value is a 6-byte value in the format <xxxx.xxxx.xxxx>.</p> <p>Important: After you configure the SPBM nickname and enable remote IS-IS, if you change the system ID, you must also change the nickname. However, changing the nickname could impact the naming convention or configuration. To maintain the same nickname with a different system ID, see Job aid on page 1266.</p>

Name	Description
DynamicallyLearnedArea	For FAN members, specifies the remote IS-IS area that is dynamically learned from the neighbor's Hello PDU if the node does not have the IS-IS manual area configured.
MAOperState	Specifies the Multi-area SPB operational state: <ul style="list-style-type: none"> • Disabled – Multi-area SPB is disabled on the switch. • Init – the switch is initializing to function as a boundary node. • Designated – the switch is up and running as a designated boundary node. • Non-designated – the switch is up and running as a non-designated boundary node. The default is disabled.
AreaName	Configures the remote area name.
DesignatedBNHostName Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the hostname of the designated boundary node.
DesignatedBNHomeSystemId Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the system ID of the designated boundary node in the home area.
DesignatedBNRemoteSystemId Note: Exception: only applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.	Specifies the system ID of the designated boundary node in the remote area.

Configure Remote IS-IS Interfaces

About This Task

Perform this procedure to configure the remote IS-IS interfaces. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Interfaces** tab.
4. Select **Insert** to create a remote IS-IS instance.

5. Configure the IS-IS interface parameters.
6. Select **Insert**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Index	Specifies the identifier of the circuit, unique within the Intermediate System. This value is for SNMP Indexing purposes only and need not have any relation to any protocol value.
IfIndex	Specifies the interface on which the circuit is configured (port or MLT).
Type	Specifies the IS-IS circuit type. Only the point-to-point (PtToPt) interface type is supported.
AdminState	Specifies the administrative state of the circuit: on or off.
OperState	Specifies the operational state of the circuit.

Name	Description
AuthType	<p>Specifies the authentication type:</p> <ul style="list-style-type: none"> • none • simple: If selected, you must also specify a key value but the key id is optional. Simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5: If selected, you must also specify a key value, but the key-id is optional. MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. There is an optional key ID. • hmac-sha-256: If selected, you must also specify a key value, but the key-id is optional. With SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. If the digests match, the packet is accepted. If the digests do not match, the receiving switch discards the packet. There is an optional key ID. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default is none.</p>
AuthKey	Specifies the authentication key.
KeyId	Specifies the authentication key ID.

Name	Description
LevelType	Specifies the router type globally: <ul style="list-style-type: none"> level1: Level-1 router type level 1and2: Level-1/2 router type. This type is not supported. The default value is level1.
NumAdj	Specifies the number of adjacencies on this circuit.
NumUpAdj	Specifies the number of adjacencies that are up.
AutoNniEnable	Enable to have the node create an IS-IS interface, attach the interface to an SPBM instance, and then enable IS-IS on the port interface. This field displays on the Insert Interfaces dialog box and applies to port interfaces only.
Origin	Specifies the origin of the IS-IS circuit configuration on the port, either manually configured through CLI or EDM or dynamically configured through Auto-sense.
PortMembers	Displays the members of an MLT.

Configure Remote IS-IS Interface Level Parameters

About This Task

Perform this procedure to configure remote IS-IS interface level parameters. SPBM uses IS-IS to discover network topology, build shortest path trees between network nodes, and communicate network information in the control plane.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Interfaces Level** tab.
4. Configure the remote IS-IS interface level parameters.
5. Select **Apply**.

Interfaces Level Field Descriptions

Use the data in the following table to use the **Interfaces Level** tab.

Name	Description
Index	Specifies the identifier of the circuit, the value is unique within the remote IS-IS. The value is for SNMP Indexing purposes only and does not relate to any protocol value.
Level	Specifies the router type globally: <ul style="list-style-type: none"> • l1: Level1 router type • l12: Level1/Level2 router type. This type is not supported. The default value is l1.
ISPriority	Specifies an integer sub-range for remote IS-IS priority. The default is 64.
HelloTimer	Configures the level 1 hello interval. Specifies the maximum period, in milliseconds, between IS-IS Hello Packets (IIH) PDUs on multiaccess networks at this level for LANs. The value at Level1 is used as the period between Hellos on Level1/Level2 point to point circuits. Setting this value at Level 2 on an Level1/Level2 point-to-point circuit results in an error of InconsistentValue. The default value is 9000 milliseconds or 9 seconds.
HelloMultiplier	Configures the level 1 hello multiplier. The default value is 3 seconds.
DRHelloTimer	Indicates the period, in milliseconds, between Hello PDUs on multiaccess networks when this Intermediate System is the Designated Intermediate System. The default is 3.

Configure Remote SPBM on an Interface

About This Task

Perform this procedure to configure remote Shortest Path Bridging MAC (SPBM) on an interface.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Interfaces** tab.
4. Select the **SPBM** button.
5. Select **Insert**
6. Configure the interfaces SPBM parameters.
7. Select **Insert**.

SPBM Field Descriptions

Use data in the following table to use the **Interfaces SPBM** tab.

Name	Description
Index	Specifies an Index value for the remote SPBM interface.
SpbmId	Specifies the remote SPBM ID.
State	Specifies whether the remote SPBM interface is enabled or disabled.
Type	Configures the remote SPBM instance interface-type on the IS-IS interface on the specific port or MLT, the values are point-to-point (ptpt) or broadcast (bcast). The system supports the ptpt interface type only.
LlMetric	Configures the remote SPBM instance ll-metric on the IS-IS interface on the specific port or MLT. The default value is 10.
Origin	Specifies the source of the remote SPBM instance configuration. <ul style="list-style-type: none"> config - configure manually through CLI or EDM. dynamic - configure dynamically through Auto-sense.

View Statistical Graph of Remote IS-IS Interface Counters

About This Task

Perform this procedure to view statistical graph of the remote IS-IS interface counters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select **Graph**.

Interface Counters Field Descriptions

The following table describes the fields in the **Interface Counters** tab.

Name	Description
InitFails	Indicates the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures.
RejAdjs	Indicates the number of times an adjacency has been rejected on this circuit.

Name	Description
IDFieldLenMismatches	Indicates the number of times an Intermediate-System-to-Intermediate-System (IS-IS) control PDU with an ID field length different from that for this system has been received.
MaxAreaAddrMismatches	Indicates the number of times an IS-IS control PDU with a max area address field different from that for this system has been received.
AuthFails	Indicates the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Indicates the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

View Statistical Graph of Remote IS-IS Interface Sending Control Packets

About This Task

Perform this procedure to view statistical graph of the remote IS-IS interface sending control packets.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select the **Graph** button.
6. Select the **Interface Sending Control Packets** tab.

Interface Sending Control Packets Field Descriptions

Use data in the following table to use the **Interface Sending Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello (IIH) PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on Layer 1 or L1L2 circuits, and at Layer 2 otherwise.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

View Statistical Graph of Remote IS-IS Interface Receiving Control Packet

About This Task

Perform this procedure to view statistical graph of the remote IS-IS interface receiving control packets.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select **Graph**.
6. Select the **Interface Receiving Control Packets** tab.

Interface Receiving Control Packets Field Descriptions

Use data in the following table to use the **Interface Receiving Control Packets** tab.

Name	Description
Hello	Indicates the number of IS-IS Hello PDUs seen in this direction at this level. Point-to-Point IIH PDUs are counted at the lowest enabled level: at L1 on L1 or L1L2 circuits, and at L2 otherwise.
LSP	Indicates the number of IS-IS link-state packet (LSP) frames seen in this direction at this level.

Name	Description
CSNP	Indicates the number of IS-IS Complete Sequence Number Packet (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNPs) seen in this direction at this level.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.

Configure Remote IS-IS Manual Area

About This Task

Perform this procedure to configure a remote IS-IS manual area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Manual Area** tab.
4. Select **Insert**.
5. In the **AreaAddr** field, type the Area Address.
6. Select **Insert**.

Manual Area Field Descriptions

Use the data in the following table to use the **Manual Area** tab.

Name	Description
AreaAddr	Specifies the remote IS-IS manual-area in the range of 1 to 13 bytes.

View Level 1 Remote Area Information

About This Task

Perform this procedure to display Level 1 remote area information. Remote Intermediate-System-to-Intermediate-System (IS-IS) provides support for hierarchical routing, which enables the system to partition large routing domains into smaller areas. Remote IS-IS uses a two-level hierarchy, dividing the domain into multiple Level 1 areas

and one Level 2 area. The Level 2 area serves as backbone of the domain, connecting to all the Level 1 areas.

**Important**

The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **L1 Area** tab.

L1 Area Remote Field Descriptions

Use the data in the following table to use the **L1 Area Remote** tab.

Name	Description
AreaAddr	Specifies an area address reported in a Level 1 link-state packets (LSP) that the system generates or receives through the Intermediate System.

*View Remote LSP Summary Information***About This Task**

Perform this procedure to view the remote link-state packet (LSP) summary information. Link State Packets (LSP) contain information about the state of adjacencies or defined and distributed static routes. Intermediate-System-to-Intermediate-System (IS-IS) exchanges this information with neighboring IS-IS routers at periodic intervals.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **LSP Summary** tab.

LSP Summary Field Descriptions

Use the data in the following table to use the **LSP Summary** tab.

Name	Description
Level	Specifies the level at which the system displays the remote LSP.
ID	Specifies the 8-byte LSP ID, consisting of the system ID, circuit ID, and Fragment Number.
Seq	Specifies the sequence number for the LSP.
Checksum	Specifies the 16-bit Fletcher Checksum for the LSP.

Name	Description
LifetimeRemain	Specifies the remaining lifetime in seconds for the LSP.
HostName	Specifies the host name.

View Remote IS-IS Adjacencies

About This Task

Perform this procedure to view remote IS-IS adjacency information. The platform sends IS-IS Hello (IIH) packets to discover IS-IS neighbors and establish and maintain the IS-IS adjacency. The platform continues to send IIH packets to maintain the adjacencies that the system establishes. For two nodes to form an adjacency the B-VLAN pairs for the primary B-VLAN and secondary B-VLAN must match.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Adjacency** tab.

Adjacency Field Descriptions

Use the data in the following table to use the **Adjacency** tab.

Name	Description
Interface	Specifies the remote IS-IS interface on which the adjacency is found.
Level	Indicates the level of the remote IS-IS interface (Level 1 [default] or Level 2).
State	Specifies the state of the adjacency: <ul style="list-style-type: none"> • down • initializing • up • failed
LastUpTime	Indicates when the adjacency most recently entered the state up , measured in hundredths of a second since the last re-initialization of the network management subsystem. Displays 0 if the adjacency has never been in state up .
NeighPriority	Specifies the priority of the neighboring Intermediate System to transition to the Designated Intermediate System.
HoldTimer	Specifies the holding time in seconds for this adjacency. This value is based on received IS-IS Hello (IIH) PDUs and the elapsed time since receipt.
NeighSysID	Specifies the system ID of the neighboring Intermediate System.

Name	Description
AdjHostName	Specifies the host name listed in the LSP.
ParallelActive	Specifies if the current adjacency among all the parallel adjacencies between two nodes is active. The values are: <ul style="list-style-type: none"> • true • false
PortMembers	Displays the members of an MLT.

Configure IS-IS Area Virtual Node

About This Task



Note

As a best practice, use the derived values for the virtual node to avoid misconfiguration between peer boundary nodes. However, you can use this procedure to override the derived values.

The following list describes how the switch derives the virtual node parameters:

- The switch derives the virtual node nickname from the last 5 nibbles of the manual area address. For example, area ID 49.0001 generates a nickname of 9.00.01.
- The switch derives the system ID and system name from the nickname. For example, a nickname of 9.00.01 generates system ID 92:00:01:ff:ff:f0.

Perform this procedure to configure the area virtual node parameters like system name, system ID, and nickname on the Intermediate-System-to-Intermediate-System (IS-IS) instance.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. Select the **Area VNode** tab.
4. Configure the parameters.
5. Select **Apply**.

Area VNode Field Descriptions

Use the data in the following table to use the **Area VNode** tab.

Name	Description
VNodeNickName	Specifies the nickname of the virtual node on the IS-IS instance.
VNodeSysId	Specifies the system ID of the virtual node on the IS-IS instance.
VNodeSysName	Specifies the system name of the virtual node on the IS-IS instance.

Configure Remote IS-IS Area Virtual Node

About This Task



Note

As a best practice, use the derived values for the virtual node to avoid misconfiguration between peer boundary nodes. However, you can use this procedure to override the derived values.

The following list describes how the switch derives the virtual node parameters:

- The switch derives the virtual node nickname from the last 5 nibbles of the manual area address. For example, area ID 49.0001 generates a nickname of 9.00.01.
- The switch derives the system ID and system name from the nickname. For example, a nickname of 9.00.01 generates system ID 92:00:01:ff:ff:f0.

Perform this procedure to configure the area virtual node parameters like system name, system ID, and nickname on the remote Intermediate-to-Intermediate-System (IS-IS) instance.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS Remote**.
3. Select the **Area VNode** tab.
4. Configure the parameters.
5. Select **Apply**.

Area VNode Field Descriptions

Use the data in the following table to use the **Area VNode** tab.

Name	Description
VNodeNickName	Specifies the nickname of the virtual node on remote IS-IS instance.
VNodeSysId	Specifies the system ID of the virtual node on remote IS-IS instance.
VNodeSysName	Specifies the system name of the virtual node on remote IS-IS instance.

Configure Remote SPBM Parameters

About This Task

Perform this procedure to configure remote SPBM global parameters. SPBM uses the remote Intermediate-System-to-Intermediate-System (IS-IS) link state routing protocol to provide a loop free Ethernet topology that creates a shortest path topology from every node to every other node in the network based on node MAC addresses.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **SPBM** tab.

4. Configure the remote SPBM parameters.
5. Select **Apply**.

SPBM Field Descriptions

Use data in the following table to use the **SPBM** tab.

Name	Description
Id	Displays the remote SPBM instance ID.
NodeNickName	Specifies a nickname for the remote SPBM instance globally. The value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Displays the primary SPBM B-VLANs to add to the remote SPBM instance.
Vlans	Displays the SPBM B-VLANs to add to the remote SPBM instance.

Configure Remote SPBM Parameters on the Interface

About This Task

Perform this procedure to configure the remote SPBM interface parameters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **Interfaces SPBM** tab.
4. Select **Insert** to create a remote SPBM instance on the interface.
5. Configure the remote SPBM interface parameters.
6. Select **Insert**.

Interfaces SPBM Field Descriptions

Use the data in the following table to use the **Interfaces SPBM** tab.

Name	Description
Index	Specifies an Index value for the remote SPBM interface.
SpbmId	Specifies the remote SPBM ID.
State	Specifies whether the remote SPBM interface is enabled or disabled.
Type	Configures the remote SPBM instance interface-type on the IS-IS interface on the specific port or MLT, the values are point-to-point (ptpt) or broadcast (bcast). The system supports the ptpt interface type only.

Name	Description
l1Metric	Configures the remote SPBM instance l1-metric on the IS-IS interface on the specific port or MLT. The default value is 10.
Origin	Specifies the source of the remote SPBM instance configuration. <ul style="list-style-type: none"> • config - configure manually through CLI or EDM. • dynamic - configure dynamically through Auto-sense.

Create a Virtual NNI Link Between Multi-Area Boundary Nodes

About This Task

Perform this procedure to create a virtual NNI link between Multi-area boundary nodes.

The assumption is that you are creating the CLIP and the logical interface for the first time. You can also use the table-based tab to apply the configuration to an existing interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select **Insert**.
5. In the Interface field, add a CLIP interface number.
6. Type the IP address.
7. Type the network mask.
8. Select the type of loopback interface for the virtual NNI link:
 - **circuitlessIPMAVirtualLinkHome**
 - **circuitlessIPMAVirtualLinkRemote**



Note

If you configure the loopback interface for the virtual NNI link in the home area, you must configure IS-IS on the logical interface in the remote area, and vice versa.

9. (Optional) For **Name**, type a name for this interface.
10. Select **Insert**.
11. In the navigation pane, expand **Configuration > Fabric**.
12. Select **IS-IS**.
13. Select **Logical Interfaces** tab.
14. Select **Insert**.
15. For **Id**, type the index number that uniquely identifies this logical interface.

16. (Optional) For **Name**, type the name of this logical interface.
17. For **Type**, select **ip**.
18. For **DestIPAddr**, type the destination IP address for the logical interface.
19. Select **Multi-area Virtual Link** to configure a Multi-area virtual link on this interface.

**Note**

If you configure the loopback interface for the virtual NNI link in the home area, you must configure IS-IS on the logical interface in the remote area, and vice versa.

20. Select **Insert**.

What to Do Next

Configure IS-IS and SPBM on the logical interface.

Circuitless IP Field Descriptions

Use the data in the following table to use the **Circuitless IP** tab.

Name	Description
Interface	Specifies the number assigned to the interface. You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.
Ip Address	Specifies the IP address of the CLIP. You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.
Net Mask	Specifies the network mask. You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.
Name	Specifies a name assigned to the IPv4 CLIP address. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 5720 Series, 7520 Series, and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link</code> .

Name	Description
IfType	<p>Specifies the interface type.</p> <p>Note: Exception: circuitlessIPMAVirtualLinkHome and circuitlessIPMAVirtualLinkRemote apply to 5720 Series, 7520 Series, and 7720 Series only.</p> <p>You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.</p>
AccessType Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	<p>Specifies to which segment of a Segmented VRF the interface belongs. If you use the <i>trusted</i> parameter, the interface is part of the trusted segment and is only accessible from hosts in the trusted and unrestricted segments. Use the <i>unrestricted</i> parameter to make the interface accessible from hosts in all three segments. The default is trusted.</p> <p>EDM only displays this field if you launch the VRF context view. This field does not apply to the GlobalRouter.</p> <p>You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.</p>

Logical Interfaces Field Descriptions

Use the data in the following table to use the **Logical Interfaces** tab and the **Insert Logical Interfaces** dialog. The available fields in the dialog differ depending on the type of core you select: **layer 2** or **ip**.

Name	Description
Id	<p>Specifies the index number that uniquely identifies this logical interface.</p> <p>This field displays on the Insert Logical Interfaces dialog only.</p>
IfIndex	<p>Specifies the index number that uniquely identifies this logical interface. This field is read-only.</p> <p>This field displays on the Logical Interfaces tab only.</p>
Name	<p>Specifies a name associated with this logical interface, which can be up to 64 characters.</p> <p>Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as <i>Virtual-link-x</i>, where <i>x</i> is the logical interface ID.</p>

Name	Description
Type	Specifies the type of logical interface to create: <ul style="list-style-type: none"> Specify layer 2 for a Layer 2 core network that the tunnel will traverse. Specify ip for a Layer 3 core network that the tunnel will traverse.
DestIPAddr	Specifies the destination IP address for the IP-type logical interface.
DestIfIndex	Specifies the physical port or MultiLink Trunking (MLT) that the Layer 2 logical interface is connected to.
Vids	Specifies the list of VLANs that are associated with this logical interface.
PrimaryVid	Specifies the primary tunnel VLAN ID associated with this Layer 2 Intermediate-System-to-Intermediate-System (IS-IS) logical interface.
CircIndex	Identifies the IS-IS circuit created under the logical interface. This field displays on the Logical Interfaces tab only.
NextHopVrf	Displays the next-hop VRF name to reach the logical tunnel destination IP. This field displays on the Logical Interfaces tab only. You can use this field to specify the VRF to reach the logical tunnel destination IP associated with a parallel tunnel.
ISIS Mtu	Specifies the Maximum Transmission Unit (MTU) size in bytes for IS-IS packets that use this logical interface. The default value is 1600.
BfdEnable	Enables or disables BFD on an IS-IS Logical Interface.
SrcIPAddr	Configures an additional source address to use as the parallel tunnel to create a backup adjacency. Note: To use an IPsec-encrypted tunnel as the parallel tunnel ensure that you configure the same source IP address on the logical IS-IS interface and in the Fabric IPsec Gateway virtual machine.
Origin	Specifies the origin of the IS-IS logical interface configuration, either through Zero Touch Fabric Configuration (ZTF) or manual configuration (config) through CLI or EDM.
Multi-Area Virtual Link Note: Exception: only applies to 7520 Series and 7720 Series.	Enables (true) or disables (false) a Multi-area boundary node virtual NNI link on an IS-IS Logical Interface. The default is disabled.

View Remote SPBM I-SID Information

About This Task

Perform this procedure to view the SPBM Service Instance Identifier (I-SID) information. The SPBM BMAC header includes an I-SID with a length of 24 bits. This I-SID can be used to identify and transmit any virtualized traffic in an encapsulated SPBM frame.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **I-SID** tab.
4. (Optional) Select the **Filter** button to filter the rows on basis of specific criteria.

I-SID Field Descriptions

Use data in the following table to use the **I-SID** tab.

Name	Description
SysId	Indicates the system identifier.
Vlan	Indicates the B-VLAN where the remote I-SID is configured or discovered.
Isid	Indicates the remote IS-IS SPBM I-SID identifier.
NickName	Indicates the nickname of the node where the remote I-SID is configured or discovered.
HostName	Indicates the host name listed in the LSP, or the IS-IS system name. The IS-IS system name uses global system name, if you do not configure the host name.
Type	Indicates the remote SPBM I-SID type. The values are discover or configure.

View Remote SPBM Nicknames

About This Task

Perform this procedure to view remote SPBM nicknames on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **Nick Names** tab.

Nick Names Field Descriptions

Use the data in the following table to use the **Nick Names** tab.

Name	Description
Level	Indicates the level at which the system displays this LSP.
ID	Indicates the 8 byte LSP ID, consisting of the System ID, Circuit ID, and Fragment Number.
LifetimeRemain	Indicates the remaining lifetime in seconds for the LSP.
NickName	Indicates the nickname for the remote SPBM node.
HostName	Indicates the hostname that the system displays in the LSP, or the system name if the host name is not configured.

View Remote IP Unicast FIB

In SPBM, each node has a System ID, which also serves as Backbone MAC address (BMAC) of the switch. The system populates these Backbone MAC addresses into the SPBM VLAN Forwarding Information Base (FIB). When the system discovers a network topology and stores it in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information that it receives from SPBM into the FIB for forwarding purposes.

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If you only enable IP Shortcuts on the Backbone Edge Bridges, I-SIDs are never exchanged in the network as IP Shortcuts allows for Global Routing Table (GRT) IP networks to be transported across IS-IS.

The system adds the routes within the same VSN to the LSDB with a default preference of 7 and the inter-VSN routes with a route preference of 200. With IS-IS accept policies you can change the route preference for incoming routes. If the system learns the same route from multiple sources with different route preferences, then the routes are not equal cost multipath (ECMP) routes. The system prefers the route with the lowest route preference. In Layer 2, if there is a tie-break between routes from multiple sources, the tie-breaking is based on cost and hop count.

About This Task

Perform this procedure to display the remote IP unicast FIB. You can view the IP routes from remote Backbone Edge Bridges (BEBs)

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **IP Unicast FIB** tab.
4. (Optional) Select the **Filter** button to filter the rows on basis of specific criteria.

IP Unicast FIB Field Descriptions

Use the data in the following table to use the **IP Unicast FIB** tab.

Name	Description
VrfId	Specifies the VRF ID of the remote IP unicast FIB entry, 0 indicates NRE.
DestinationIpAddrType	Specifies the address type of the destination IP address.
DestinationIpAddr	Specifies the destination IP Address of the remote IP unicast FIB entry.
DestinationMask	Specifies the destination IP mask of the remote IP unicast FIB entry
NextHopBmac	Specifies the nexthop BMAC of the remote IP unicast FIB entry.
DestIsid	Specifies the destination I-SID of the remote IP unicast FIB entry.
Vlan	Specifies the VLAN of the remote IP unicast FIB entry.
Isid	Specifies the I-SID of the remote IP unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the remote IP unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the remote IP unicast FIB entry.
PrefixCost	Specifies the prefix cost of the remote IP unicast FIB entry.
SpbmCost	Specifies the BMAC cost of the remote IP unicast FIB entry.
Preference	Specifies the IP Route preference of the remote IP unicast FIB entry
MetricType	Specifies the IP Metric Type of the remote IP unicast FIB entry.

View Remote IPv6 Unicast FIB

In SPBM, each node has a System ID, which also serves as Backbone MAC address (BMAC) of the switch. The system populates these Backbone MAC addresses into the SPBM VLAN Forwarding Information Base (FIB). When the system discovers a network topology and stores it in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information that it receives from SPBM into the FIB for forwarding purposes.

I-SIDs are only used for virtual services (Layer 2 VSNs and Layer 3 VSNs). If you only enable IP Shortcuts on the Backbone Edge Bridges, I-SIDs are not exchanged in the network as IP Shortcuts allows for Global Routing Table (GRT) IP networks to be transported across IS-IS.

About This Task

Perform this procedure to display the remote IPv6 unicast FIB. You can view the IPv6 routes from remote Backbone Edge Bridges (BEBs).

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **IPv6 Unicast FIB** tab.
4. (Optional) Select the **Filter** button to filter the rows on basis of specific criteria.

IPv6 Unicast FIB Field Descriptions

Use the data in the following table to use the **IPv6 Unicast FIB** tab.

Name	Description
Vrflid	Specifies the VRF ID of the remote IPv6 unicast FIB entry, 0 indicates NRE.
DestinationIpAddrType	Specifies the address type of the destination IPv6 address.
DestinationIpAddr	Specifies the destination IPv6 Address of the remote IPv6 unicast FIB entry.
DestinationMask	Specifies the destination IPv6 mask of the remote IPv6 unicast FIB entry
NextHopBmac	Specifies the nexthop BMAC of the remote IPv6 unicast FIB entry.
DestIsid	Specifies the destination I-SID of the remote IPv6 unicast FIB entry.
Vlan	Specifies the VLAN of the remote IPv6 unicast FIB entry.
Isid	Specifies the I-SID of the IPv6 unicast FIB entry.
NextHopName	Specifies the nexthop hostname of the remote IPv6 unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the remote IPv6 unicast FIB entry.
PrefixCost	Specifies the prefix cost of the remote IPv6 unicast FIB entry.
SpbmCost	Specifies the BMAC cost of the remote IPv6 unicast FIB entry.
MetricType	Specifies the Metric Type of the remote IPv6 unicast FIB entry.

View Remote Unicast FIB

In SPBM, the IS-IS link-state database carries the BMAC addresses. To do this, SPBM supports an IS-IS TLV that advertises the I-SID and BMAC information across the network. Each node has a System ID, which also serves as Backbone MAC address

(BMAC) of the switch. These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB).

When the system discovers a network topology and stores it in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information that it receives from SPBM into the FIB for forwarding purposes.

About This Task

Perform this procedure to view the remote unicast FIB.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **Unicast FIB** tab.

Unicast FIB Remote Field Descriptions

Use the data in the following table to use the **Unicast FIB Remote** tab.

Name	Description
SysId	Specifies the system ID of the node where the remote unicast FIB entry originates.
Vlan	Specifies the VLAN of the remote unicast FIB entry.
DestinationMacAddr	Specifies the destination MAC Address of the remote unicast FIB entry.
OutgoingPort	Specifies the outgoing port of the remote unicast FIB entry.
HostName	Specifies the host name of the node where the remote unicast FIB entry originates.
Cost	Specifies the cost of the remote unicast FIB entry.

View Remote Multicast FIB

In SPBM, the IS-IS link-state database carries the BMAC addresses. To do this, SPBM supports an IS-IS TLV that advertises the I-SID and BMAC information across the network. Each node has a System ID, which also serves as Backbone MAC address (BMAC) of the switch. The system populates these Backbone MAC addresses into the SPBM VLAN Forwarding Information Base (FIB).

When the system discovers a network topology and stores it in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information that it receives from SPBM into the FIB for forwarding purposes.

The system produces the remote multicast FIB only after you configure the virtual services and the system learns them.

About This Task

Perform this procedure to view the remote multicast FIB.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **Multicast FIB** tab.

Multicast FIB Field Descriptions

Use the data in the following table to use the **Multicast FIB** tab.

Name	Description
SysId	Specifies the system ID of the node where the remote multicast FIB entry originates.
Vlan	Specifies the VLAN ID of the remote multicast FIB entry.
McastDestMacAddr	Specifies the multicast destination MAC Address of the remote multicast FIB entry.
Isid	Specifies the I-SID value in the remote multicast FIB entry.
Isid Name	Specifies the name that the system assigns to the I-SID.
HostName	Specifies the host name of the node where the multicast FIB entry originates.
OutgoingInterfaces	Specifies the switched UNI port outgoing interface of the remote multicast FIB entry.
IncomingInterface	Specifies the incoming interface (port or MLT) of the remote multicast FIB entry.

View Remote IP Multicast over Fabric Connect Routes

About This Task

Perform this procedure to view the remote IP Multicast over Fabric Connect routes.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **IpMcastRoutes** tab.

IpMcastRoutes Field Descriptions

Use the data in the following table to use the **IpMcastRoutes** tab.

Name	Description
VsnIsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Source	Specifies the IP address where the IP Multicast over Fabric Connect route originated.
SourceBeb	Specifies the source BEB for the IP multicast route.
VlanId	Specifies the ID for the C-VLAN.
VrfName	Specifies the VRF name.
Datalsid	Specifies the data I-SID for the IP Multicast over Fabric Connect route. A a BEB receives IP multicast data from a sender, a BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
Type	Specifies the type for the IP Multicast over Fabric Connect route.
Bvlan	Specifies the B-VLAN for the IP Multicast over Fabric Connect route.
NniInterfaces	Specifies the network-to-network interface (NNI) ports for the IP multicast route. SPBM runs in the core on the ports that connect to the core. These ports are NNI ports. Ports that face a customer VLAN are user-to-network interface (UNI) ports.

View UNI Ports for remote IP Multicast Routes

About This Task

Perform this procedure to view the UNI ports associated with particular remote IP multicast routes.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM Remote**.
3. Select the **IpMcastRoutes** tab.
4. Select an entry.
5. Select the **UNI Ports** button.

IpMcastRoutes-Uni Ports Field Descriptions

Use data in the following table to use the **IpMcastRoutes-Uni Ports** tab.

Name	Description
VsnIsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Source	Specifies the IP address where the IP Multicast over Fabric Connect route originates.
NickName	Specifies the nickname.
EgressVlan	Specifies the egress VLAN of the IP multicast route.
EgressActivePorts	Specifies the egress active ports.

View Remote IS-IS System Statistics

About This Task

Perform this procedure to view the remote Intermediate-System-to-Intermediate-System (IS-IS) system statistics.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Stats Remote**.
3. Select the **System Stats** tab.

System Stats Field Descriptions

Use the data in the following table to use the **System Stats** tab.

Name	Description
CorrLSPs	Indicates the number of corrupted in-memory link-state packets (LSPs) detected. LSPs received from the wire with a bad checksum are silently dropped and not counted.
AuthFails	Indicates the number of authentication key failures recognized by this Intermediate System.
LSPDbaseOloads	Indicates the number of times the LSP database has become overloaded.
ManAddrDropFromAreas	Indicates the number of times a manual address has been dropped from the area.
AttmptToExMaxSeqNums	Indicates the number of times the IS has attempted to exceed the maximum sequence number.

Name	Description
SeqNumSkips	Indicates the number of times a sequence number skip has occurred.
OwnLSPPurges	Indicates the number of times a zero-aged copy of the system's own LSP is received from some other node.
IDFieldLenMismatches	Indicates the number of times a PDU is received with a different value for ID field length to that of the receiving system.
PartChanges	Indicates partition changes.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.

View Remote IS-IS Interface Counters

About This Task

Perform this procedure to view the remote Intermediate-System-Intermediate-System (IS-IS) interface counters.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Stats Remote**.
3. Select the **Interface Counters** tab.

Interface Counters Field Descriptions

Use the data in the following table to use the **Interface Counters** tab.

Name	Description
Index	Shows a unique value identifying the IS-IS interface.
Level	Shows the type of circuit that discovered the interface counters. The point to point Hello PDU includes both Layer 1 and Layer 2, and IS from a single adjacency on point to point links, therefore combining counts on point to point links into one group.
AdjChanges	Shows the number of times an adjacency state change has occurred on this circuit.

Name	Description
InitFails	Shows the number of times initialization of this circuit has failed. This counts events such as PPP NCP failures. Failures to form an adjacency are counted by isisCircRejAdjs.
RejAdjs	Shows the number of times an adjacency has been rejected on this circuit.
IDFieldLenMismatches	Shows the number of times an IS-IS control PDU with an ID field length different to that for this system has been received.
MaxAreaAddrMismatches	Shows the number of times an IS-IS control PDU with a max area address field different to that for this system has been received.
AuthFails	Shows the number of times an IS-IS control PDU with the correct auth type has failed to pass authentication validation.
LANDesISChanges	Shows the number of times the Designated IS has changed on this circuit at this level. If the circuit is point to point, this count is zero.

View Remote IS-IS Interface Control Packets

About This Task

Perform this procedure to view the remote Intermediate-System-to-Intermediate-System (IS-IS) interface control packets.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **Stats Remote**.
3. Select the **Interface Control Packets** tab.

Interface Control Packets Field Descriptions

Use the data in the following table to use the **Interface Control Packets** tab.

Name	Description
Index	Shows a unique value identifying the Intermediate-System-to-Intermediate-System (IS-IS) interface.
Level	Specifies the level.
Direction	Indicates whether the switch is sending or receiving the PDUs.
Hello	Indicates the number of IS-IS Hello frames seen in this direction at this level.
LSP	Indicates the number of IS-IS LSP frames seen in this direction at this level.

Name	Description
CSNP	Indicates the number of IS-IS Complete Sequence Number Packets (CSNP) frames seen in this direction at this level.
PSNP	Indicates the number of IS-IS Partial Sequence Number Packets (PSNP) frames seen in this direction at this level.

View the Remote I-SID Forwarding Database

About This Task

Perform this procedure to view the remote I-SID forwarding database (FDB) on the switch.



Note

To view the T-UNI I-SID FDB entries filtered on a port that is part of an MLT, you must mention the MLT ID in the option for the port.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID Remote**.
3. Select the **FDB** tab.
4. (Optional) Select the **Filter** button to filter the rows on basis of specific criteria.

FDB Field Descriptions

Use data in the following table to use the **FDB** tab.

Name	Description
IsidId	Specifies the service interface identifier (I-SID).
Address	Specifies the MAC address of the port assigned to the specific I-SID or CMAC learned on the particular I-SID.
Status	Specifies the learning status of the associated MAC.
Port	Specifies the port on which the MAC is learned for the specific I-SID.
PortType	Specifies if the MAC address is learned locally or on a network-to-network interface (NNI) port from a remote destination.
RemoteMacDestAddr	Specifies the virtual B-MAC address or system-ID of the remote destination.
RemoteMacBVlanId	Specifies the B-VLAN ID on which the remote destination was discovered.
RemoteMacDestSysName	Specifies the remote destination system name.
Cvid	Specifies the customer VLAN ID of the associated Switched UNI port.

Multi-area SPB Configuration Example

The following figure is an example of the devices connecting to each other in the Multi-area SPB network.



Note
The example applies to 5520 Series, 5720 Series, 7520 Series, and 7720 Series.




Figure 89: Multi-area SPB network example

The following sections display an example to configure two boundary nodes in the Multi-area SPB network (BG1 and BG2), to enable remote Intermediate-System-to-Intermediate-System (IS-IS), and to configure the Layer 2 and Layer 3 unicast redistribution.



Note
You must first configure basic SPBM and IS-IS infrastructure. For more information, see [SPBM Configuration Examples](#) on page 1293.

Boundary Node BG1 Configuration

```
enable
configure terminal
router isis remote
manual-area 49.0050
spbm 1 nick-name 1.21.15
exit
router isis remote enable
show isis area-vnode
=====
ISIS SPBM Multi-Area VNode Info
=====
VNODE      VNODE      VNODE      REPRESENTED  REPRESENTED  VNODE
SYSTEM-ID  NICK-NAME  HOST-NAME  AREA-ADDRESS AREA          USED-IN-AREA
-----
9200.30ff.fff0  9.00.30    vn-30      49.0030      HOME         REMOTE
9200.50ff.fff0  9.00.50    vn-50      49.0050      REMOTE       HOME
-----

show isis remote
=====
ISIS Remote Area Info
=====
AdminState : enabled
System ID : d887.66f6.e486
Num of Interfaces : 0
Num of Area Addresses : 1
Dynamically Learned Area :
Multi-Area OperState : disabled
```

```

interface gigabitEthernet 1/4,1/9
isis remote
isis remote spbm 1
isis remote enable
exit
show isis adjacencies
=====
ISIS Adjacencies
=====

```

INTERFACE	L	STATE	UPTIME	PRI	HOLDTIME	SYSID	HOST-NAME	STATUS	AREA	AREA-NAME
Port1/3	1	UP	20:33:27	127	19	b0ad.aa4d.b884	2118-50	ACTIVE	HOME	area-9.00.30
PortVirtual	1	UP	20:33:27	127	0	9200.50ff.fff0	area-9.00.50	ACTIVE	HOME	area-9.00.30
Port1/9	1	UP	00:01:16	127	27	d884.66f6.e484	BN2116	ACTIVE	HOME	area-9.00.30
Port1/4	1	UP	20:33:27	127	20	b0ad.aa40.1484	2119-50	ACTIVE	REMOTE	area-9.00.50
PortVirtual	1	UP	20:33:27	127	0	9200.30ff.fff0	area-9.00.30	ACTIVE	REMOTE	area-9.00.50
Port1/9	1	UP	00:01:00	127	21	d884.66f6.e486	BN2116	ACTIVE	REMOTE	area-9.00.50

```

Home: 3 out of 3 interfaces have formed an adjacency
Remote: 3 out of 3 interfaces have formed an adjacency
=====

router isis
multi-area l2 redistribute i-sid permit-all
exit
isis multi-area l2 apply redistribute i-sid
show isis multi-area l2 redistribute i-sid
=====
MULTI AREA L2 ISID REDIST POLICY
=====
Permit Except List Name
-----
permit-all -
-----

router isis
multi-area ip redistribute unicast
exit
isis multi-area ip apply redistribute unicast
show isis multi-area ip redistribute unicast
=====
ISIS Multiarea Redistribute List for ip unicast - GlobalRouter
=====
DIRECTION ENABLE RPOLICY
-----
home-to-remote TRUE
remote-to-home TRUE
-----

```

Boundary Node BG2 Configuration

```

enable
configure terminal
router isis remote
manual-area 49.0050
spbm 1 nick-name 1.21.16
exit
router isis remote enable
show isis area-vnode
=====
ISIS SPBM Multi-Area VNode Info
=====

```

VNODE SYSTEM-ID	VNODE NICK-NAME	VNODE HOST-NAME	REPRESENTED AREA-ADDRESS	REPRESENTED AREA	VNODE USED-IN-AREA
9200.30ff.fff0	9.00.30	vn-30	49.0030	HOME	REMOTE
9200.50ff.fff0	9.00.50	vn-50	49.0050	REMOTE	HOME

```

show isis remote
=====
ISIS Remote Area Info
=====
AdminState : enabled
System ID : d884.66f6.e486
Num of Interfaces : 0
Num of Area Addresses : 1
Dynamically Learned Area :
Multi-Area OperState : disabled

interface gigabitEthernet 1/4,1/9
isis remote
isis remote spbm 1

```

```
isis remote enable
exit
show isis adjacencies
=====
                        ISIS Adjacencies
=====
INTERFACE          L STATE   UPTIME    PRI  HOLDDTIME  SYSID          HOST-NAME      STATUS  AREA   AREA-NAME
-----
Port1/3             1 UP      20:33:27 127   19          b0ad.aa4d.b884  2118-50        ACTIVE  HOME   area-9.00.30
PortVirtual         1 UP      20:33:27 127    0          9200.50ff.fff0  area-9.00.50   ACTIVE  HOME   area-9.00.30
Port1/9             1 UP      00:01:16 127   27          d887.66f6.e484  BN2117         ACTIVE  HOME   area-9.00.30
Port1/4             1 UP      20:33:27 127   20          b0ad.aa40.1484  2119-50        ACTIVE  REMOTE  area-9.00.50
PortVirtual         1 UP      20:33:27 127    0          9200.30ff.fff0  area-9.00.30   ACTIVE  REMOTE  area-9.00.50
Port1/9             1 UP      00:01:00 127   21          d887.66f6.e486  BN2117         ACTIVE  REMOTE  area-9.00.50
-----
Home:    3 out of 3 interfaces have formed an adjacency
Remote:  3 out of 3 interfaces have formed an adjacency
-----

router isis
multi-area l2 redistribute i-sid permit-all
exit
isis multi-area l2 apply redistribute i-sid
show isis multi-area l2 redistribute i-sid
=====
                        MULTI AREA L2 ISID REDIST POLICY
=====
Permit          Except List Name
-----
permit-all     -
-----

multi-area ip redistribute unicast
exit
isis multi-area ip apply redistribute unicast
show isis multi-area ip redistribute unicast
=====
                        ISIS Multiarea Redistribute List for ip unicast - GlobalRouter
=====
DIRECTION      ENABLE    RPOLICY
-----
home-to-remote  TRUE
remote-to-home  TRUE
-----
```

Layer 2 VSN Configuration

Table 140: Layer 2 VSN product support

Feature	Product	Release introduced
Equal Cost Trees (ECT)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 140: Layer 2 VSN product support (continued)

Feature	Product	Release introduced
Layer 2 Virtual Service Network (VSN)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Switched UNI	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Transparent Port UNI (T-UNI)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Layer 2 VSN Configuration Fundamentals

This section provides fundamentals concepts for Layer 2 VSN.

SPBM Layer 2 VSN

SPBM supports Layer 2 VSN functionality where customer VLANs (C-VLANs) are bridged over the SPBM core infrastructure.

At the Backbone Edge Bridges (BEBs), customer VLANs (C-VLAN) are mapped to I-SIDs based on the local service provisioning. Outgoing frames are encapsulated in a MAC-in-MAC header, and then forwarded across the core to the far-end BEB, which strips off the encapsulation and forwards the frame to the destination network based on the I-SID-to-C-VLAN provisioning.

In the backbone VLAN (B-VLAN), Backbone Core Bridges (BCBs) forward the encapsulated traffic based on the B-MAC-DA, using the shortest path topology learned using IS-IS.

The following figure shows a sample campus SPBM Layer 2 VSN network.




Figure 90: SPBM L2 VSN in a campus

One of the key advantages of the SPBM Layer 2 VSN is that network virtualization provisioning is achieved by configuring only the edge of the network (BEBs). As a result, the intrusive core provisioning that other Layer 2 virtualization technologies require is not needed when new connectivity services are added to the SPBM network. For example, when new virtual server instances are created and need their own VLAN instances, they are provisioned at the network edge only and do not need to be configured throughout the rest of the network infrastructure.

Based on its I-SID scalability, this solution can scale much higher than any 802.1Q tagging based solution. Also, due to the fact that there is no need for Spanning Tree in the core, this solution does not need any core link provisioning for normal operation.

Redundant connectivity between the C-VLAN domain and the SPBM infrastructure can be achieved by operating two SPBM switches in switch clustering (SMLT) mode. This

allows the dual homing of any traditional link aggregation capable device into an SPBM network.

Fabric Connect Service Types

The Fabric Connect technology delivers Layer 2 and Layer 3 virtualization. These virtualized Layer 2 and Layer 3 instances are referred to as Virtual Service Networks (VSNs). A Service Identifier (I-SID) is used to uniquely distinguish these service instances network-wide, and a User Network Interface (UNI) is the boundary or demarcation point between the “service layer” of traditional networks, that is VLANs and VRFs, and the Fabric Connect “service layer”, that is Layer 2 & Layer 3 VSNs.

- Layer 2 VSNs are virtual broadcast domains interconnecting UNI members that share the same Layer 2 VSN I-SID. MAC learning/aging is applied to all Layer 2 VSNs.
- Layer 3 VSNs are virtual routed Layer 3 networks (Layer 3 VPN) leveraging IS-IS as the routing protocol between VRFs that share the same Layer 3 VSN I-SID.

Fabric Connect uses the User-Network-Interface (UNI) to denote the capabilities and attributes of the service interfaces. Fabric connect devices support the following UNI types:

- VLAN UNI (C-VLAN) — a device-specific VLAN-ID maps to a Layer 2 VSN I-SID – all device physical ports that are associated with the VLAN are therefore associated with the UNI.
- Flex UNI — it has the following sub-types:
 - Switched UNI — a VLAN-ID and a given port (VID, port) maps to a Layer 2 VSN I-SID. With this UNI type, VLAN-IDs can be reused on other ports and therefore mapped to different I-SIDs.
 - Transparent Port UNI — a physical port maps to a Layer 2 VSN I-SID (all traffic through that port, 802.1Q tagged or untagged, ingress and egress is mapped to the I-SID). Note: All VLANs on a Transparent Port UNI interface now share the same single MAC learning table of the Transparent Port UNI I-SID.
- E-Tree UNI — it extends Private VLANs beyond one Switch to form a network-wide E-Tree service infrastructure. An E-Tree UNI is a Layer 2 VSN where broadcast traffic flows from Hub sites to Spokes sites, and from Spokes to Hubs, but not between Spoke sites. E-Tree Hubs can be formed with any VLAN UNI, while E-Tree Spokes must be configured as Private VLAN UNIs.
- Layer 3 VSN UNI — a device-specific VRF maps to an I-SID, and the control plane exchanges the Layer 3 routes belonging to the same I-SID. All VRFs in a network sharing the same Layer 3 I-SID effectively form an Layer 3 VPN. Layer 3 VSNs can be configured to simultaneously support both IP Unicast and IP Multicast.

Transparent Port UNI

Use a Transparent Port User-Network-Interface (Transparent Port UNI or T-UNI) to map an entire port or an MLT to an I-SID. CMAC learning is done against the I-SID. T-UNI configures a transparent port where all traffic is MAC switched on an internal virtual port using the assigned I-SID. No VLAN is involved in this process. Devices switch tagged and untagged traffic in the assigned I-SID regardless of the VLAN ID. The T-UNI port or MLT can be either static or LACP and is not a member of any VLAN or Spanning Tree Group (STG). The T-UNI port or MLT is always in the forwarding state.

You can map multiple ports to a T-UNI I-SID. Multiple ports on the same switch and on other BEBs can use the common I-SID to switch traffic.

T-UNI is a point to point service and all traffic that ingress the UNI egress from the remote UNI end-point

For information about QoS re-marking, see [QoS Re-marking on a Transparent Port UNI](#) on page 3009.

Transparent

T-UNI is transparent because the MAC learning occurs within the I-SID, and packets that ingress from any CVLAN are processed in an identical manner. Devices switch tagged and untagged traffic in the assigned I-SID. Devices switch control protocols, such as BPDU, LACP, LLDP, and others, in the assigned I-SID, rather than forwarding to the CP.

The service classification of packets that are received on a T-UNI port, is independent of the VLAN ID values present in those packets. All data packets received on a T-UNI port are classified into the same service. When data packets enter and exit the T-UNI service, no VLAN tag modifications are performed on the data packets.

T-UNI based MAC learning

When a packet ingresses a port or MLT associated with a T-UNI I-SID, the system performs MAC lookup based on the I-SID. A packet that ingresses a T-UNI port on a BEB can transfer through the SPB network, or can egress out another T-UNI port configured to the same I-SID.

When a packet ingresses a network-to-network interface (NNI) port, before egressing a T-UNI port, the system performs a MAC Destination Address (DA) lookup based on the I-SID. If the DA lookup fails, the packet floods to all T-UNI ports.

Considerations

Consider the following design requirements when you configure a T-UNI:

- Only ELAN based T-UNI is supported. All T-UNI I-SID end points for a given I-SID become members of the same shared ELAN service. If an E-LINE type of service is required, provision T-UNI at the two end points comprising the point-to-point service.
- You cannot configure a T-UNI on the same I-SID as a C-VLAN.
- A port or MLT associated with a T-UNI I-SID cannot be part of any VLAN and does not belong to any STG.
- Ensure that you always associate a T-UNI LACP MLT with a VLAN (even if it is the default VLAN) before adding it to a T-UNI I-SID. Otherwise, traffic is not forwarded on the T-UNI LACP MLT.
- No Layer 3 processing takes place on packets ingressing on a T-UNI port.
- Pause frames do not switch through the T-UNI I-SID.
- Multiple ports or MLTs can be associated with same T-UNI I-SID.
- One port or MLT cannot be part of multiple T-UNI I-SIDs.

- An I-SID mapped to a T-UNI service must not be mapped to any other service, such as Layer 2 VSN and Layer 3 VSN, on any of the remote BEBs in the SPBM network.
- Any Spanning Tree Protocol implementation is disabled on the port or MLT associated with the T-UNI I-SID. The port will always be in a Forwarding state.
- No additional IS-IS TLVs are added to advertise or withdraw T-UNI I-SID services. Extreme Networks makes use of the existing IS-IS TLV-144 and sub TLV-3 to carry I-SID information.
- MACs are learned against the combination of the I-SID and port or MLT.
- The MAC address limit is supported on a per-I-SID basis. For example, the MAC addresses learned on the T-UNI I-SID can be limited.

**Note**

MAC learning limit for T-UNI service is not supported on all hardware platforms.

- Static MAC is not supported for a T-UNI port.
- IP traffic and control packets are transparently bridged over T-UNI endpoints.
- Untagged traffic ingressing on the T-UNI port will use COS 0. B-TAG and I-TAG priorities are derived from the best effort queue that is assigned. If the T-UNI port is set as a Layer 2 untrusted port, a best-effort queue is assigned.
- The 802.1p bits of the incoming traffic are used to derive the B-TAG and I-TAG priorities for tagged traffic.
- LACP, VLACP and LLDP PDUs are extracted to the CP and all other control packets are transparently bridged over the T-UNI port or MLT.

This feature handles control PDUs in the following manner:

All the Layer 2 and Layer 3 control packets are transparently bridged over the T-UNI port or MLT with the exception of LACP, VLACP and LLDP PDUs. LACP PDUs, VLACP PDUs and LLDP PDUs are not transparently bridged over the T-UNI port or MLT if LACP, VLACP or LLDP is enabled on the port or MLT.

- If an LACP MLT is associated with a T-UNI I-SID, LACP PDUs are extracted to CP and processed locally.
- If LACP is not enabled globally and LACP MLT is not associated with the T-UNI I-SID, LACP PDUs are transparently bridged across the T-UNI port or MLT.
- If a VLACP enabled port is added to a T-UNI I-SID, VLACP PDUs are extracted to the CP for local processing. If a port that is not VLACP enabled is added to the T-UNI I-SID, VLACP PDUs are transparently bridged across T-UNI port.
- If a LLDP enabled port is added to a T-UNI I-SID, LLDP PDUs are extracted to the CP for local processing.
- If LLDP is not enabled on the port or MLT interface associated to TUNI I-SID , LLDP PDUs are transparently bridged across the T-UNI port or MLT.

The following list of control packet types are transparently bridged across the T-UNI I-SID:

- SLPP
- VRRP

- OSPF
- RIP
- BGP
- IS-IS
- CFM
- STP
- SONMP

Use T-UNI when either of the following apply:

- All tagged and untagged traffic on a port must be classified into the same broadcast domain.
- You want to offer a transparent provider solution.

An example of an application for T-UNI is a typical Ethernet provider deployment with port-based classification and transparent forwarding.

Transparent Port UNI over vIST

Virtual IST (vIST) provides the ability to dual-home hosts, servers and other network devices to a pair of Multi-Chassis Link Aggregation (MC-LAG) enabled devices. The system displays the MC-LAG nodes to the connected devices as one link-aggregated group. So, although the physical connection is spread between two individual network nodes, logically the system displays them as a single connection.

Transparent Port UNI (T-UNI) over vIST peers extends the capability of dual-home hosts on the SPB cloud to achieve higher network resiliency. The MACs learnt on the T-UNI interface of any one vIST peer is synchronized with the other peer through MAC synchronization.

In the following figure, the T-UNI access switch ACCESS-1 is dual-homed into vIST peer hosts VIST-PEER 1 and VIST-PEER 2. At ACCESS-1, a link aggregation is created to connect to the SPBM cluster. On the VIST peers, an SMLT is created towards ACCESS-1. Depending on the link aggregation hashing logic, traffic is hashed on to VIST-PEER 1 and VIST-PEER 2. The MACs learnt on the T-UNI interfaces of either host is synchronized with the other host.




Figure 91: Example of Transparent Port UNI over vIST

If one of the links between ACCESS-1 and the vIST cluster goes down, all traffic is serviced through the other link. The same applies when any of the vIST peers go down.

Since MAC learning on both peers are synchronized, both peers can switch traffic with the same efficiency.

Single-homed T-UNI service on a vIST-enabled node

If you configure a T-UNI service as a single-homed service on a vIST-enabled node, you must configure the same I-SID service without port/MLT being mapped to I-SID, on the other vIST peer node. Failure to perform this configuration on the vIST peer node can result in the loss of traffic to the single-homed T-UNI service in various scenarios.

Switched UNI

Switched User Network Interface (S-UNI) allows the association of local endpoints to I-SIDs based on local port and VLAN together. With Switched UNI, the same VLAN can be used on one port to create an endpoint to one I-SID, and on another port to create an endpoint to another I-SID.

Switched UNI summary:

- Switched UNI is a VLAN and ports associated with I-SIDs.
- Local significance on the ports.
- You can re-use the same VLAN to associate different ports with different I-SIDs.
- You can use a different VLAN to the same ports, or you can assign different ports to the same I-SID.
- Supports VLAN mapping on the local switch.
- To accept untagged traffic, the port needs to be configured as untagged-traffic in the I-SID.

Use Switched UNI when either of the following apply:

- Vlan ID (VID) reuse is required. The same VID is used on different broadcast domains (multi-tenant applications).
- Multiple VLANs must be part of the same broadcast domain.
- VID translation is required.

An example of an application for Switched UNI is a typical host and provider deployment, with a port and VID-based classification.

Switched UNI based MAC learning

MAC learning is done on I-SID MAC. When a packet ingresses on a port or MLT which is associated with Switched UNI I-SID, the system performs MAC look up based on the I-SID. Switched UNI operates on Any-To-Any (ELAN) mode, there can be one or more ports associated to a Switched UNI I-SID. A packet that ingresses to a Switched UNI port on a BEB can transfer through the SPBM cloud, or can egress out another Switched UNI port configured to the same I-SID.

When a packet ingresses an network-to-network interface (NNI) port, before egressing a Switched UNI port, the system performs a MAC Destination Address (DA) lookup based on the I-SID. If the DA lookup fails, the packet floods to all Switched UNI ports in the I-SID.

Considerations

Consider the following when you configure a Switched UNI:

- The VLAN tag is removed before the traffic egresses out on the untagged-traffic port or MLT.
- VLAN priority received on the packet is maintained across VLAN IDs.
- Spanning tree is disabled on all Switched UNI ports, and the ports remain in forwarding state.
- The Switched UNI I-SID is advertised to the SPBM cloud.
- The Broadcast and unknown Unicast packets are flooded to all ports in the I-SID.

Limitations

- You cannot change from one UNI type to another dynamically. The I-SID has to be deleted and created with new UNI type (Customer VLAN (C-VLAN), Transparent port user-network-interface (T-UNI), ELAN).
- I-SID cannot be used by IPVPN, MVPN, SPBM dynamic multicast range, or Transparent Port UNI.
- If the port is a member of MLT, the entire MLT has to be added to the VID.
- The port is always in the forwarding state.
- The same VID, port, or MLT cannot be member of more than one I-SID.
- Static MAC, Static ARP and static IGMP group are not supported on Switched UNI enabled ports.

BPDUs handling on S-UNI port/MLT

The switch handles Bridge Protocol Data Units (BPDUs) according to whether or not you configure a platform VLAN.

- When you configure a platform VLAN:
 - BPDUs are forwarded to the CPU by default.
 - For both the ingress and egress ports, BPDUs are not flooded in the S-UNI I-SID associated with the platform VLAN.



Note

If the platform VLAN is configured for the S-UNI port, you cannot enable BPDU forwarding.

- When you DO NOT configure a platform VLAN:
 - BPDUs received on untagged-traffic ports are dropped by default.
 - To flood BPDUs in its I-SID, enable BPDU forwarding under S-UNI I-SID using the command **untagged-traffic port <port no> bpdu enable**.

SPBM Sample Operation—Layer 2 VSN

The following section shows how a SPBM network is established, in this case, a Layer 2 VSN.

1. Discover network topology




Figure 92: SPBM topology discover

IS-IS runs on all nodes of the SPBM domain. Since IS-IS is the basis of SPBM, the IS-IS adjacency must be formed first. After the neighboring nodes see hellos from each other they look for the same Level (Level 1) and the same area (for example, Area 2f.8700.0000.00). After the hellos are confirmed both nodes send Link State Protocol Data Units, which contain connectivity information for the SPBM node. These nodes also send copies of all other LSPs they have in their databases. This establishes a network of connectivity providing the necessary information for each node to find the best and proper path to all destinations in the network.

Each node has a system ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (BMAC), which is used as the source and destination MAC address in the SPBM network.

2. Each IS-IS node automatically builds trees from itself to all other nodes

When the network topology is discovered and stored in the IS-IS link state database (LSDB), each node calculates shortest path trees for each source node. A unicast path now exists from every node to every other node

With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes. Multicast FIB is not produced until Layer 2 VSN services are configured and learned.

3. IS-IS advertises new service communities of interest

When a new service is provisioned, its membership is flooded throughout the topology with an IS-IS advertisement.




Figure 93: SPBM BMAC and I-SID population

BMAC and I-SID information is flooded throughout the network to announce new I-SID memberships. In this case, VLAN 20 is mapped to I-SID 100.



Note

I-SIDs are only used for virtual services (Layer 2 and Layer 3 VSNs). If IP Shortcuts only is enabled on the BEBs, I-SIDs are never exchanged in the network as IP Shortcuts allow for IP networks to be transported across IS-IS.

Each node populates its FDB with the BMAC information derived from the IS-IS shortest path tree calculations. Thus there is no traditional flooding and learning mechanism in place for the B-VLAN, but FDBs are programmed by the IS-IS protocol.

4. When a node receives notice of a new service AND is on the shortest path, it updates the FDB

In this scenario, where there are three source nodes having a membership on I-SID 100, there are three shortest path trees calculated (not counting the Equal Cost Trees (ECTs)).




Figure 94: Shortest path tree for source node A

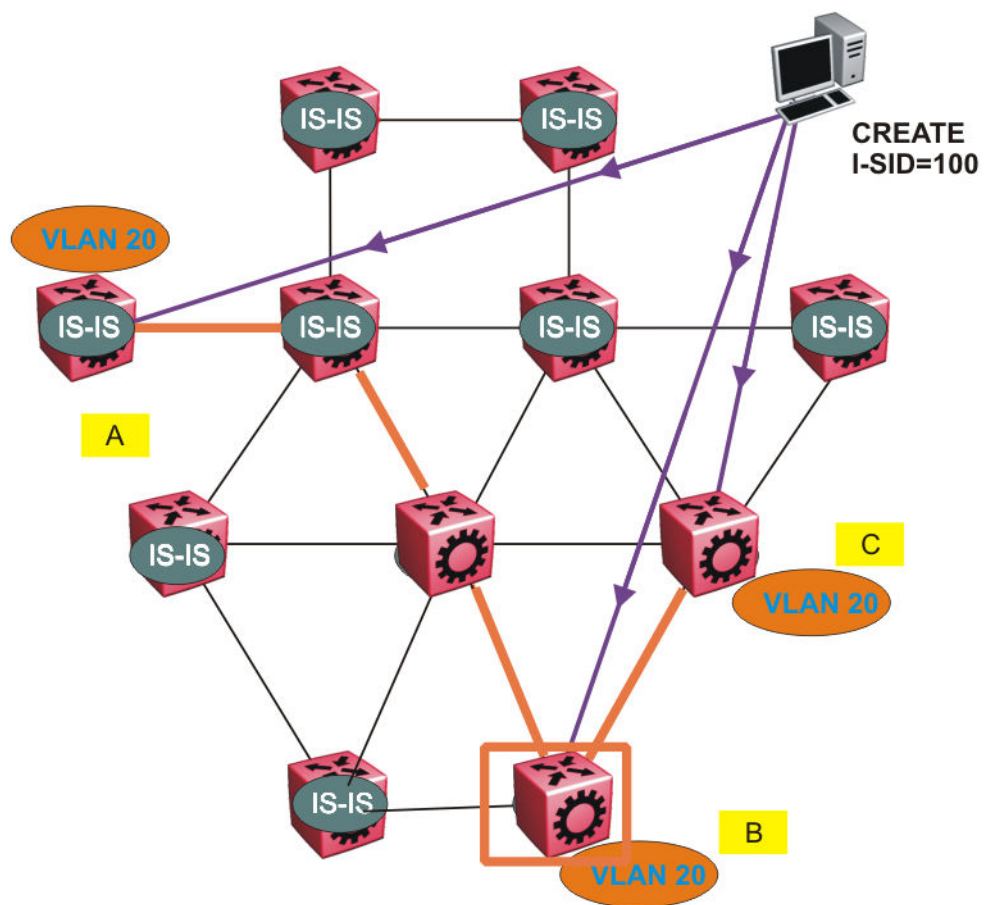


Figure 95: Shortest path tree for source node B




Figure 96: Shortest path tree for source node C

The paths between any two nodes are always the shortest paths. Also, the paths in either direction are congruent, thus a bidirectional communication stream can be monitored easily by mirroring ingress and egress on a link to a network analyzer.

VLAN traffic arriving on switch A and VLAN 20 is forwarded following the blue path, traffic arriving on switch B and VLAN 20 the orange path and on switch C VLAN 20 traffic is following the green path.

If the destination CMAC is unknown at the SPBM ingress node or the traffic is of type broadcast or multicast, then it is flooded to all members of the topology which spans VLAN 20. If the destination CMAC is already known, then the traffic is only forwarded as a unicast to the appropriate destination. In the SPBM domain, the traffic is switched on the BMAC header only. The bridge filtering database (FDB) at the VLAN to I-SID boundary (backbone edge bridge BEB), maintains a mapping between CMACs and corresponding BMACs.

For example, Switch B learns all CMACs which are on VLAN 20 connected to switch A with the BMAC of A in its FDB and the CMACs which are behind C are learned with the BMAC of C.

Layer 2 VSN Configuration using the CLI

This section provides procedures to configure Layer 2 VSN using the command line interface (CLI).

Configure SPBM Layer 2 VSN

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the customer VLANs (C-VLANs) and add slots/ports.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID):

```
vlan i-sid <1-4059> <0-16777215> [force]
```



Important

When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

The switch reserves I-SID 0x00ffffff. The switch uses this I-SID to advertise the virtual BMAC in an SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

3. Display C-VLAN information:

```
show vlan i-sid
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#vlan i-sid 10 100
Switch:1>show vlan i-sid
```

```
=====
                        Vlan I-SID
=====
VLAN_ID      I-SID      VLAN NAME      I-SID NAME
-----
1
10           100       VLAN-10        Hospital-Server-10
90           1000      VLAN-90        ISID-1000
3 out of 3 Total Num of Vlans displayed
```

Variable Definitions

The following table defines parameters for the **vlan i-sid** command.

Variable	Value
<1-4059>	Specifies the primary VLAN ID. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<0-16777215>	Specifies the service instance identifier (I-SID). Note: The switch reserves I-SID 0x00ffffff. The switch uses this I-SID to advertise the virtual BMAC in an SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service. This value is the same for the primary and secondary VLANs. Note: When you add or remove an I-SID on a VLAN IP interface for 5520 Series or 5720 Series, the switch disables and enables routing on the VLAN interface.
<i>force</i>	Specifies the software must replace the existing VLAN-to-I-SID mapping, if one exists.

Configure a Global I-SID Name

About This Task

Use this procedure to provide a descriptive name for the Service Identifier (I-SID).

You can configure a service name for I-SIDs, loopback interfaces, and static routes. You can configure the service name before or after you create the I-SID for the following services:

- Layer 2 VSN
- Layer 3 VSN
- ELAN I-SID or Switched UNI I-SID
- ELAN transparent I-SID or Transparent UNI I-SID
- IPv4 and IPv6 static routes
- IPv4 and IPv6 loopback CLIP interface



Note

The service name for I-SIDs does not support the following special characters:
" " # \$ % ' / [\] ^ { | } ~ @ &.

By default, the service name is ISID-x, where x correlates to the I-SID number of the service.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enter a name for the global I-SID:

```
i-sid name <1-6777215> WORD<1-64>
```
3. Display I-SID names for all configured I-SIDs:

```
show i-sid name
```
4. Display I-SID name by I-SID:

```
show i-sid name <1-6777215>
```

Examples

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#i-sid name 1 ExtremeServer1
Switch:1(config)#i-sid name 20 ExtremeServer7
```

View the configured I-SID names:

```
Switch:1>show i-sid name
=====
I-SID Name
=====
I-SID      I-SID NAME      TYPE
-----
1          ExtremeServer1  adminName
2          ExtremeServer2  adminName
3          ExtremeServer3  config adminName
4          ISID-4          config
23         ISID-23         config
25         ExtremeServer4  config adminName

Total number of I-SID Name entries: 6.
```

View the configured I-SID by number:

```
Switch:1>show i-sid name 1
=====
I-SID Name
=====
I-SID      I-SID NAME      TYPE
-----
1          ExtremeServer1  adminName

Switch:1>show i-sid name 20
=====
I-SID Name
=====
I-SID      I-SID NAME      TYPE
-----
20         ExtremeServer7  adminName
```

Variable Definitions

Use the data in the following table to use the **i-sid name** command.

Variable	Value
<1-6777215>	Specifies the I-SID number.
WORD<1-64>	Specifies the name of the I-SID. The I-SID can be named before or after the I-SID is created. By default, for an I-SID in use, the service is named ISID-x, where x correlates to the I-SID number of the service.

Display C-VLAN I-SID Information

Use the following procedure to display C-VLAN I-SID information.

Procedure

1. Display the C-VLAN to I-SID associations:

```
show vlan i-sid <1-4059>
```
2. Display the IS-IS SPBM multicast-FIB calculation results by I-SID:

```
show isis spbm i-sid {all|config|discover} [vlan <1-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
```
3. Discover where entries are learned:

```
show vlan mac-address-entry [spbm-tunnel-as-mac]
```
4. Display the VLAN remote MAC table for a C-VLAN:

```
show vlan remote-mac-table <1-4059>
```

Examples

```
Switch:1>show vlan i-sid
=====
Vlan I-SID
=====
VLAN_ID    I-SID    VLAN NAME    I-SID NAME
-----
1
10         100     VLAN-10     Hospital-Server-10
90         1000    VLAN-90     ISID-1000

3 out of 3 Total Num of Vlans displayed

Switch# show isis spbm i-sid all
=====
SPBM ISID INFO
=====
ISID      SOURCE NAME    VLAN    SYSID      TYPE      HOST_NAME      ISID NAME    AREA    AREA NAME
-----
101001    1.11.16        4051    0200.10ff.fff0    discover    area-0.00.10    ISID-101001    HOME    area-0.00.20
101003    1.11.16        4051    0200.10ff.fff0    discover    area-0.00.10    ISID-101003    HOME    area-0.00.20
101005    1.11.16        4051    0200.10ff.fff0    discover    area-0.00.10    ISID-101005    HOME    area-0.00.20
101007    1.11.16        4051    0200.10ff.fff0    discover    area-0.00.10    ISID-101007    HOME    area-0.00.20
101009    1.11.16        4051    0200.10ff.fff0    discover    area-0.00.10    ISID-101009    HOME    area-0.00.20
101011    1.11.16        4051    0200.10ff.fff0    discover    area-0.00.10    ISID-101011    HOME    area-0.00.20

-----
Total number of SPBM ISID entries configed: 0
-----
Total number of SPBM ISID entries discovered: 6
-----
Total number of SPBM ISID entries: 6
-----

Switch:# show vlan mac-address-entry
=====
Vlan Fdb
```

```

=====
VLAN      MAC
ID  STATUS ADDRESS      INTERFACE      SMLT
                                REMOTE  TUNNEL
-----
1    learned 00:1d:42:6b:10:03 Port-1/9      false SwitchB
1    learned 00:80:2d:22:ac:46 Port-1/15     false SwitchB
2    self   a4:25:1b:51:48:84 103.103.103.103 false -
2    self   02:01:03:ff:ff:ff Tunnel_to_HQ   false -
5    learned 00:00:00:00:00:1a MLT-1         false SwitchB
10   self   00:00:00:00:49:50 Port-1/9      false -
10   self   00:00:00:50:00:50 Port-1/9      false -

Switch# show vlan remote-mac-table 100
=====
Vlan Remote Mac Table
=====
VLAN STATUS MAC-ADDRESS      DEST-MAC      BVLAN DEST-SYSNAME  PORTS      SMLTREMOTE
-----
100 learned 00:15:40:af:d2:00 00:74:00:00:00:00 20 Switch-6005 MLT-2 false
100 learned b4:a9:5a:04:c8:83 b4:a9:5a:04:c8:65 3 Switch-174 103.103.103.103 true
100 learned b4:a9:5a:04:c8:84 b4:a9:5a:04:c8:66 3 Switch-175 Tunnel_to_HQ true

3 of 3 matching entries out of total of 3 Remote Mac entries in all fdb(s) displayed.
=====

```

Variable definitions

The following table defines parameters for the **show vlan** commands.

Variable	Value
<i>i-sid</i> <1-4059>	Displays I-SID information for the specified C-VLAN.
<i>mac-address-entry</i> [<i>spbm-tunnel-as-mac</i>]	Displays the bridging forwarding database. Use the optional parameter, <i>spbm-tunnel-as-mac</i> to display the BMAC in the TUNNEL column. If you do not use this optional parameter, the TUNNEL column displays the host name. If an entry is not learned in the SPBM network, the TUNNEL column will be empty (-).
<i>remote-mac-table</i> <1-4059>	Displays C-VLAN remote-mac-table information.

The following table defines parameters for the **show isis** commands.

Variable	Value
<i>spbm i-sid</i> { <i>all</i> <i>config</i> <i>discover</i> }	<ul style="list-style-type: none"> all: displays all I-SID entries config: displays configured I-SID entries discover: displays discovered I-SID entries
<i>vlan</i> <1-4059>	Displays I-SID information for the specified SPBM VLAN.
<i>id</i> <1-16777215>	Displays I-SID information for the specified I-SID.
<i>nick-name</i> <x.xx.xx>	Displays I-SID information for the specified nickname.

Configure an SPBM Layer 2 Transparent Port UNI

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes creating the SPBM B-VLANs.

- Associate a T-UNI LACP MLT with a VLAN before mapping the LACP MLT to a T-UNI I-SID.



Caution

In the case of T-UNI LACP SMLT, before you configure SMLT on switch peers, ensure that the T-UNI LACP MLT on each peer is always associated with a VLAN, even if it is the default VLAN, and that it is added to a T-UNI I-SID. Otherwise, traffic is not forwarded on the T-UNI LACP MLT.

About This Task

Use this procedure to configure a Transparent Port UNI or ELAN Transparent service.

You can configure Transparent Port UNI when either of the following apply:

- You want all tagged and untagged traffic on a port to be classified into the same broadcast domain.
- You want to offer a transparent provider solution.



Note

If you configure a T-UNI to terminate on a port or MLT on a switch in a vIST switch cluster, you must also configure the T-UNI I-SID on the other switch of the vIST switch cluster. You must configure the T-UNI I-SID on both switches of a vIST pair. You do not need to assign an actual port or MLT to the T-UNI on the second switch.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a Transparent Port UNI (ELAN-Transparent based service):

```
i-sid <1-16777215> elan-transparent
```

This command automatically transitions you to the ELAN-Transparent I-SID Configuration mode.

3. Add ports to the ELAN-Transparent based service:

```
port {slot/port[-slot/port]}[,...][slot/all][all]}
```

A warning message displays to indicate that adding a port to a T-UNI I-SID removes the port from all VLANs. Type *y* to continue.

4. Add an MLT to the Elan-Transparent based service:

```
mlt <1-512>
```

A warning message displays to indicate that adding an MLT to a Transparent Port UNI I-SID removes the MLT from all VLANs. Type *y* to continue.

5. To verify the Transparent Port UNI configuration, enter:

```
show i-sid <1-16777215>
```

Example

Configure a Transparent Port UNI I-SID (ELAN-Transparent based service).

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#i-sid 3 elan-transparent
Switch:1(elan-tp:3)#port 1/25
Adding Ports to Transparent UNI i-sid removes it from all VLANs.
Do you wish to continue (y/n) ? y
Switch:1(elan-tp:3)#mlt 1
Adding MLTs to Transparent UNI i-sid removes it from all VLANs.
Do you wish to continue (y/n) ? y
```

Verify Transparent Port UNI or ELAN-Transparent based service configuration.

```
Switch:1(config)#show i-sid 3
=====
Isid Info
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES
-----
3         ELAN_TR   N/A       -         -           CONFIG     ISID-3
```

Variable definitions

The following table defines parameters for the **i-sid** command.



Note
When SPB is enabled, I-SID IDs 16000000 (0xF42400) and greater, up to 16,777,215 (0xFFFFF), are reserved for dynamic i-sid allocation and used to support IP Multicast traffic over SPB and other advanced Fabric services.

Variable	Value
<i>i-sid</i> <1-16777215> <i>elan-transparent</i>	Creates an ELAN-Transparent based service. The service interface identifier (I-SID) range is 1 to 16777215.
<i>port</i> { <i>slot/port</i> [- <i>slot/port</i>] [<i>,...</i>][<i>slot/all</i>][<i>all</i>]}	Add ports to the ELAN-Transparent based service.
<i>mlt</i> <1-512>	Add MLTs to the ELAN-Transparent based service. The MLT range is 1 to 512.

View all Configured I-SIDs

About This Task

View all configured I-SIDs (both CVLAN and T-UNI). View also the I-SID types and the ports or MLTs that are assigned to each I-SID.

Procedure

1. Enter Privileged EXEC mode:
enable

2. View all configured I-SIDs.

This command displays both CVLAN and T-UNI based I-SIDs.

```
show i-sid
```

3. View all T-UNI (ELAN-Transparent) I-SIDs.

```
show i-sid [elan-transparent]
```

4. View information for a particular T-UNI I-SID.

```
show i-sid [<1-16777215>]
```

5. View all IS-IS SPBM I-SID information by I-SID ID:

```
show isis spbm i-sid {all|config|discover} [vlan <2-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
```

Examples

View all configured I-SIDs.

```
Switch:1(config)#show i-sid
=====
Isid Info
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES   NAME
-----
15999999  ELAN      4048      -         -           C --- - --- - -   Onboarding I-SID
16777001  ELAN      N/A       -         -           C --- - --- - -   FAN-ISID

c: customer vid    u: untagged-traffic

All 2 out of 2 Total Num of i-sids displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense; R: multi-area redistrib
l: discover by local switch  r: discover by remote VIST switch
```

View T-UNI (ELAN-Transparent) I-SIDs.

```
Switch:1 (config)#show i-sid elan-transparent
=====
Isid Info
=====
ISID      ISID      PORT      MLT      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES NAME
-----
2         ELAN_TR   N/A       -         -       ExtremeServer2
25        ELAN_TR   N/A       1/2-1/8,8/11 25      ExtremeServer4

All 1 out of 1 Total Num of elan-tp i-sids displayed
```

View MLT or port information for a particular T-UNI I-SID.

```
Switch:1(config)#show i-sid 111
=====
Isid Info
=====
ISID      ISID      PORT      MLT      ORIGIN      ISID
ID        TYPE      VLANID    INTERFACES INTERFACES   NAME
-----
111       ELAN_TR   N/A       1/2-1/8,8/11 111        CONFIG      ISID-111
```

View all IS-IS SPBM I-SID information:

```
Switch# show isis spbm i-sid all
=====
SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID      TYPE      HOST_NAME      ISID NAME      AREA      AREA NAME
-----
101001    1.11.16     4051  0200.10ff.fff0 discover    area-0.00.10   ISID-101001    HOME      area-0.00.20
```

```

101003  1.11.16      4051  0200.10ff.fff0    discover  area-0.00.10    ISID-101003    HOME    area-0.00.20
101005  1.11.16      4051  0200.10ff.fff0    discover  area-0.00.10    ISID-101005    HOME    area-0.00.20
101007  1.11.16      4051  0200.10ff.fff0    discover  area-0.00.10    ISID-101007    HOME    area-0.00.20
101009  1.11.16      4051  0200.10ff.fff0    discover  area-0.00.10    ISID-101009    HOME    area-0.00.20
101011  1.11.16      4051  0200.10ff.fff0    discover  area-0.00.10    ISID-101011    HOME    area-0.00.20

```

```
-----
Total number of SPBM ISID entries configed: 0
-----
```

```
-----
Total number of SPBM ISID entries discovered: 6
-----
```

```
-----
Total number of SPBM ISID entries: 6
-----
```

View all IS-IS SPBM I-SID information by I-SID ID:

```
Switch:1#show isis spbm i-sid all id 300
```

```

=====
SPBM ISID INFO
=====
ISID  SOURCE NAME  VLAN  SYSID          TYPE  HOST_NAME  ISID NAME  AREA  AREA NAME
-----
300   7.15.16      20    a425.1b51.9484 config  Switch1    ISID-300   HOME  area-0.00.20
300   4.01.18      10    b4a9.5a2a.d065 discover Switch2    ISID-300   HOME  area-0.00.20

```

```
-----
Total number of SPBM ISID entries configured: 1
-----
```

```
-----
Total number of SPBM ISID entries discovered: 1
-----
```

```
-----
Total number of SPBM ISID entries: 2
-----
```

Variable Definitions



Note

When SPB is enabled, I-SID IDs 16777216 and greater are reserved for internal I-SID and SPB multicast.

The following table defines parameters for the **show i-sid** command.

Variable	Value
<1-16777215>	Specifies the service interface identifier (I-SID).
<i>elan-transparent</i>	Displays only all the ELAN-Transparent (T-UNI based) I-SIDs.

The following table defines parameters for the **show isis spbm i-sid** command.

Variable	Value
{all config discover}	<ul style="list-style-type: none"> all: displays all I-SID entries config: displays configured I-SID entries discover: displays discovered I-SID entries
vlan <2-4059>	Displays I-SID information for the specified SPBM VLAN.
id <1-16777215>	Displays I-SID information for the specified I-SID.
nick-name <x.xx.xx>	Displays I-SID information for the specified nickname.

View CMACs Learned on T-UNI Ports for an I-SID

About This Task

Perform this procedure to view the I-SID bridge forwarding database.

The **show i-sid mac-address-entry** command displays the CMACs learned on T-UNI I-SIDs. It also displays the CMACs learned on T-UNI I-SIDs for a specific I-SID, MAC address, port or port list or remote MAC address.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View CMACs learned on the T-UNI I-SIDs:

```
show i-sid mac-address-entry [<1-16777215>] [home] [mac
<0x00:0x00:0x00:0x00:0x00:0x00>] [non-local] [port {slot/port[/sub-
port] [-slot/port[/sub-port]] [,...][slot/all][all]}] [remote]
```

Examples

View CMACs learned on all T-UNI I-SIDs.

```
Switch:1#show i-sid mac-address-entry
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
100    learned  cc:f9:54:ae:28:81  Port-1/16  LOCAL  00:00:00:00:00:00  0      HOME          area-20.0020
4      learned  cc:f9:54:ae:2c:18  MLT-6      LOCAL  00:00:00:00:00:00  0      HOME          area-20.0020
252    learned  cc:f9:54:ae:38:64  Port-1/15  NON-LOCAL 00:13:0a:0c:d3:e0  128    DIST-1B      REMOTE      area-30.0030

All 3 out of 3 Total Num of i-sid FDB Entries displayed
```

View CMACs learned on a specific T-UNI I-SID.

```
Switch:1#show i-sid mac-address-entry 100
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
100    learned  cc:f9:54:ae:28:81  Port-1/16  LOCAL  00:00:00:00:00:00  0      HOME          area-20.0020

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

```
Switch:1#show i-sid mac-address-entry 252
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
252    learned  cc:f9:54:ae:38:64  Port-1/15  NON-LOCAL 00:13:0a:0c:d3:e0  128    DIST-1B      REMOTE      area-30.0030

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View CMACs learned on a T-UNI I-SID for a specific MAC address.

```
Switch:1#show i-sid mac-address-entry mac cc:f9:54:ae:38:64
```

```
=====
I-SID Fdb Table
=====
I-SID  STATUS  MAC-ADDRESS  INTERFACE  TYPE  DEST-MAC  BVLAN  DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
252    learned  cc:f9:54:ae:38:64  Port-1/15  NON-LOCAL 00:13:0a:0c:d3:e0  128    DIST-1B      REMOTE      area-30.0030

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```


View CMACs learned on aT-UNI I-SID for a specific port.

```
Switch:1#show i-sid mac-address-entry port 1/15
```

```
=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE      DEST-MAC      BVLAN DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
252   learned  cc:f9:54:ae:38:64 Port-1/15  NON-LOCAL  00:13:0a:0c:d3:e0 128   DIST-1B      REMOTE     area-30.0030

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View CMACs learned on a T-UNI I-SID as a remote MAC address.

```
Switch:1#show i-sid mac-address-entry remote
```

```
=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE      DEST-MAC      BVLAN DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
252   learned  cc:f9:54:ae:38:64 Port-1/15  NON-LOCAL  00:13:0a:0c:d3:e0 128   DIST-1B      REMOTE     area-30.0030

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View CMACs learned on a T-UNI I-SID as a home MAC address.

```
Switch:1#show i-sid mac-address-entry home
```

```
=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE      DEST-MAC      BVLAN DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
100   learned  cc:f9:54:ae:28:81 Port-1/16  LOCAL      00:00:00:00:00:00 0       HOME        area-20.0020

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

View CMACs learned on a T-UNI I-SID as a non-local MAC address.

```
Switch:1#show i-sid mac-address-entry non-local
```

```
=====
I-SID Fdb Table
=====
I-SID STATUS  MAC-ADDRESS      INTERFACE  TYPE      DEST-MAC      BVLAN DEST-SYSNAME  AREA-ROLE  AREA-NAME
-----
252   learned  cc:f9:54:ae:38:64 Port-1/15  NON-LOCAL  00:13:0a:0c:d3:e0 128   DIST-1B      REMOTE     area-30.0030

All 1 out of 1 Total Num of i-sid FDB Entries displayed
```

Variable Definitions

The following table defines parameters for the **show i-sid mac-address-entry** command.

Variable	Value
<1-16777215>	Displays the MAC address learned on the service interface identifier (I-SID).
<i>home</i>	Filters the command output to show only MAC addresses learned in the home area.
<i>mac</i> <0x00:0x00:0x00:0x00:0x00:0x00>	Displays the I-SID FDB details for the specified MAC address.
<i>non-local</i>	Filters the command output to show only MAC addresses learned from other nodes; not local nodes.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	Displays the MAC address learned on the specified port or port list.
<code>remote</code>	Filters the command output to show only MAC addresses learned in the remote area.

Configure an SPBM Layer 2 Switched UNI on an MLT

Before You Begin

You must configure the required SPBM and IS-IS infrastructure.

About This Task

To configure a Switched UNI on an MLT, you must create a Switched UNI I-SID, and map an MLT to the Switched UNI I-SID.



Note

When you configure Switched UNI, Spanning tree is disabled on all Switched UNI ports.

Procedure

1. Enter MLT Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface mlt <1-512>
```

2. Enable S-UNI on MLT:

```
flex-uni enable
```



Note

You cannot enable Switched UNI on EAPoL enabled interface.

3. Configure a Switched UNI Service Instance Identifier (I-SID):

```
i-sid <1-16777215> [elan]
```

This command automatically transitions you to the ELAN I-SID Configuration mode.

4. Add an MLT to a Switched UNI I-SID:

```
c-vid <1-4094> mlt <1-512>
```



Note

You can run this command again to map a Switched UNI MLT to multiple I-SIDs.

5. Add untagged traffic to a Switched UNI I-SID:

```
untagged-traffic mlt <1-512> [bpdu enable]
```

6. Display the Switched UNI information:

```
show mlt i-sid
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#mlt 10
Switch:1(config)#interface mlt 10
Switch:1(config-mlt)#flex-uni enable
Switch:1(config-mlt)#i-sid 100
Switch:1(elan:100)#c-vid 20 mlt 10
Switch:1(elan:100)#untagged-traffic mlt 10 bpdu enable
Switch:1(elan:100)#show mlt i-sid
=====
MLT Isid Info
=====
```

ISID		ISID		ISID		BPDU	
MLTID	IFINDEX	ID	VLANID	C-VID	TYPE	ORIGIN	NAME
10	6153	100	N/A	20	ELAN	C --- - --- -	EXTR
11	6154	100	N/A	11	ELAN	C --- - --- -	ISID-100

```
-----
2 out of 2 Total Num of i-sid endpoints displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense
l: discover by local switch r: discover by remote VIST switch
```

Variable definitions

The following table defines parameters for the **i-sid** command to configure a Switched UNI.

Variable	Value
<i>i-sid</i> <1-16777215>	Specifies the service interface identifier (I-SID).
<i>c-vid</i> <1-4094>	Specifies the customer VLAN ID.
<i>mlt</i> <1-512>	Specifies the MLT ID.

Configure an SPBM Layer 2 Switched UNI on a Port

Before You Begin

- You must configure the required SPBM and IS-IS infrastructure.
- To configure an untagged Switched UNI (S-UNI) on an EAP-enabled interface, you must associate the untagged S-UNI with a platform VLAN.

About This Task



Note

EAP and FA can coexist on the same port. EAP and FA can be enabled in any order; however, EAP must have Flex UNI enabled to function on an FA-enabled port. If EAP is currently enabled, FA can only be enabled if the port is a Flex UNI-enabled port.

To configure an S-UNI on a port, you must create an S-UNI I-SID, and map the port to the S-UNI I-SID.



Note

When you configure S-UNI, Spanning tree is disabled on all S-UNI ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable S-UNI on a port:

```
flex-uni enable
```

3. Configure an S-UNI Service Instance Identifier (I-SID):

```
i-sid <1-16777215> [elan]
```

This command automatically transitions you to the ELAN I-SID Configuration mode.

4. Add ports to an S-UNI I-SID:

```
c-vid <1-4094> port {slot/port[-slot/port] [,...][slot/all][all]}
```

5. Add untagged traffic to an S-UNI I-SID:

```
untagged-traffic port {slot/port[-slot/port] [,...][slot/all][all]}
[bpdu enable]
```

6. Display the S-UNI information:

```
show interfaces gigabitethernet i-sid {slot/port[-slot/port] [,...][slot/all][all]}
```

Examples

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/1,1/2
Switch:1(config-if)#flex-uni enable
```

```
Switch:1(config-if)#i-sid 100
Switch:1(elan:100)#c-vid 10 port 1/1,1/2
Switch:1(elan:100)#untagged-traffic port 1/1,1/2 bpdu enable
```

```
Switch:1#show interfaces gigabitEthernet i-sid
```

```
=====
                        PORT Isid Info
=====
```

PORTNUM	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU	MAC SUNI
1/1	192	27	N/A	4000	ELAN	C --- - --- -	ISID-27		FALSE
1/1	192	270	N/A	4001	ELAN	C --- - --- -	ISID-270		FALSE
1/1	192	309	N/A	309	ELAN	C --- - --- -	ISID-309		FALSE
1/1	192	401	N/A	401	ELAN	C --- - --- -	ISID-401		FALSE
1/1	192	1001	N/A	1001	ELAN	C --- - --- -	ISID-1001		FALSE
1/1	192	1111	N/A	1111	ELAN	C --- - --- -	ISID-1111		FALSE
1/1	192	1121	N/A	1121	ELAN	C --- - --- -	ISID-1121		FALSE
1/1	192	1201	N/A	1201	ELAN	C --- - --- -	ISID-1201		FALSE
1/1	192	2001	N/A	2001	ELAN	C --- - --- -	ISID-2001		FALSE
1/2	193	38	N/A	4000	ELAN	C --- - --- -	ISID-38		FALSE
1/2	193	310	N/A	310	ELAN	C --- - --- -	ISID-310		FALSE
1/2	193	380	N/A	4001	ELAN	C --- - --- -	ISID-380		FALSE
1/2	193	402	N/A	402	ELAN	C --- - --- -	ISID-402		FALSE

```
13 out of 152 Total Num of i-sid endpoints displayed
```

ORIGIN Legend:

C: manually configured; D: discovered by FA or EPT

M: FA management; E: discovered by EAP; A: auto-sense

l: discover by local switch r: discover by remote VIST switch

Variable Definitions

The following table defines parameters for the **i-sid** command to configure a Switched UNI.

Variable	Value
<i>i-sid</i> <1-16777215>	Specifies the service interface identifier (I-SID).
<i>c-vid</i> <1-4094>	Specifies the customer VLAN ID.
{ <i>slot/port[/sub-port]</i> [- <i>slot/port[/sub-port]</i>][, ...] [<i>slot/all</i>][<i>all</i>]}	Specifies the ports associated with the S-UNI I-SID.
<i>bpdu enable</i>	Enables bridge packet data units (BPDU) for an untagged end point.

View All Configured Switched UNI I-SIDs

Perform this procedure to view all the configured Switched UNI I-SIDs including their types, ports, and MLTs.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View all configured CVLAN, T-UNI, and S-UNI based I-SIDs:
show i-sid

3. View all S-UNI I-SIDs.

```
show i-sid [elan]
```

4. View all associated MLT on the S-UNI I-SID.

```
show mlt i-sid [MLT ID <1-512>]
```

5. View all associated ports on the S-UNI I-SID.

```
show interfaces gigabitethernet i-sid {slot/port[/sub-port] [-slot/
port[/sub-port]][,...][slot/all][all]}
```

6. View all IS-IS SPBM multicast FIB entries.

```
show isis spbm multicast-fib detail
```

Examples

View all configured I-SIDs.

```
Switch:1#show i-sid
=====
Isid Info
=====
```

ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN	ISID NAME
15999999	ELAN	4048	-	-	C --- - - - -	Onboarding I-SID
16777001	ELAN	N/A	-	-	C --- - - - -	FAN-ISID

```

c: customer vid    u: untagged-traffic

All 2 out of 2 Total Num of i-sids displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
M: FA management; E: discovered by EAP; A: auto-sense; R: multi-area redistrib
l: discover by local switch  r: discover by remote VIST switch
```

View all S-UNI I-SIDs.

```
Switch:1>show i-sid elan
=====
Isid Info
=====
```

ISID ID	ISID TYPE	VLANID	PORT INTERFACES	MLT INTERFACES	ORIGIN	ISID NAME
27	ELAN	N/A	c4000:1/1,2/1	-	C --- - - - -	ISID-27
38	ELAN	N/A	c4000:1/2,2/2	-	C --- - - - -	ISID-38
270	ELAN	N/A	c4001:1/1,2/1	-	C --- - - - -	ISID-270
307	ELAN	N/A	c307:1/5,2/1	-	C --- - - - -	ISID-307
308	ELAN	N/A	c308:1/6,2/1	-	C --- - - - -	ISID-308
309	ELAN	N/A	c309:1/1,2/1	-	C --- - - - -	ISID-309
310	ELAN	N/A	c310:1/2,2/2	-	C --- - - - -	ISID-310
311	ELAN	N/A	c311:1/3,2/2	-	C --- - - - -	ISID-311
312	ELAN	N/A	c312:1/4,2/2	-	C --- - - - -	ISID-312
317	ELAN	N/A	c317:1/7,2/1	-	C --- - - - -	ISID-317
318	ELAN	N/A	c318:1/8,2/1	-	C --- - - - -	ISID-318
319	ELAN	N/A	c319:1/9,2/1	-	C --- - - - -	ISID-319
320	ELAN	N/A	c320:1/10,2/1	-	C --- - - - -	ISID-320

```

--More-- (q = quit)

c: customer vid    u: untagged-traffic

13 out of 77 Total Num of Elan displayed

ORIGIN Legend:
C: manually configured; D: discovered by FA or EPT
```

M: FA management; E: discovered by EAP; A: auto-sense
 l: discover by local switch r: discover by remote VIST switch

View all associated MLT on the S-UNI I-SID.

Switch:1>show mlt i-sid

MLT Isid Info									
MLTID	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU	
3	6146	3	N/A	33	ELAN	C --- - --- -	ISID-3		

1 out of 1 Total Num of i-sid endpoints displayed

ORIGIN Legend:

C: manually configured; D: discovered by FA or EPT

M: FA management; E: discovered by EAP; A: auto-sense

l: discover by local switch r: discover by remote VIST switch

View all associated ports on the S-UNI I-SID.

Switch:1#show interfaces gigabitEthernet i-sid

PORT Isid Info										
PORTNUM	IFINDEX	ISID ID	VLANID	C-VID	ISID TYPE	ORIGIN	ISID NAME	BPDU	MAC	SUNI
1/1	192	27	N/A	4000	ELAN	C --- - --- -	ISID-27		FALSE	
1/1	192	270	N/A	4001	ELAN	C --- - --- -	ISID-270		FALSE	
1/1	192	309	N/A	309	ELAN	C --- - --- -	ISID-309		FALSE	
1/1	192	401	N/A	401	ELAN	C --- - --- -	ISID-401		FALSE	
1/1	192	1001	N/A	1001	ELAN	C --- - --- -	ISID-1001		FALSE	
1/1	192	1111	N/A	1111	ELAN	C --- - --- -	ISID-1111		FALSE	
1/1	192	1121	N/A	1121	ELAN	C --- - --- -	ISID-1121		FALSE	
1/1	192	1201	N/A	1201	ELAN	C --- - --- -	ISID-1201		FALSE	
1/1	192	2001	N/A	2001	ELAN	C --- - --- -	ISID-2001		FALSE	
1/2	193	38	N/A	4000	ELAN	C --- - --- -	ISID-38		FALSE	
1/2	193	310	N/A	310	ELAN	C --- - --- -	ISID-310		FALSE	
1/2	193	380	N/A	4001	ELAN	C --- - --- -	ISID-380		FALSE	
1/2	193	402	N/A	402	ELAN	C --- - --- -	ISID-402		FALSE	

13 out of 152 Total Num of i-sid endpoints displayed

ORIGIN Legend:

C: manually configured; D: discovered by FA or EPT

M: FA management; E: discovered by EAP; A: auto-sense

l: discover by local switch r: discover by remote VIST switch

View all IS-IS SPBM multicast FIB entries.

Switch:1#show isis spbm multicast-fib detail

SPBM MULTICAST FIB ENTRY DETAIL INFO									
MCAST DA	ISID	BVLAN	SYSID	HOST-NAME	OUTGOING-INTERFACES	INCOMING INTERFACE	CVLAN	AREA	AREA-NAME
03:00:20:00:00:65	101	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	101	HOME	area-10.0010
03:00:20:00:00:67	103	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	103	HOME	area-10.0010
03:00:20:00:00:69	105	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	105	HOME	area-10.0010
03:00:20:00:00:6b	107	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	107	HOME	area-10.0010
03:00:20:00:00:6d	109	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	109	HOME	area-10.0010
03:00:20:00:00:6f	111	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	111	HOME	area-10.0010
03:00:20:00:00:71	113	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	113	HOME	area-10.0010

```

03:00:20:00:00:73 115 4051 0200.20ff.fff0 vn-area-20.0020 1/32 Virtual 115 HOME area-10.0010
03:00:20:00:00:75 117 4051 0200.20ff.fff0 vn-area-20.0020 1/32 Virtual 117 HOME area-10.0010
03:00:20:00:00:77 119 4051 0200.20ff.fff0 vn-area-20.0020 1/32 Virtual 119 HOME area-10.0010
03:00:20:00:00:79 121 4051 0200.20ff.fff0 vn-area-20.0020 1/32 Virtual 121 HOME area-10.0010
03:00:20:00:00:7b 123 4051 0200.20ff.fff0 vn-area-20.0020 1/32 Virtual 123 HOME area-10.0010
13:00:11:00:00:65 101 4051 a425.1b53.7084 BEB1-7200 1/32 MLT-2 101 HOME area-10.0010

```

```

-----
13 out of 2354 Total Num of SPBM MULTICAST FIB entries displayed
Remote: 0 out of 0 Total Num of SPBM MULTICAST FIB entries displayed
-----

```

Variable Definitions

The following table defines parameters for the **i-sid** command.

Variable	Value
elan	Displays only all the ELAN (S-UNI based) I-SIDs.
MLT ID <1-512>	Specifies the MLT associated with the Switched UNI I-SID.
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}	Specifies the ports associated with the Switched UNI I-SID.

Display C-VLAN and Switched UNI I-SID Information

Use the following procedure to display C-VLAN and Switched UNI (S-UNI) I-SID information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the C-VLAN to I-SID associations:

```
show vlan i-sid <1-4059>
```
3. Display I-SID information and Switched UNI to I-SID associations:

```
show i-sid <1-16777215>
```
4. Display the IS-IS SPBM multicast-FIB calculation results by I-SID:

```
show isis spbm i-sid {all|config|discover} [vlan <1-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
```
5. Display all ELAN I-SIDs:
 - ```
show i-sid elan
```
6. Display I-SID configured on MLT:
  - ```
show mlt i-sid
```
7. Display I-SID configured on port:
 - ```
show interfaces gigabitethernet i-sid
```

#### Examples

```

Switch:1>show isis spbm i-sid all
=====
 SPBM ISID INFO
=====
ISID SOURCE NAME VLAN SYSID TYPE HOST_NAME ISID NAME AREA AREA NAME

101001 1.11.16 4051 0200.10ff.fff0 discover area-0.00.10 ISID-101001 HOME area-0.00.20

```



```

101003 1.11.16 4051 0200.10ff.fff0 discover area-0.00.10 ISID-101003 HOME area-0.00.20
101005 1.11.16 4051 0200.10ff.fff0 discover area-0.00.10 ISID-101005 HOME area-0.00.20
101007 1.11.16 4051 0200.10ff.fff0 discover area-0.00.10 ISID-101007 HOME area-0.00.20
101009 1.11.16 4051 0200.10ff.fff0 discover area-0.00.10 ISID-101009 HOME area-0.00.20
101011 1.11.16 4051 0200.10ff.fff0 discover area-0.00.10 ISID-101011 HOME area-0.00.20

```

```

Total number of SPBM ISID entries configed: 0

```

```

Total number of SPBM ISID entries discovered: 6

```

```

Total number of SPBM ISID entries: 6

```

```
Switch:1>show i-sid elan
```

#### Isid Info

```

=====
ISID ISID PORT MLT ORIGIN ISID
ID TYPE VLANID INTERFACES INTERFACES NAME

27 ELAN N/A c4000:1/1,2/1 - C --- - --- - ISID-27
38 ELAN N/A c4000:1/2,2/2 - C --- - --- - ISID-38
270 ELAN N/A c4001:1/1,2/1 - C --- - --- - ISID-270
307 ELAN N/A c307:1/5,2/1 - C --- - --- - ISID-307
308 ELAN N/A c308:1/6,2/1 - C --- - --- - ISID-308
309 ELAN N/A c309:1/1,2/1 - C --- - --- - ISID-309
310 ELAN N/A c310:1/2,2/2 - C --- - --- - ISID-310
311 ELAN N/A c311:1/3,2/2 - C --- - --- - ISID-311
312 ELAN N/A c312:1/4,2/2 - C --- - --- - ISID-312
317 ELAN N/A c317:1/7,2/1 - C --- - --- - ISID-317
318 ELAN N/A c318:1/8,2/1 - C --- - --- - ISID-318
319 ELAN N/A c319:1/9,2/1 - C --- - --- - ISID-319
320 ELAN N/A c320:1/10,2/1 - C --- - --- - ISID-320

```

```
--More-- (q = quit)
```

```
c: customer vid u: untagged-traffic
```

```
13 out of 77 Total Num of Elan displayed
```

ORIGIN Legend:

C: manually configured; D: discovered by FA or EPT

M: FA management; E: discovered by EAP; A: auto-sense

l: discover by local switch r: discover by remote VIST switch

```
Switch:1>show mlt i-sid
```

#### MLT Isid Info

```

=====
MLTID IFINDEX ISID VLANID C-VID ISID ORIGIN ISID BPDU
 ID TYPE

3 6146 3 N/A 33 ELAN C --- - --- - ISID-3

```

```
1 out of 1 Total Num of i-sid endpoints displayed
```

ORIGIN Legend:

C: manually configured; D: discovered by FA or EPT

M: FA management; E: discovered by EAP; A: auto-sense

l: discover by local switch r: discover by remote VIST switch

```
Switch:1#show interfaces gigabitEthernet i-sid
```

#### PORT Isid Info

```

=====
ISID ISID ISID MAC

```

| PORTNUM | IFINDEX | ID   | VLANID | C-VID | TYPE | ORIGIN        | NAME      | BPDU | SUNI  |
|---------|---------|------|--------|-------|------|---------------|-----------|------|-------|
| 1/1     | 192     | 27   | N/A    | 4000  | ELAN | C --- - --- - | ISID-27   |      | FALSE |
| 1/1     | 192     | 270  | N/A    | 4001  | ELAN | C --- - --- - | ISID-270  |      | FALSE |
| 1/1     | 192     | 309  | N/A    | 309   | ELAN | C --- - --- - | ISID-309  |      | FALSE |
| 1/1     | 192     | 401  | N/A    | 401   | ELAN | C --- - --- - | ISID-401  |      | FALSE |
| 1/1     | 192     | 1001 | N/A    | 1001  | ELAN | C --- - --- - | ISID-1001 |      | FALSE |
| 1/1     | 192     | 1111 | N/A    | 1111  | ELAN | C --- - --- - | ISID-1111 |      | FALSE |
| 1/1     | 192     | 1121 | N/A    | 1121  | ELAN | C --- - --- - | ISID-1121 |      | FALSE |
| 1/1     | 192     | 1201 | N/A    | 1201  | ELAN | C --- - --- - | ISID-1201 |      | FALSE |
| 1/1     | 192     | 2001 | N/A    | 2001  | ELAN | C --- - --- - | ISID-2001 |      | FALSE |
| 1/2     | 193     | 38   | N/A    | 4000  | ELAN | C --- - --- - | ISID-38   |      | FALSE |
| 1/2     | 193     | 310  | N/A    | 310   | ELAN | C --- - --- - | ISID-310  |      | FALSE |
| 1/2     | 193     | 380  | N/A    | 4001  | ELAN | C --- - --- - | ISID-380  |      | FALSE |
| 1/2     | 193     | 402  | N/A    | 402   | ELAN | C --- - --- - | ISID-402  |      | FALSE |

13 out of 152 Total Num of i-sid endpoints displayed

ORIGIN Legend:  
 C: manually configured; D: discovered by FA or EPT  
 M: FA management; E: discovered by EAP; A: auto-sense  
 l: discover by local switch r: discover by remote VIST switch

### Variable Definitions

The following table defines parameters for the **show vlan i-sid** commands.

| Variable | Value                                                                             |
|----------|-----------------------------------------------------------------------------------|
| <1-4059> | Displays I-SID information for the specified C-VLAN. You can specify the VLAN ID. |

The following table defines parameters for the **show i-sid** commands.

| Variable     | Value                                                     |
|--------------|-----------------------------------------------------------|
| <1-16777215> | Displays I-SID information. You can specify the I-SID ID. |

The following table defines parameters for the **show isis** commands.

| Variable                                | Value                                                                                                                                                                                     |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>spbm i-sid {all config discover}</i> | <ul style="list-style-type: none"> <li>all: displays all I-SID entries</li> <li>config: displays configured I-SID entries</li> <li>discover: displays discovered I-SID entries</li> </ul> |

## Layer 2 VSN configuration using EDM

This section provides procedures to configure Layer 2 Virtual Services Networks (VSNs) using Enterprise Device Manager (EDM).

*Configure SPBM Layer 2 VSN***Before You Begin**

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.

**Procedure**

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. To map a C-VLAN to a Service instance identifier (I-SID), in the **I-sid** field, specify the I-SID to associate with the specified VLAN.
5. Select **Apply**.

**Important**

- When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.
- The switch reserves I-SID 0x00ffffff. The switch uses this I-SID to advertise the virtual B-MAC in an SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

*Display the remote MAC table for a C-VLAN*

Use the following procedure to view a the remote MAC table for a C-VLAN.

**Procedure**

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Remote MAC** tab.

**Remote MAC Field Descriptions**

Use the data in the following table to use the **Remote MAC** tab.

| Name                  | Description                                                                                          |
|-----------------------|------------------------------------------------------------------------------------------------------|
| <b>VlanId</b>         | Indicates the VLAN ID for this MAC address.                                                          |
| <b>Addr</b>           | Indicates the customer MAC address for which the bridge has forwarding and/or filtering information  |
| <b>DestAddr</b>       | Indicates the provider MAC address for which the bridge has forwarding and/or filtering information. |
| <b>PrimaryBVlanId</b> | Indicates the primary B-VLAN ID for this MAC address.                                                |

| Name                        | Description                                                                                                                                                         |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PrimaryDestSysName</b>   | Indicates the primary system name of the node where the MAC address entry comes from.                                                                               |
| <b>PrimaryPort</b>          | Either displays the value 0, or indicates the primary port on which a frame came from.                                                                              |
| <b>SecondaryBVlanId</b>     | Indicates the secondary B-VLAN ID for this MAC address                                                                                                              |
| <b>SecondaryDestSysName</b> | Indicates the secondary system name of the node where the MAC address entry comes from.                                                                             |
| <b>SecondaryPort</b>        | Either displays the value 0, or indicates the secondary port on which a frame came from.                                                                            |
| <b>SmltRemote</b>           | Indicates the MAC address entry for the remote vIST peer.                                                                                                           |
| <b>Status</b>               | Indicates the status of this entry: <ul style="list-style-type: none"> <li>• other</li> <li>• invalid</li> <li>• learned</li> <li>• self</li> <li>• mgmt</li> </ul> |

### Configure UNI

#### Before You Begin

- You must enable Flex UNI to create a Switched UNI (S-UNI) service.
- To configure an untagged Switched UNI (S-UNI) on an EAP-enabled interface, you must associate the untagged S-UNI with a platform VLAN.

#### About This Task

Use the following procedure to configure a Transparent Port UNI or Switched UNI (S-UNI) by mapping an I-SID to a port or MLT and VLAN together.



#### Note

If you configure a T-UNI to terminate on a port or MLT on a switch in a vIST switch cluster, you must also configure the T-UNI I-SID on the other switch of the vIST switch cluster. You must configure the T-UNI I-SID on both switches of a vIST pair. You do not need to assign an actual port or MLT to the T-UNI on the second switch.

You must first create a type of service instance identifier (I-SID) to create the different types of services available. After you create an I-SID you can add members (ports or MLTs) to the I-SID to create end-points for the service.

#### Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. Select the **Service** tab.

4. To create a Transparent Port UNI service:
  - a. Select **Insert**.
  - b. Select **elan Transparent** in the **Type** field.
  - c. Enter the I-SID in the **Id** field.
5. To create an S-UNI service:
  - a. Select **Insert**.
  - b. Select **elan** in the **Type** field.
  - c. Enter the I-SID in the **Id** field.
6. Select **Insert**.

### Service Field Descriptions

Use the data in the following table to use the **Service** tab.

| Name                  | Description                                                                    |
|-----------------------|--------------------------------------------------------------------------------|
| <b>ID</b>             | Specifies a unique value to identify the service associated with this entry.   |
| <b>Type</b>           | Specifies the type of service associated with this entry.                      |
| <b>MacLimitEnable</b> | Indicates whether the MAC limit is enabled (true) or disabled (false).         |
| <b>MaxMacLimit</b>    | Indicates the maximum learned value of the MAC address for each service I-SID. |
| <b>Action</b>         | Specifies I-SID related actions.                                               |
| <b>OriginBitMap</b>   | Specifies the origin of the I-SID.                                             |
| <b>VnId</b>           | Identifies the VXLAN service associated with this I-SID.                       |
| <b>Name</b>           | Specifies the name of the I-SID.                                               |

*Associate a Port and MLT with an I-SID for ELAN Transparent*

### Before You Begin

- You must configure Transparent Port UNI. For more information, see [Configuring Transparent UNI](#).
- You must associate a T-UNI LACP MLT with a VLAN before mapping the LACP MLT to a T-UNI I-SID.



### Caution

Ensure that a T-UNI LACP MLT is always associated with a VLAN (even if it is the default VLAN) before adding it to a T-UNI I-SID. Otherwise, traffic is not forwarded on the T-UNI LACP MLT.

### About This Task

Use the following procedure to associate a port and MLT with an I-SID.

### Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. On the **Service** tab, select a row with the type configured as elanTransparent.
4. Select **ELAN**.
5. Select port members.
6. Select MLT IDs.
7. Select **Apply**.

### Service Field Descriptions

Use the data in the following table to use the **Service** tab.

| Name        | Description                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| PortMembers | The set of ports that are members of the ELAN Transparent service type. From the available ports, you can select a single or multiple ports.        |
| MltIds      | The set of bits that represent the MLT IDs. From the available MLTs, you can select any, or all of the MLTs to be a part of ELAN Transparent I-SID. |

*View the I-SID forwarding database*

View the I-SID forwarding database (FDB).



#### Note

To view the T-UNI I-SID FDB entries filtered on a port that is part of an MLT, you must mention the MLT ID in the option for the port.

### Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. Select the **FDB** tab.
4. (Optional) Select the **Filter** button to filter rows based on specific filter criteria.

### FDB Field Descriptions

Use data in the following table to use the **FDB** tab.

| Name     | Description                                                                                                                |
|----------|----------------------------------------------------------------------------------------------------------------------------|
| IsidId   | Specifies the service interface identifier (I-SID).                                                                        |
| Address  | Specifies the MAC address of the port assigned to the specific I-SID or CMAC learned on the particular I-SID.              |
| Status   | Specifies the learning status of the associated MAC.                                                                       |
| Port     | Specifies the port on which the MAC is learned for the specific I-SID.                                                     |
| PortType | Specifies if the MAC address is learned locally or on a network-to-network interface (NNI) port from a remote destination. |

| Name                        | Description                                                                 |
|-----------------------------|-----------------------------------------------------------------------------|
| <b>RemoteMacDestAddr</b>    | Specifies the virtual B-MAC address or system-ID of the remote destination. |
| <b>RemoteMacBVlanId</b>     | Specifies the B-VLAN ID on which the remote destination was discovered.     |
| <b>RemoteMacDestSysName</b> | Specifies the remote destination system name.                               |
| <b>Cvid</b>                 | Specifies the customer VLAN ID of the associated Switched UNI port.         |

*Associate a Port and MLT with an I-SID for ELAN*

### About This Task

Use the following procedure to associate a port and MLT with an I-SID.

You must configure Switched UNI.

### Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. On the **Service** tab, select a row with the type configured as elan.
4. Select **Switched Uni**.
5. Select **Insert**.
6. Enter the VLAN ID in the **Cvid** field.
7. Select **Port** or **Mlt** to update the interface index in the **IfIndex** field.
8. Select **Insert**.

### Switched Uni Field Descriptions

Use the data in the following table to use the **Switched Uni** tab.

| Name                | Description                                                                                      |
|---------------------|--------------------------------------------------------------------------------------------------|
| <b>Isid</b>         | Displays the I-SID.                                                                              |
| <b>Cvid</b>         | Specifies the customer VLAN identifier.                                                          |
| <b>IfIndex</b>      | Specifies the interface index of the ELAN end point.                                             |
| <b>Bpdu</b>         | Enables or disables for an untagged end point. The default is disabled.                          |
| <b>OriginBitmap</b> | Specifies the origin information of the service associated with the I-SID ELAN end point.        |
| <b>MacBased</b>     | Shows if the current entry is associated to a MAC-based Switched User Network Interface (S-UNI). |

*View the I-SID interface***Procedure**

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. Select the **Interface** tab.
4. (Optional) Select **Filter** to filter rows on specific filter criteria.

**Interface Field Descriptions**

Use the data in the following table to use the **Interface** tab.

| Name                | Description                                                              |
|---------------------|--------------------------------------------------------------------------|
| <b>IfIndex</b>      | Specifies the interface index                                            |
| <b>Isid</b>         | Specifies the service interface identifier (I-SID).                      |
| <b>Isid Name</b>    | Specifies the service interface identifier name.                         |
| <b>Vlan</b>         | Specifies the platform VLAN.                                             |
| <b>Cvid</b>         | Specifies the customer VID.                                              |
| <b>Type</b>         | Specifies the type of service associated with the I-SID interface.       |
| <b>OriginBitMap</b> | Specifies the origin of the service associated with the I-SID interface. |
| <b>Bpdu</b>         | Specifies the BPDU forward option for the untagged traffic port.         |
| <b>MacBased</b>     | Specifies the Switched UNI MAC address.                                  |

*Modify Global I-SID Name***About This Task**

Use this procedure to modify the assigned name for the Service Identifier (I-SID).

**Procedure**

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **ISID**.
3. Select **Global Name**.
4. Double-click the name of the I-SID in the **ISID Name** field, and then type a new name.
5. Select **Apply**.



## Global Name Field Descriptions

Use the data in the following table to use the **Global Name** tab.

| Name              | Description                                                                                                                                                                         |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ISID Id</b>    | Specifies the index number that uniquely identifies the I-SID.                                                                                                                      |
| <b>ISID Name</b>  | Specifies the name of the I-SID, which can be up to 64 characters.                                                                                                                  |
| <b>UsedByType</b> | Specifies the I-SIDs that are in use as services. An I-SID can have one base type or a combination of base types so that multiple services can use the same I-SID at the same time. |


## Layer 2 VSN Configuration Examples

This section provides configuration examples to configure Layer 2 VSNs.

Examples and network illustrations in this document may not include your specific platform. Unless otherwise noted, the concept illustrated applies to all supported platforms. For information about high-level feature support, see *Fabric Engine and VOSS Feature Matrix*.

### Layer 2 VSN Configuration Example

The following figure shows a sample Layer 2 VSN deployment.



**Figure 97: Layer 2 VSN**

The following sections show the steps required to configure the Layer 2 VSN parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see: [SPBM Configuration Examples](#) on page 1293.

### SwitchC

```
VLAN CONFIGURATION

vlan create 10 type port-mstprstp 1
vlan members 10 1/1 portmember
vlan i-sid 10 12990010
```

**SwitchD**

```
VLAN CONFIGURATION

vlan create 10 type port-mstprstp 1
vlan members 10 1/1 portmember
vlan i-sid 10 12990010
```

*Verify Layer 2 VSN operation*

The following sections show how to verify the Layer 2 VSN operation in this example.

**SwitchC**

```
Switch# show isis spbm i-sid all
=====
SPBM ISID INFO
=====
ISID SOURCE NAME VLAN SYSID TYPE HOST_NAME ISID NAME AREA AREA NAME

101001 1.11.16 4051 0200.10ff.fff0 discover SWITCHD ISID-101001 HOME area-0.00.20
101003 1.11.16 4051 0015.e89f.e3df config SWITCHC ISID-101003 HOME area-0.00.20

Total number of SPBM ISID entries configed: 1

Total number of SPBM ISID entries discovered: 1

Total number of SPBM ISID entries: 2

SwitchC:1# show isis spbm multicast-fib
=====
SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA ISID BVLAN SYSID HOST-NAME OUTGOING-INTERFACES INCOMING AREA AREA-NAME
 INTERFACE

f3:30:14:c6:36:3a 101001 4000 0200.10ff.fff0 SwitchD 1/1 MLT-1 HOME area-0.00.20
f3:30:13:c6:36:3a 101003 4000 0015.e89f.e3df SwitchC 1/30,1/1 MLT-1 HOME area-0.00.20

```

**SwitchD**

```
SwitchD:1# show isis spbm i-sid all
=====
SPBM ISID INFO
=====
ISID SOURCE NAME VLAN SYSID TYPE HOST_NAME ISID NAME AREA AREA NAME

12990010 f.30.14 4000 0014.0da0.13df config SwitchD ISID-12990010 HOME area-0.00.20
12990010 f.30.13 4000 0015.e89f.e3df discover SwitchC ISID-12990010 HOME area-0.00.20

SwitchD:1# show isis spbm multicast-fib
=====
SPBM MULTICAST FIB ENTRY INFO
=====
MCAST DA ISID BVLAN SYSID HOST-NAME OUTGOING-INTERFACES INCOMING AREA AREA-NAME
 INTERFACE

f3:30:14:c6:36:3a 12990010 4000 0014.0da0.13df SwitchD MLT-1,1/1 1/30 HOME area-0.00.20
f3:30:13:c6:36:3a 12990010 4000 0015.e89f.e3df SwitchC 1/1 1/30 HOME area-0.00.20

```

**SwitchC — verifying with CFM**

```
SwitchC:1# 12 tracetree 4000 12990010

Please wait for 12tracetree to complete or press any key to abort

12tracetree to f3:30:13:c6:36:3a, vlan 4000 i-sid 12990010 nickname f.30.13 hops 64
1 SwitchC 00:15:e8:9f:e3:df -> SwitchG 00:0e:62:25:a3:df
2 SwitchG 00:0e:62:25:a3:df -> SwitchD 00:14:0d:a0:13:df
```

**SwitchD — verifying with CFM**

```
SwitchD:1# 12 tracetree 4000 12990010

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to f3:30:14:c6:36:3a, vlan 4000 i-sid 12990010 nickname f.30.14 hops 64
1 SwitchD 00:14:0d:a0:13:df -> SwitchG 00:0e:62:25:a3:df
2 SwitchG 00:0e:62:25:a3:df -> SwitchC 00:15:e8:9f:e3:df
```

**SwitchC — verifying FDB**

```
SwitchC:1# show vlan mac-address-entry 10
=====
Vlan Fdb
=====
VLAN MAC
ID STATUS ADDRESS INTERFACE SMLT
REMOTE TUNNEL

10 learned 00:00:00:00:00:01 Port-1/1 false SwitchD
10 learned 00:00:00:00:00:02 Port-1/1 false SwitchD

2 out of 4 entries in all fdb(s) displayed.

SwitchC:1# show vlan remote-mac-table 10
=====
Vlan Remote Mac Table
=====
VLAN STATUS MAC-ADDRESS DEST-MAC BVLAN DEST-SYSNAME PORTS

10 learned 00:00:00:00:00:02 00:14:0d:a0:13:df 0014.0da0.13df SwitchD 1/30

Total number of VLAN Remote MAC entries 1

```

**SwitchD — verifying FDB**

```
SwitchD:1# show vlan mac-address-entry 10
=====
Vlan Fdb
=====
VLAN MAC
ID STATUS ADDRESS INTERFACE SMLT
REMOTE TUNNEL

10 learned 00:00:00:00:00:01 Port-1/1 false SwitchC
10 learned 00:00:00:00:00:02 Port-1/1 false SwitchC

2 out of 4 entries in all fdb(s) displayed.

SwitchD:1# show vlan remote-mac-table 10
=====
Vlan Remote Mac Table
=====
VLAN STATUS MAC-ADDRESS DEST-MAC DEST-SYSID DEST-SYSNAME PORTS


10 learned 00:00:00:00:00:01 00:15:e8:9f:e3:df 0015.e89f.e3df SwitchC MLT-1

Total number of VLAN Remote MAC entries 1

```

### Layer 2 VSN Example with VLAN ID Translation

The following figure shows a sample Layer 2 VSN deployment where the C- VLAN IDs are different at each end. You must first configure basic SPBM and IS-IS infrastructure. For more information, see [SPBM Configuration Examples](#) on page 1293.



**Figure 98: Layer 2 VSN with different VLAN IDs**

The following sections show the steps required to configure the Layer 2 VSN parameters in this example.

#### SwitchC

```
VLAN CONFIGURATION

vlan create 9 type port 1
vlan members 9 1/1 portmember
vlan i-sid 9 9
```

#### SwitchD

```
VLAN CONFIGURATION

vlan create 19 type port 1
vlan members 19 1/1 portmember
vlan i-sid 19 9
```


## Inter-VSN Routing Configuration

**Table 141: Inter-VSN Routing product support**

| Feature                  | Product         | Release introduced |
|--------------------------|-----------------|--------------------|
| Inter-VSN routing (IPv4) | 4220 Series     | Fabric Engine 9.2  |
|                          | 5320 Series     | Fabric Engine 8.6  |
|                          | 5420 Series     | VOSS 8.4           |
|                          | 5520 Series     | VOSS 8.2.5         |
|                          | 5720 Series     | Fabric Engine 8.7  |
|                          | 7520 Series     | Fabric Engine 8.10 |
|                          | 7720 Series     | Fabric Engine 8.10 |
|                          | 7830 Series     | Fabric Engine 9.3  |
|                          | VSP 4900 Series | VOSS 8.1           |
|                          | VSP 7400 Series | VOSS 8.0           |
| Inter-VSN routing (IPv6) | 4220 Series     | Not Supported      |
|                          | 5320 Series     | Fabric Engine 8.6  |
|                          | 5420 Series     | VOSS 8.4           |
|                          | 5520 Series     | VOSS 8.2.5         |
|                          | 5720 Series     | Fabric Engine 8.7  |
|                          | 7520 Series     | Fabric Engine 8.10 |
|                          | 7720 Series     | Fabric Engine 8.10 |
|                          | 7830 Series     | Not Supported      |
|                          | VSP 4900 Series | VOSS 8.1           |
|                          | VSP 7400 Series | VOSS 8.0           |

## Inter-VSN Routing Fundamentals

Inter-VSN routing with SPBM allows routing between Layer 2 VLANs with different I-SIDs.



**Figure 99: Inter-VSN routing**

Inter-VSN routing provides a routing hub for Layer 2 Virtual Services Network edge devices, Layer 3 devices, routers, or hosts connected to the SPBM cloud using the SPBM Layer 2 VSN service. To go between a routed network, a Layer 2 VSN termination point provides the routing services to hop onto another Layer 2 VSN, using I-SID.



**Note**

The Layer 2 VLANs must be in the same VRF. You cannot route traffic between two different VRFs with Inter-VSN routing.

In this example, the C-VLANs are associated with I-SIDs on the BEBs using SPBM Layer 2 VSN. With Inter-VSN routing enabled, BCB C can route traffic between VLAN 11 (I-SID 2011) and VLAN 12 (I-SID 2012).

IP interfaces are where the routing instance exists. In this case, on Switch-20.



**Note**

The switch does not support IP multicast over Fabric Connect routing on inter-VSN routing interfaces.

## Inter-VSN routing configuration using the CLI

This section provides a procedure to configure Inter-VSN routing using the CLI.

### Configure SPBM Inter-VSN Routing

#### Before You Begin

- You must configure the required SPBM and IS-IS infrastructure.

#### About This Task

Configure Inter-VSN routing to route between IP networks on Layer 2 VLANs with different I-SIDs. Inter-VSN routing is typically used only when you have to extend a VLAN as a Layer 2 Virtual Services Network (VSN) for applications such as vMotion. As a best practice, use IP Shortcuts or Layer 3 VSNs to route traffic. You must configure both the Backbone Edge Bridges (BEBs) and the Backbone Core Bridge (BCB). This procedure includes the steps for both roles.

To enable inter-VSN routing, you must configure an IP interface where the routing instance exists.

### Procedure

1. Enter Global Configuration mode:
 

```
enable
```

```
configure terminal
```
2. Complete the following substeps on the Backbone Edge Bridges (BEBs) that contain the VSNs you want to route traffic between.

- a. Create a customer VLAN (C-VLAN) by port:

```
vlan create <2-4059> type port-mstprstp <0-63>
```

- b. Add ports in the C-VLAN:

```
vlan members add <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
```

- c. Map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID):

```
vlan i-sid <1-4059> <0-16777215> [force]
```



#### Important

When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

3. On the Backbone Core Bridge (BCB), create a VRF and add a VLAN for each VSN:
  - a. Create a VRF:

```
ip vrf WORD<1-16> vrfid <1-511>
```



#### Note

The supported VRF ID range on the 5320-16P-2MXT-2X differs from other Fabric Engine platforms. 5320-16P-2MXT-2X supports a range of 1-14. To verify the supported range on a specific model, use the CLI contextual help (?).

- b. Create a VLAN to associate with each VSN:

```
vlan create <2-4059> type port-mstprstp <0-63>
```

- c. Enter VLAN Interface Configuration mode:

```
interface vlan <1-4059>
```

- d. Add a VLAN to the VRF you created in step a:

```
vrf WORD<1-16>
```

- e. Associate an I-SID with the VLAN:

```
vlan i-sid <1-4059> <0-16777215> [force]
```



### Important

When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

The switch reserves I-SID 0x00ffffff. The switch uses this I-SID to advertise the virtual BMAC in an SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.

- f. Configure an IP address for the VLAN:

```
ip address {A.B.C.D/X}
```

- g. Repeat steps b to f for every VLAN you want to route traffic between.

### Variable Definitions

The following table defines parameters for the **vlan create** command.

| Variable                                                   | Value                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <2-4059>                                                   | Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. |
| <i>type port-mstprstp</i> <0-63><br><i>[color</i> <0-32> ] | Creates a VLAN by port: <ul style="list-style-type: none"> <li>• &lt;0-63&gt; is the STP instance ID.</li> <li>• <i>color</i> &lt;0-32&gt; is the color of the VLAN.</li> </ul>                                                                                                                                                                                                    |



The following table defines parameters for the **vlan members add** command.

| Variable                                                             | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <1-4059>                                                             | Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.                                                                                                         |
| {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]} | Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> <li>• a single slot and port (slot/port)</li> <li>• a range of slots and ports (slot/port-slot/port)</li> <li>• a series of slots and ports (slot/port,slot/port,slot/port)</li> <li>• all ports on the same slot (slot/all)</li> <li>• all ports on the switch (all)</li> </ul> <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> |

The following table defines parameters for the **vlan i-sid** command.

| Variable     | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <1-4059>     | Specifies the primary VLAN ID.<br>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. |
| <0-16777215> | Specifies the service instance identifier (I-SID). You cannot use I-SID 0x00ffffff. The system reserves this I-SID to advertise the virtual BMAC in an SMLT dual-homing environment.<br>This value is the same for the primary and secondary VLANs.<br><br><b>Note:</b><br>When you add or remove an I-SID on a VLAN IP interface for 5520 Series/5720 Series, the switch disables and enables routing on the VLAN interface.                                |
| <i>force</i> | Specifies the software must replace the existing VLAN-to-I-SID mapping, if one exists.                                                                                                                                                                                                                                                                                                                                                                       |

The following table defines parameters for the **ip vrf** command.

| Variable             | Value                                                    |
|----------------------|----------------------------------------------------------|
| <i>WORD</i> <1-16>   | Create the VRF and specify the name of the VRF instance. |
| <i>vrfid</i> <1-511> | Specifies the VRF instance by number.                    |

The following table defines parameters for the **vrf** command.

| Variable           | Value                                               |
|--------------------|-----------------------------------------------------|
| <i>WORD</i> <1-16> | Specifies the VRF name. Associates a port to a VRF. |

The following table defines parameters for the **ip address** command.

| Variable    | Value                                  |
|-------------|----------------------------------------|
| {A.B.C.D/X} | Configures an IP address for the VLAN. |

## Inter-VSN routing configuration using EDM

This section provides procedures to configure Inter-VSN routing using Enterprise Device Manager (EDM).

### Configure BEBs for Inter-VSN Routing

#### Before You Begin

You must configure the required SPBM and IS-IS infrastructure.

#### About This Task

Use Inter-VSN routing to route between IP networks on Layer 2 VLANs with different I-SIDs. Inter-VSN routing is typically used only to extend a VLAN as a Layer 2 Virtual Services Network (VSN) for applications such as vMotion. Use IP Shortcuts or Layer 3 VSNs to route traffic. You must configure both the Backbone Edge Bridges (BEBs) and the Backbone Core Bridge (BCB).

Follow these steps on the BEBs that contain the VSNs you want to route traffic between. To configure the BCB, see [Configure BCBs for Inter-VSN Routing](#) on page 1565.



#### Note

To enable inter-VSN routing, you must configure the IP interface where the routing instance exists.

## Procedure

1. Create a customer VLAN (C-VLAN) by port and add ports in the C-VLAN:
  - a. In the navigation pane, expand **Configuration > VLAN**.
  - b. Select **VLANs**.
  - c. On the **Basic** tab, select **Insert**.
  - d. For **Id**, type an unused VLAN ID, or use the ID provided.
  - e. For **Name**, type the VLAN name, or use the name provided.
  - f. For **Color Identifier**, select a color from the list, or use the color provided.
  - g. For **Type**, select **byPort**.
  - h. For **PortMembers**, select the ellipsis (...).
  - i. Select the ports to add as member ports.  
 The ports that are selected are recessed, while the nonselected ports are not recessed. Port numbers that are dimmed cannot be selected as VLAN port members.
  - j. Select **OK**.
  - k. Select **Insert**.
2. Map a C-VLAN to an I-SID:
  - a. From the same **Configuration > VLAN > VLANs** navigation path, select the **Advanced** tab.
  - b. For **Isid**, specify the I-SID to associate with the specified VLAN.  
 The switch reserves I-SID 0x00ffffff. The switch uses this I-SID to advertise the virtual BMAC in a SMLT dual-homing environment. The platform clears the receive and transmit bit of this I-SID, therefore I-SID 0x00ffffff cannot be used for any other service.
  - c. Select **Apply**.



### Important

When a protocol VLAN is created, all ports are added to the VLAN including SPBM ports. To configure a protocol-based VLAN as a C-VLAN, you must first remove the SPBM-enabled ports from the protocol based VLAN, and then configure the protocol-based VLAN as a C-VLAN.

3. Configure the Backbone Core Bridge (BCB) for Inter-VSN Routing. For more information, see [Configure BCBs for Inter-VSN Routing](#) on page 1565.

## Configure BCBs for Inter-VSN Routing

Inter-VSN allows you to route between IP networks on Layer 2 VLANs with different I-SIDs. Inter-VSN routing is typically used only when you have to extend a VLAN as a Layer 2 Virtual Services Network (VSN) for applications such as vMotion. Use IP Shortcuts to route traffic. You must configure both the Backbone Edge Bridges (BEBs) and the Backbone Core Bridge (BCB).



### Note

To enable inter-VSN routing, you must configure the IP interface where the routing instance exists.

## Before You Begin

- You must configure the required SPBM and IS-IS infrastructure.
- You must configure the BEBs that contain the VSNs you want to route traffic between. For more information, see [Configure BEBs for Inter-VSN Routing](#) on page 1564.

## About This Task

Follow these steps to configure the BCB for inter-VSN routing.

## Procedure

1. On the BCB, create a VRF:
  - a. In the navigation pane, expand **Configuration > VRF**.
  - b. Select **VRF**.
  - c. Select **Insert**.
  - d. For **Id**, specify the VRF ID.
  - e. Name the VRF instance.
  - f. Configure the other parameters as required.
  - g. Select **Insert**.
2. Create a VLAN to associate with each VSN:
  - a. In the navigation pane, expand **Configuration > VLAN**.
  - b. Select **VLANs**.
  - c. On the **Basic** tab, select **Insert**.
  - d. For **Id**, type an unused VLAN ID, or use the ID provided.
  - e. For **Name**, type the VLAN name, or use the name provided.
  - f. For **Color Identifier**, select a color from the list, or use the color provided.
  - g. For **Type**, select **byPort**.
  - h. For **PortMembers**, select the ellipsis (...).
  - i. Select the ports to add as member ports.

The ports that are selected are recessed, while the nonselected ports are not recessed. Port numbers that are dimmed cannot be selected as VLAN port members.
  - j. Select **OK**.
  - k. Select **Insert**.
3. Associate the VLAN with an I-SID:
  - a. From the same **Configuration > VLAN > VLANs** navigation path, select the **Advanced** tab.
  - b. For **Isid**, specify the I-SID to associate with the specified VLAN.
  - c. Select **Apply**.
4. Configure a circuitless IP interface (CLIP):
  - a. In the navigation pane, expand **Configuration > IP**.
  - b. Select **IP**.
  - c. Select the **Circuitless IP** tab.

- d. Select **Insert**.
- e. For **Interface**, type a CLIP interface number.
- f. Provide the IP address.
- g. Provide the network mask.
- h. Select **Insert**.


## Inter-VSN Routing Configuration Example

The following topics provide a configuration example for Inter-VSN routing.

Examples and network illustrations in this document may not include your specific platform. Unless otherwise noted, the concept illustrated applies to all supported platforms. For information about high-level feature support, see *Fabric Engine and VOSS Feature Matrix*.

### Inter-VSN Routing with SPBM Configuration Example

The following figure shows a sample Inter-VSN deployment.



**Figure 100: Inter-VSN routing**

The following sections show the steps required to configure the Inter-VSN parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see: [SPBM Configuration Examples](#) on page 1293.

Note that the IP interfaces are configured where the routing instance exists, namely, on Switch-20 (BCB C).

#### BEB A

```
VLAN CONFIGURATION

vlan create 11 type port-mstprstp 1
vlan members 11 1/2 portmember
vlan i-sid 11 2011
```

#### BCB C

```
VRF CONFIGURATION

ip vrf blue vrfid 100
```

## VLAN CONFIGURATION

```
vlan create 11 type port-mstprstp 1
vlan i-sid 11 2011
interface Vlan 11
vrf blue
ip address 192.0.2.1 255.255.255.0
exit
```

## VLAN CONFIGURATION

```
vlan create 12 type port-mstprstp 1
vlan i-sid 12 2012
interface Vlan 12
vrf blue
ip address 198.51.100.1 255.255.255.0
exit
```

**BEB D**

## VLAN CONFIGURATION

```
vlan create 12 type port-mstprstp 1
vlan members 12 1/2 portmember
vlan i-sid 12 2012
```

*Verifying Inter-VSN Routing Operation*

The following sections show how to verify Inter-VSN Routing operation in this example.

**BCBC**

```
Switch:1# show ip route vrf blue
```

```
=====
 IP Route - VRF blue
=====
```

| DST          | MASK          | NEXT         | NH<br>VRF/ISID | COST | INTER<br>FACE | PROT | AGE | TYPE | PRF | ACCESS  |
|--------------|---------------|--------------|----------------|------|---------------|------|-----|------|-----|---------|
| 192.0.2.0    | 255.255.255.0 | 192.0.2.1    | -              | 1    | 11            | LOC  | 0   | DB   | 0   | TRUSTED |
| 198.51.100.0 | 255.255.255.0 | 198.51.100.1 | -              | 1    | 12            | LOC  | 0   | DB   | 0   | TRUSTED |

```
Switch:1# show ip arp vrf blue
```

```
=====
 IP Arp - VRF blue
=====
```

| IP_ADDRESS     | MAC_ADDRESS       | VLAN | PORT | TYPE  | TTL(10 Sec) | TUNNEL | L3 | ACCESS       |
|----------------|-------------------|------|------|-------|-------------|--------|----|--------------|
| 192.0.2.1      | 00:0e:62:25:a2:00 | 11   | -    | LOCAL | 2160        |        |    | GlobalRouter |
| 192.0.2.255    | ff:ff:ff:ff:ff:ff | 11   | -    | LOCAL | 2160        |        |    | GlobalRouter |
| 198.51.100.1   | 00:0e:62:25:a2:01 | 12   | -    | LOCAL | 2160        |        |    | GlobalRouter |
| 198.51.100.255 | ff:ff:ff:ff:ff:ff | 12   | -    | LOCAL | 2160        |        |    | GlobalRouter |

```
=====
 IP Arp Extn - VRF blue
=====
```

| MULTICAST-MAC-FLOODING | AGING (Minutes) | ARP-THRESHOLD |
|------------------------|-----------------|---------------|
| disable                | 360             | 500           |

```
4 out of 50 ARP entries displayed
```

**BCB C**

```
Switch:1# show vlan mac-address-entry 11
=====
 Vlan Fdb
=====
VLAN MAC
ID STATUS ADDRESS INTERFACE SMLT
REMOTE TUNNEL

11 learned 00:00:00:00:01:02 Port-1/2 false
11 self 00:0e:62:25:a2:00 Port-cpp false
Switch-10
-

2 out of 4 entries in all fdb(s) displayed.

Switch:1# show vlan mac-address-entry 12
=====
 Vlan Fdb
=====
VLAN MAC
ID STATUS ADDRESS INTERFACE SMLT
REMOTE TUNNEL

12 learned 00:00:00:00:02:02 Port-1/2 false
12 self 00:0e:62:25:a2:01 Port-cpp false
Switch-30
-

2 out of 4 entries in all fdb(s) displayed.
```

**BEB A**

```
Switch:1# show vlan mac-address-entry 11
=====
 Vlan Fdb
=====
VLAN MAC
ID STATUS ADDRESS INTERFACE SMLT
REMOTE TUNNEL

11 learned 00:00:00:00:01:02 Port-1/2 false
11 learned 00:0e:62:25:a2:00 Port-1/2 false
Switch-30
Switch-30

2 out of 2 entries in all fdb(s) displayed.
```

**BEB D**

```
Switch:1# show vlan mac-address-entry 12
=====
 Vlan Fdb
=====
VLAN MAC
ID STATUS ADDRESS INTERFACE SMLT
REMOTE TUNNEL


12 learned 00:00:00:00:02:02 Port-1/2 false
12 learned 00:0e:62:25:a2:01 Port-1/2 false
Switch-10
Switch-10

2 out of 2 entries in all fdb(s) displayed.
```

## SBPM Reference Architectures

SBPM has a straightforward architecture that simply forwards encapsulated CMACs across the backbone. Because the BMAC header stays the same across the network, there is no need to swap a label or perform a route lookup at each node. This architecture allows the frame to follow the most efficient forwarding path from end to end.

The following figure shows the MAC-in-MAC SPBM domain with BEBs on the boundary and BCBs in the core.




**Figure 101: SPBM basic architecture**

Provisioning an SPBM core is as simple as enabling SPBM and IS-IS globally on all the nodes and on the core facing links. To migrate an existing edge configuration into an SPBM network is just as simple.

The boundary between the MAC-in-MAC SPBM domain and the 802.1Q domain is handled by the BEBs. At the BEBs, VLANs or VRFs are mapped into I-SIDs based on the local service provisioning. Services (whether Layer 2 or Layer 3 VSNs) only need to be configured at the edge of the SPBM backbone (on the BEBs). There is no provisioning needed on the core SPBM nodes.

The following figure illustrates an existing edge that connects to an SPBM core.





**Figure 102: Access to the SPBM Core**

All BEBs that have the same I-SID configured can participate in the same VSN. That completes the configuration part of the migration and all the traffic flows return to normal operation.

For Layer 3 virtualized routing (Layer 3 VSN), map IPv4-enabled VLANs to VRFs, create an IP VPN instance on the VRF, assign an I-SID to the VRF, and then configure the desired IP redistribution of IP routes into IS-IS.


For Layer 2 virtualized bridging (Layer 2 VSN), identify all the VLANs that you want to migrate into SPBM and assign them to an I-SID on the BEB.

## Campus Architecture

For migration purposes, you can add SPBM to an existing network that has SMLT configured. In fact, if there are other protocols already running in the network, such as Open Shortest Path First (OSPF), you can leave them in place too. SPBM uses IS-IS, and operates independently from other protocols. However, as a best practice, eliminate SMLT in the core and eliminate other unnecessary protocols. This reduces the complexity of the network and makes it much simpler to maintain and troubleshoot.

Whether you configure SMLT in the core, the main point to remember is that SPBM separates services from the infrastructure. For example, in a large campus, a user may need access to other sites or data centers. With SPBM you can grant that access by associating the user to a specific I-SID. With this mechanism, the user can work without getting access to confidential information of another department.

The following figure depicts a topology where the BEBs in the edge and data center distribution nodes are configured in SMLT clusters. Prior to implementing SPBM, the core nodes would also have been configured as SMLT clusters. When migrating SPBM onto this network design, it is important to note that you can deploy SPBM over the existing SMLT topology without network interruption. After the SPBM infrastructure is in place, you can create VSN services over SPBM or migrate them from the previous end-to-end SMLT-based design.



**Figure 103: SPBM campus without SMLT**

After you migrate all services to SPBM, the customer VLANs (C-VLANs) will exist only on the BEB SMLT clusters at the edge of the SPBM network. The C-VLANs will be assigned to an I-SID instance and then associated with either a VLAN in an Layer 2 VSN or terminated into a VRF in an Layer 3 VSN. You can also terminate the C-VLAN into the default router, which uses IP shortcuts to IP route over the SPBM core.

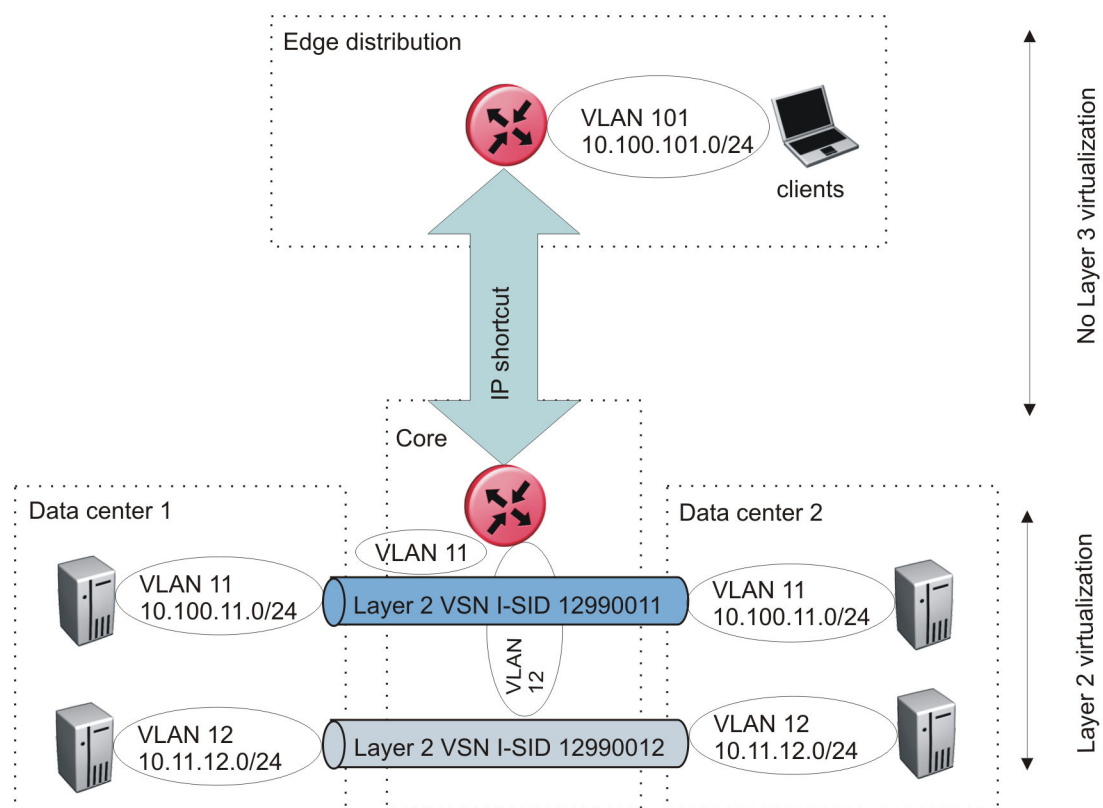
In an SPBM network design, the only nodes where it makes sense to have an SMLT cluster configuration is on the BEB nodes where VSN services terminate. These are the SPBM nodes where C-VLANs exist and these C-VLANs need to be redundantly extended to non-SPBM devices such as Layer 2 edge stackable switches. On the BCB core nodes where no VSNs are terminated and no Layer 2 edge stackables are connected, there is no longer any use for the SMLT clustering functionality. Therefore, in the depicted SPBM design, the SMLT/VIST configuration can be removed from the core nodes because they now act as pure BCBs that simply transport VSN traffic and the only control plane protocol they need to run is IS-IS.

Because SMLT BEB nodes exist in this design (the edge BEBs) and it is desirable to use equal cost paths to load balance VSN traffic across the SPBM core, all SPBM nodes in the network are configured with the same two B-VIDs.

Where the previous figure shows the physical topology, the following two figures illustrate a logical rendition of the same topology. In both of the following figures, you can see that the core is almost identical. Because the SPBM core just serves as a transport mechanism that transmits traffic to the destination BEB, all the provisioning is performed at the edge.

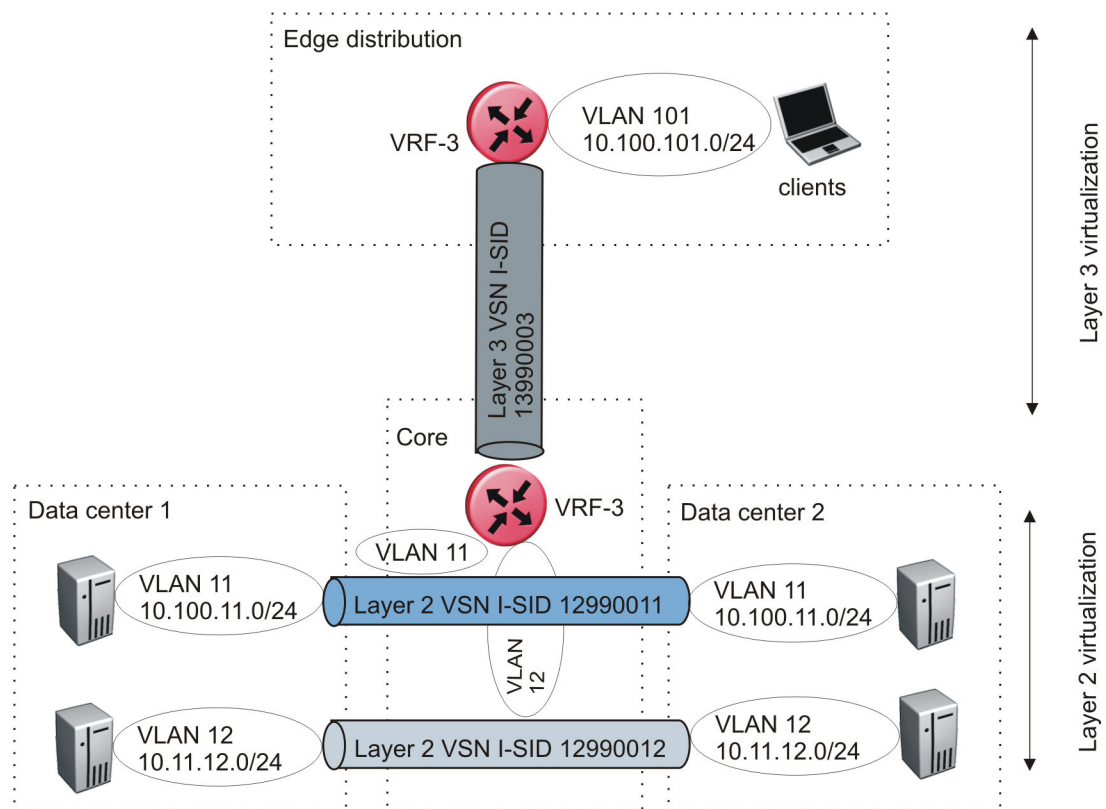
In the data center, VLANs are attached to Inter-VSNs that transmit the traffic across the SPBM core between the data center on the left and the data center on the right. A common application of this service is VMotion moving VMs from one data center to another.

The following figure uses IP shortcuts that route VLANs. There is no I-SID configuration and no Layer 3 virtualization between the edge distribution and the core. This is normal IP forwarding to the BEB.



**Figure 104: IP shortcut scenario to move traffic between data centers**

The following figure uses Layer 3 VSNs to route VRFs between the edge distribution and the core. The VRFs are attached to I-SIDs and use Layer 3 virtualization.



**Figure 105: VRF scenario to move traffic between data centers**

## Large data center architecture


SPBM supports data centers with IP shortcuts, Layer 2 VSNs, or Layer 3 VSNs. If you use vMotion, you must use Layer 2 between data centers (Layer 2 VSN). With Layer 2 VSNs, you can add IP addresses to the VLAN on both data centers and run Virtual Router Redundancy Protocol (VRRP) between them to allow the ESX server to route to the rest of the network.

The following figure shows an SPBM topology of a large data center. This figure represents a full-mesh data center fabric using SPBM for storage over Ethernet. This topology is optimized for storage transport because traffic never travels more than two hops.



### Note

As a best practice, use a two-tier, full-mesh topology for large data centers.




**Figure 106: SPBM data center—full mesh**

#### *Traditional data center routing of VMs*

In a traditional data center configuration, the traffic flows into the network to a VM and out of the network in almost a direct path.

The following figure shows an example of a traditional data center with VRRP configured. Because end stations are often configured with a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks. VRRP eliminates the single point of failure that can occur when the single static default gateway router for an end station is lost.






**Figure 107: Traditional routing before moving VMs**

A VM is a virtual server. When you move a VM, the virtual server is moved as is. This action means that the IP addresses of that server remain the same after the server is moved from one data center to the other. This in turn dictates that the same IP subnet (and hence VLAN) exist in both data centers.

In the following figure, the VM moved from the data center on the left to the data center on the right. To ensure a seamless transition that is transparent to the user, the VM retains its network connections through the default gateway. This method works, but it adds more hops to all traffic. As you can see in the figure, one VM move results in a complicated traffic path. Multiply this with many moves and soon the network look like a tangled mess that is very inefficient, difficult to maintain, and almost impossible to troubleshoot.



**Figure 108: Traditional routing after moving VMs**

#### *Optimized data center routing of VMs*

Two features make a data center optimized:


- VLAN routers in the Layer 2 domain (green icons)
- VRRP BackupMaster

The VLAN routers use lookup tables to determine the best path to route incoming traffic (red dots) to the destination VM.

VRRP BackupMaster solves the problem of traffic congestion on the vIST. Because there can be only one VRRP Master, all other interfaces are in backup mode. In this case, all traffic is forwarded over the vIST link towards the primary VRRP switch. All traffic that arrives at the VRRP backup interface is forwarded, so there is not enough bandwidth on the vIST link to carry all the aggregated riser traffic. VRRP BackupMaster overcomes this issue by ensuring that the vIST trunk is not used in such a case for primary data forwarding. The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. All traffic is directly routed to the destined subnetwork and not through Layer 2 switches to the VRRP Master. This avoids potential limitation in the available vIST bandwidth.

The following figure shows a solution that optimizes your network for bidirectional traffic flows. However, this solution turns two SPBM BCB nodes into BEBs where MAC and ARP learning will be enabled on the Inter-VSN routing interfaces. If you do not


care about top-down traffic flows, you can omit the Inter-VSN routing interfaces on the SPBM BCB nodes. This makes the IP routed paths top-down less optimal, but the BCBs remain pure BCBs, thus simplifying core switch configurations.



**Figure 109: Optimized routing before moving VMs**

In the traditional data center, chaos resulted after many VMs were moved. In an optimized data center as shown in the following figure, the incoming traffic enters the Layer 2 domain where an edge switch uses Inter-VSN routing to attach an I-SID to a VLAN. The I-SID bridges traffic directly to the destination. With VRRP BackupMaster, the traffic no longer goes through the default gateway; it takes the most direct route in and out of the network.





**Figure 110: Optimized routing after moving VMs**



# Fabric Layer 3 Services

[IP Shortcuts Configuration](#) on page 1580

[Layer 3 VSN Configuration](#) on page 1654

The topics in this section provide conceptual and procedural information about Layer 3 Services:

- IP Shortcuts
- Layer 3 VSN

## IP Shortcuts Configuration

**Table 142: IP Shortcuts product support**

| Feature                            | Product         | Release introduced                                                                                                                                                                                                                                                                           |
|------------------------------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Shortcut routing including ECMP | 4220 Series     | Fabric Engine 9.2<br>The 4220 Series switches support only a single active VRF with IP configuration.                                                                                                                                                                                        |
|                                    | 5320 Series     | Fabric Engine 8.6<br>For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"><li>• 5320-16P-2MXT-2X</li><li>• 5320-24T-4X-XT</li><li>• 5320-24T-24S-4XE-XT</li><li>• 5320-48P-8XE</li><li>• 5320-48T-8XE</li></ul> |
|                                    | 5420 Series     | VOSS 8.4                                                                                                                                                                                                                                                                                     |
|                                    | 5520 Series     | VOSS 8.2.5                                                                                                                                                                                                                                                                                   |
|                                    | 5720 Series     | Fabric Engine 8.7                                                                                                                                                                                                                                                                            |
|                                    | 7520 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                           |
|                                    | 7720 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                           |
|                                    | 7830 Series     | Fabric Engine 9.3                                                                                                                                                                                                                                                                            |
|                                    | VSP 4900 Series | VOSS 8.1                                                                                                                                                                                                                                                                                     |
|                                    | VSP 7400 Series | VOSS 8.0                                                                                                                                                                                                                                                                                     |

**Table 142: IP Shortcuts product support (continued)**

| Feature               | Product         | Release introduced                                                                                                                                                                                                                                                                                 |
|-----------------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv6 Shortcut routing | 4220 Series     | Not Supported                                                                                                                                                                                                                                                                                      |
|                       | 5320 Series     | Fabric Engine 8.6<br>For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> <li>• 5320-16P-2MXT-2X</li> <li>• 5320-24T-4X-XT</li> <li>• 5320-24T-24S-4XE-XT</li> <li>• 5320-48P-8XE</li> <li>• 5320-48T-8XE</li> </ul> |
|                       | 5420 Series     | VOSS 8.4                                                                                                                                                                                                                                                                                           |
|                       | 5520 Series     | VOSS 8.2.5                                                                                                                                                                                                                                                                                         |
|                       | 5720 Series     | Fabric Engine 8.7                                                                                                                                                                                                                                                                                  |
|                       | 7520 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|                       | 7720 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|                       | 7830 Series     | Fabric Engine 9.4                                                                                                                                                                                                                                                                                  |
|                       | VSP 4900 Series | VOSS 8.1                                                                                                                                                                                                                                                                                           |
|                       | VSP 7400 Series | VOSS 8.0                                                                                                                                                                                                                                                                                           |

**Table 142: IP Shortcuts product support (continued)**

| Feature                    | Product         | Release introduced                                                                                                                                                                                                                                                                                 |
|----------------------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv4 IS-IS accept policies | 4220 Series     | Fabric Engine 9.2<br>The 4220 Series switches support only a single active VRF with IP configuration.                                                                                                                                                                                              |
|                            | 5320 Series     | Fabric Engine 8.6<br>For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> <li>• 5320-16P-2MXT-2X</li> <li>• 5320-24T-4X-XT</li> <li>• 5320-24T-24S-4XE-XT</li> <li>• 5320-48P-8XE</li> <li>• 5320-48T-8XE</li> </ul> |
|                            | 5420 Series     | VOSS 8.4                                                                                                                                                                                                                                                                                           |
|                            | 5520 Series     | VOSS 8.2.5                                                                                                                                                                                                                                                                                         |
|                            | 5720 Series     | Fabric Engine 8.7                                                                                                                                                                                                                                                                                  |
|                            | 7520 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|                            | 7720 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|                            | 7830 Series     | Fabric Engine 9.4                                                                                                                                                                                                                                                                                  |
|                            | VSP 4900 Series | VOSS 8.1                                                                                                                                                                                                                                                                                           |
|                            | VSP 7400 Series | VOSS 8.0                                                                                                                                                                                                                                                                                           |

**Table 142: IP Shortcuts product support (continued)**

| Feature                    | Product         | Release introduced                                                                                                                                                                                                                                                                                 |
|----------------------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPv6 IS-IS accept policies | 4220 Series     | Not Supported                                                                                                                                                                                                                                                                                      |
|                            | 5320 Series     | Fabric Engine 8.6<br>For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> <li>• 5320-16P-2MXT-2X</li> <li>• 5320-24T-4X-XT</li> <li>• 5320-24T-24S-4XE-XT</li> <li>• 5320-48P-8XE</li> <li>• 5320-48T-8XE</li> </ul> |
|                            | 5420 Series     | VOSS 8.4                                                                                                                                                                                                                                                                                           |
|                            | 5520 Series     | VOSS 8.2.5                                                                                                                                                                                                                                                                                         |
|                            | 5720 Series     | Fabric Engine 8.7                                                                                                                                                                                                                                                                                  |
|                            | 7520 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|                            | 7720 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|                            | 7830 Series     | Not Supported                                                                                                                                                                                                                                                                                      |
|                            | VSP 4900 Series | VOSS 8.1                                                                                                                                                                                                                                                                                           |
|                            | VSP 7400 Series | VOSS 8.0                                                                                                                                                                                                                                                                                           |

## IP Shortcuts configuration fundamentals

This section provides fundamental concepts for IP Shortcuts.


Fabric Connect supports both IPv4 Shortcuts and IPv6 Shortcuts. Because IPv6 Shortcuts depend on IPv4 Shortcuts, you should understand how IPv4 Shortcuts work (see [SPBM IP shortcuts](#)) before jumping to the IPv6 section.

### *SPBM IP Shortcuts*

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

Unlike Layer 2 VSN, with SPBM IP shortcuts, no I-SID configuration is required. Instead, SPBM nodes propagate Layer 3 reachability as “leaf” information in the IS-IS LSPs using Extended IP reachability TLVs (TLV 135), which contain routing information such as neighbors and locally configured subnets. SPBM nodes receiving the reachability information can use this information to populate the routes to the announcing nodes. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes.

The following figure shows a network running SPBM IP shortcuts.



**Figure 111: SPBM IP Shortcuts**


In this example, BEB A receives a packet with a destination IP address in the subnet of VLAN 14 and knows to forward the packet to BEB D based on the IP route propagation within IS-IS. After a route lookup, BEB A knows that BEB D is the destination for the subnet and constructs a new BMAC header with destination BMAC: D. BCBs B and C need only perform normal Ethernet switching to forward the packet to BEB D. A route lookup is only required once, at the source BEB, to identify BEB D as the node that is closest to the destination subnet.

In contrast to IP routing or Multiprotocol Label Switching (MPLS), SPBM IP shortcuts provide a simpler method of forwarding IP packets in an Ethernet network using the preestablished Ethernet FIBs on the BEBs. SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing SPT. This allows for end-to-end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

In the preceding example, the SPBM nodes in the core that are not enabled with IP shortcuts can be involved in the forwarding of IP traffic. Since SPBM nodes only forward on the MAC addresses that comprise the BMAC header, and since unknown TLVs in IS-IS are relayed to the next hop but ignored locally, SPBM nodes need not be aware of IP subnets to forward IP traffic.

With IP shortcuts, there is only one IP routing hop, as the SPBM backbone acts as a virtualized switching backplane.

The following figure shows a sample campus network implementing SPBM IP shortcuts.



**Figure 112: SPBM IP shortcuts in a campus**

To enable IP shortcuts on the BEBs, you can configure a circuitless IP address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135.

In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct, static, OSPF, RIP, or BGP routes into IS-IS. To advertise IPv6 routes from the BEBs into the SPBM network, you must enable route redistribution of IPv6 direct, IPv6 static, and OSPFv3 routes into IS-IS.

#### *SPBM IPv6 Shortcuts*

Both IPv4 and IPv6 Shortcuts use IS-IS as the Interior Gateway Protocol (IGP) and the link state packet (LSP) for reachability information. However, IPv4 Shortcuts use TLV 135 and IPv6 Shortcuts use TLV 236. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes. IS-IS transports the IPv6 reachability

information to remote BEBs and uses the shortest path, calculated by SPBM, for data forwarding.

**Note**

You only configure the IPv6 address information on the edges. There is no IPv6 in the SPBM cloud.

IS-IS transports the IPv6 routes through TLV 236 in the LSP advertisements. These routes are installed in the Global Routing Table (GRT) with the node from which the LSPs carrying the IPv6 routes are received as the next hop.

**IPv6 Shortcuts Dependency on IPv4 Shortcuts**

IPv6 Shortcuts function in a very similar manner to IPv4 Shortcuts and depends on IPv4 Shortcuts for some functions. For example, IPv6 Shortcuts use the BMAC (local and remote) information created by IPv4 Shortcuts.

**Important**

**IPv4 Shortcuts must be enabled before you enable IPv6 Shortcuts.**

An error is displayed if you try to enable IPv6 Shortcuts but do not have IPv4 Shortcuts already enabled.

IPv6 Shortcuts alone can be disabled while leaving IPv4 Shortcuts enabled. When IPv4 Shortcuts is disabled without disabling IPv6 Shortcuts disabled first, a warning or error message is displayed indicating that IPv6 should be disabled first.

**Circuitless IPv6 (CLIPv6)**

To enable IPv6 Shortcuts on the BEBs and to advertise the local BEB to other IS-IS nodes, you must configure a circuitless IPv6 address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 236.

IPv6 Shortcuts support Circuitless IPv6 (CLIPv6), which ensures uninterrupted connectivity to the switch as long as there is an actual path to reach it. This route always exists and the circuit is always up because there is no physical attachment.

**Migrating the GRT to IPv6 Shortcuts**

Use the following steps to migrate the Global Router Table (GRT) to use IPv6 Shortcuts over the SPBM core:

- Identify the nodes that should be enabled with IPv6 Shortcuts. Apply these steps to all of these nodes.
- Activate and validate basic IPv6 Shortcuts. For information, see [SPBM IPv6 Shortcuts](#) on page 1585.
- Configure IS-IS route preference to ensure that the IPv6 IGP protocol currently being used in the SPBM core is preferred over the IS-IS routes.
- Enable redistribution of direct and static IPv6 routes into IS-IS.
- Create route policies to permit only IPv6 IGP routes from the access side of the SPBM network.



- Configure redistribution of routes from the IPv6 route table from each of the IPv6 IGP protocols into IS-IS along with the appropriate route policy.
- Use the **show isis spbm ipv6-unicast-fib** command to check the IS-IS LSDB, IS-IS routes, and to verify that all the desired IPv6 routes are now in IS-IS.
- Configure redistribution of IS-IS routes from the IPv6 route table into each of the IPv6 IGP protocols in use. This redistribution does not require a route policy since IS-IS is only supported in the SPBM core.
- Change IS-IS route-preference to ensure that IS-IS routes are preferred over other IPv6 IGP routes.
- Disable/delete old IPv6 IGP in the SPBM core.



### Important

Use only one IPv6 routing protocol in the SPBM core to prevent the possibility of routing loops.

### IPv6 Shortcut Restrictions and Considerations

The following features are not supported:

- Disabling and enabling alternate routes for IPv6 routes
- Redistribution of RIP into IS-IS
- 6-in-4 tunnels are not supported when the tunnel destination IP is reachable via IPv4 Shortcuts route.

Keep the following considerations in mind when configuring IPv6 Shortcuts:

- IPv4 Shortcuts must be enabled before enabling IPv6 Shortcuts.
- IPv6 Shortcuts support Circuitless IPv6 (CLIPv6) with the following limitations:
  - Stateless address autoconfiguration (SLAAC) is not supported on IPv6 CLIP interfaces.
  - IPv6 CLIP does not support link-local address configuration.
  - To configure an IPv6 address with a prefix length from 65 to 127 on a CLIP interface, you must enable the IPv6 mode flag.
  - Neighbor discovery (ND) does not run on an IPv6 CLIP interface. Therefore, the system does not detect when you configure a duplicate IPv6 address.
  - Multiple IPv6 address configuration on an IPv6 CLIP interface is not supported.
  - You can configure a maximum of 64 IPv6 CLIP interfaces.
  - IPv6 CLIP interface is enabled by default and it cannot be disabled.
- IPv6 with vIST provides the same support as IPv4 with vIST.
- To help with debugging, CFM provides full support for both IPv4 and IPv6 addresses for the **12ping** and **12traceroute** commands.

### ECMP with IS-IS

The Equal Cost Multipath (ECMP) feature supports and complements the IS-IS protocol.

With ECMP, the switch can determine multiple equal-cost paths to the same destination prefix.

You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.

The number of multiple paths a switch can support differs by hardware platform. For more information about feature support, see *Fabric Engine Release Notes*.

### ECMP Within IS-IS routes

Equal Cost Multipath (ECMP) allows the device to determine up to eight equal cost paths to the same destination prefix. The maximum number of equal cost paths you can configure depends on the hardware platform. For more information, see *Fabric Engine Release Notes*.

If the device learns the same route from multiple sources, the information is ECMP only if the routes:

- are from the same VSN
- have the same SPBM cost
- have the same prefix cost
- have the same IP route preference

Multiple BEBs can announce the same route, either because the Layer 2 LAN connects to multiple BEBs for redundancy, or because segments of the LAN are Layer 2 bridged. In Layer 2, if the device has to tie-break between multiple sources, the tie-breaking is based on cost and hop count.

In Layer 3, hop count is not used for tie-breaking. Instead, the device uses the following precedence rules to tie-break. In the following order, the device prefers:

1. Routes that do not include nodes with the overload bit set.

When a router node runs out of system resources (memory or CPU), it alerts the other routers in the network by setting the overload bit in its link-state packets (LSPs). When this bit is set, the node is not used for transit traffic but only for traffic packets destined to the node's directly connected networks and IP prefixes.

2. IS-IS routes within the same VSN over IS-IS inter-VSN routes (inter-VRF leaked routes).

If a route is learned from the same I-SID, it is the most preferred.

3. Routes with the lowest route preference.

By default, IS-IS routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. You can however, change the route preference using IS-IS accept policies.

4. Metric type internal (type 1) over metric type external (type 2).
5. Routes with the lowest SPBM cost.

6. Routes with the lowest prefix cost.

If the metric type is internal, then the tie-break is on SPB cost first, and then on the prefix cost. Otherwise the tie-break is only on the prefix cost.

You can either change this using a route-map on the remote advertising node with the **redistribute** command, or using a route-map on the local node with the IS-IS accept policy.

7. Routes within a VSN with a lower Layer 3 VSN I-SID.

The device considers the Global Routing Table (GRT) to have an I-SID equal to zero.

When you use multiple B-VLANs in the SPBM core, multiple paths exist to reach a particular SPBM node, one on each B-VLAN; therefore, any IP prefix or IPv6 prefix that the device receives from a BEB results in multiple ECMP paths. These paths may or may not be physically diverse. SPBM supports up to two B-VLANs; a primary B-VLAN and a secondary B-VLAN.

If more ECMP paths are available than the configured number of paths, then the device adds the routes using the following order: The device selects all routes from the primary B-VLAN and orders the routes learned through that B-VLAN from lowest system ID to the highest IS-IS system ID, then the device moves on to select all routes from the secondary B-VLAN, ordering those routes from lowest IS-IS system ID to the highest IS-IS system ID until you reach the number of equal paths configured.

For example, consider an SPB core configured with two B-VLANs (primary B-VLAN 1000 and secondary B-VLAN 2000), and the device learns routes from two BEBs called BEB-A (with a lower IS-IS system ID) and BEB-B (with a higher IS-IS system ID, then the order in which the next-hop is chosen for those routes are as follows.

If a route is learned only from BEB-A with the maximum number of allowed ECMP paths configured as 8 (default), then the order in which the next-hop is chosen for that route is:

1. BEB-A B-VLAN 1000
2. BEB-A B-VLAN 2000

If routes are learned from both BEB-A and BEB-B with maximum number of allowed ECMP paths configured as 8 (default), then the order in which the next-hop is chosen for those routes are:

1. BEB-A B-VLAN 1000
2. BEB-B B-VLAN 1000
3. BEB-A B-VLAN 2000
4. BEB-B B-VLAN 2000

If ECMP is disabled, the maximum number of allowed ECMP paths is 1 and the device adds the route from the lowest system ID with the primary B-VLAN. In this example, the device adds BEB-A B-VLAN 1000.

**Note**

- ECMP is supported for IPv6 Shortcut routes.
- To add IS-IS equal cost paths in the routing table, you must enable IPv6 ECMP feature globally.

**ECMP Impact on IS-IS Route Selection for Inter-VRF Routes with vIST**

This section illustrates the impact ECMP can have on a configuration that implements user-defined VRFs in a vIST cluster and how to avoid incorrect route selection.

**Understanding the Configuration**

Imagine the following configuration:

- A vIST cluster exists with multiple VRF contexts.
- On both nodes, VRF A redistributes routes into IS-IS as external. VRF B uses an IS-IS accept policy to accept these routes.
- Each node learns three paths to the route:
  - The nodes learn one path using local inter-VRF redistribution.
  - The nodes learn the other two paths from the IST peer.
- The routes are treated as ECMP paths because the preference, metric-type, and metric are equal.

IS-IS sorts paths for the same route by source-BEB BMAC and B-VLAN ID. The primary B-VLAN ID is first installed for each BMAC, followed by the secondary B-VLAN ID for each BMAC, as long as the ECMP max-path value is not reached. On the node with the lowest BMAC, the first path listed is its own local inter-vrf route, while on the other node, the MIM path across the vIST is listed first.

If you disable ECMP, all but the first path is removed. Because IS-IS orders the paths by BMAC, each node in the vIST cluster selects the same BMAC as the nexthop. This configuration leads one of the nodes to select itself, the local inter-vrf route, while the other node selects the MIM path across the vIST to get to the inter-vrf route. This situation results in an incorrect route selection.

**Avoiding Incorrect Route Selection**

To avoid this situation, create a policy to prevent IS-IS from determining that the MIM path across the vIST and the local inter-VRF route are ECMP paths. Configure the local inter-VRF path as the preferred path, and the vIST path as the backup. The following list identifies way that you can accomplish this:

- Redistribute the VRF route into IS-IS using the internal metric-type. IS-IS will always select the local inter-VRF route. For more information about the metric type for IS-IS routes, see [Fabric Basics and Layer 2 Services](#) on page 1178.


- If an IS-IS internal metric-type is not an option, configure an IS-IS accept policy to change the preference of inter-VRF routes learned from the IST peer. The local inter-VRF route is preferred over the inter-VRF routes learned from the IST peer.

### IS-IS IP Redistribution Policies

When you connect an SPBM core using IP shortcuts to existing networks running a routing protocol such as OSPF or RIP, a redundant configuration requires two switches:

- One router redistributes IP routes from Routing Information Protocol (RIP)/Open Shortest Path First (OSPF) into IS-IS (IP).
- The second router redistributes from IS-IS (IP) into RIP or OSPF.

The following figure illustrates this configuration.



**Figure 113: Redundant OSPF or RIP Network**

In this scenario, take extra care when redistributing through both switches. By default, the preference value for IP routes generated by SPBM-IP (IS-IS) is 7. This is a higher preference than OSPF (20 for intra-area, 25 for inter-area, 120 for ext type1, 125 for ext type2) or RIP (100).




#### Note

The lower numerical value determines the higher preference.

In the preceding diagram both nodes (SwitchG and SwitchD) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to SwitchA.

As soon as the SwitchG node redistributes that IP route into IS-IS, the SwitchD node learns the same route through IS-IS from SwitchG. (The SwitchG node already has the route through OSPF or RIP). Because IS-IS has a higher preference, SwitchD replaces

its 192.168.10.0 OSPF route with an IS-IS one that points at SwitchG as the next-hop. The following figure illustrates this scenario.



**Figure 114: Redistributing Routes into IS-IS**

Clearly, this is undesirable and you must ensure that the two redistributing nodes (SwitchG and SwitchD) do not accept redistributed routes from each other. With IS-IS accept policies, you can associate an IS-IS accept policy on SwitchD to reject all redistributed IP routes received from SwitchG, and SwitchG to reject all redistribute IP routes from SwitchD.

An alternate way to solve the preceding problem with existing functionality is to reverse the problem by lowering the SPBM-IP (IS-IS) preference by configuring it to a value greater than RIP (100) or OSPF (20,25,120,125). For example, log on to Global Configuration mode and use the following command to configure a preference of 130:

```
ip route preference protocol spbm-level1 130
```



**Note**

For IPv6, the command is `ipv6 route preference protocol spbm-level1 130`

Now that the OSPF or RIP routes have a higher preference than SPBM-IP (IS-IS), the problem is temporarily solved. However, the same issue resurfaces when the IS-IS IP


routes are redistributed into OSPF or RIP in the reverse direction as shown in the following section for OSPF.



### Important

You can apply a tag for routes that BGP or OSPF redistribute into IS-IS. This configuration ensures that if one BEB redistributes a route into IS-IS from either protocol, other BEBs do not redistribute tagged IS-IS routes back into BGP or OSPF. You can also create route-map policies to match the IS-IS tag. For more information, see [IS-IS Route Tagging](#) on page 1200.

### OSPF Example



**Figure 115: Redistributing Routes into OSPF**

In the preceding figure, both SwitchG and SwitchD have an IS-IS IP route for 172.16.0.0/16 with the next hop as SwitchC. As soon as SwitchG redistributes the IS-IS route into OSPF, the SwitchD node learns that same route through OSPF from SwitchG. (The SwitchG node already has the route through IS-IS).

Because OSPF has a higher preference, SwitchD replaces its 172.16.0.0/16 IS-IS route with an OSPF one. (Note that the 172.16.0.0/16 route will be redistributed into OSPF as an AS external route, hence with preference 120 or 125 depending on whether type1 or type2 was used). In this case, however, you can leverage OSPF Accept policies, which can be configured to prevent SwitchD from accepting any AS External (LSA5) routes from SwitchG and prevent SwitchG from accepting any AS External (LSA5) routes from SwitchD. The following is a sample configuration:

```
enable
configure terminal
route-map
IP ROUTE MAP CONFIGURATION - GlobalRouter
```

```
route-map "reject" 1
no permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

OSPF CONFIGURATION - GlobalRouter

router ospf enable

OSPF ACCEPT CONFIGURATION - GlobalRouter

router ospf
accept adv-rtr {A.B.C.D}
accept adv-rtr {A.B.C.D} enable route-map "reject"
exit
```



#### Note

Disable alternative routes by issuing the command **no ip alternative-route** to avoid routing loops on the SMLT Backbone Edge Bridges (BEBs).


In the preceding figure, if SwitchA advertises 25000 OSPF routes to SwitchG and SwitchD, then both SwitchG and SwitchD install the 25000 routes as OSPF routes. Since SwitchD and SwitchG have OSPF to IS-IS redistribution enabled, they also learn these 25000 routes as IS-IS routes. IS-IS route preference is configured with a higher numerical value (130) than the OSPF route preference (125), so SwitchD and SwitchG keep IS-IS learned routes as alternative routes.

If SwitchA withdraws its 25000 OSPF routes, SwitchG and SwitchD remove the OSPF routes. While the OSPF routes are removed the routing tables of SwitchG and SwitchD activate the alternative IS-IS routes for the same prefix. Because SwitchG and SwitchD have IS-IS to OSPF redistribution enabled, SwitchA learns these routes as OSPF and this causes a routing loop. Use the **no ip alternative-route** command to disable alternative routes on SwitchG and SwitchD to avoid routing loops.

#### RIP Examples

In the preceding section, you leveraged OSPF Accept policies, which can be configured to prevent SwitchD from accepting any AS External (LSA5) routes from SwitchG and prevent SwitchG from accepting any AS External (LSA5) routes from SwitchD. In the case of a RIP access network, the preceding solution is not possible because RIP has no concept of external routes and no equivalent of accept policies. However, if you assume that a RIP network acts as an access network to an SPBM core, then it is sufficient to ensure that when IS-IS IP routes are redistributed into RIP they are aggregated into a single default route at the same time. The following figure and sample configuration example illustrates this scenario:





**Figure 116: Redistributing Routes into RIP**

```
SwitchG
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

RIP PORT CONFIGURATION

interface gigabitethernet 1/11
ip rip default-supply enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router rip
```

```
redistribute isis
redistribute isis metric 1
redistribute isis route-map "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis

SwitchA
RIP PORT CONFIGURATION

interface gigabitethernet 1/2
ip rip default-listen enable
exit
interface gigabitethernet 1/3
ip rip default-listen enable
exit

SwitchD
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

RIP PORT CONFIGURATION

interface gigabitethernet 1/11
ip rip default-supply enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router rip
redistribute isis
redistribute isis metric 1
redistribute isis route-map "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

You can control the propagation of the default route on the RIP network so that both SwitchG and SwitchD supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, SwitchA will accept the default route on its interfaces to both SwitchG and SwitchD but it will not supply the default route back to them. This will prevent the default route advertised by SwitchG from being installed by SwitchD, and vice-versa.

The preceding example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network also applies when redistributing IS-IS IP routes into OSPF if that OSPF network is an access network to an SPBM core. In this case use the following redistribution policy configuration as an example for injecting IS-IS IP routes into OSPF:

```
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

OSPF CONFIGURATION - GlobalRouter

router ospf enable
router ospf
as-boundary-router enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router ospf
redistribute isis
redistribute isis route-map "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip ospf apply redistribute isis
```

### *IS-IS Accept Policies*

You can use Intermediate-System-to-Intermediate-System (IS-IS) accept policies (for IPv4 and IPv6) to filter incoming IS-IS routes over the SPBM cloud and apply route

policies to the incoming IS-IS routes. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table.

### IS-IS Accept Policies and DvR



#### Note

IPv6 IS-IS accept policies for DvR are not supported.

When you configure DvR in an SPB network, you can leverage IS-IS accept policies to control the DvR routes learned from the DvR backbone. The DvR backbone contains the master list of all the host routes learned from various DvR domains.

You can configure accept policies on a DvR Controller or a non-DvR BEB as a filter to determine which DvR host routes to accept into the routing table, from the DvR backbone. Accept policies apply to only those backbone (or inter-domain) host routes that are not part of the Controller's own DvR enabled subnets and do not have the same domain ID as that of the Controller.

For non-DvR BEBs, all the routes present in the backbone are learned, but you can still use the accept policies to filter specific routes.

For information on DvR, see [Distributed Virtual Routing](#) on page 915.

### IS-IS Accept Policy Filters

You can filter traffic with IS-IS accept policies by:

- advertising BEB
- I-SID or I-SID list
- route-map
- backbone-route-map for IPv4 only
- a combination of route-map and backbone-route-map for IPv4 only

You can use IS-IS accept policies to apply at a global default level for all advertising Backbone Edge Bridges (BEBs) or for a specific advertising BEB.

IS-IS accept policies also allow you to use either a service instance identifier (I-SID) or an I-SID list to filter routes. The switch uses I-SIDs to define Virtual Services Networks (VSNs). I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. IS-IS accept policies can use I-SIDs or I-SID lists to filter the incoming virtualized traffic.

IS-IS accept policies can also apply route policies to determine what incoming traffic to accept into the routing table. With route policies the device can determine which routes to accept into the routing table based on the criteria you configure. You can match on the network or the route metric.

On DvR Controllers in a DvR domain, you can configure a backbone route policy to determine what host routes to accept from the DvR backbone, into the routing table. Also, just like on the route policy, you can configure match criteria, and set preferences on the backbone route policy.

To accept both IS-IS routes and host routes from the DvR backbone, you can configure both a route policy and a backbone route policy in the accept policy instance.

For more information on configuring route policies:

- For IPv4, see [IPv4 Routing Operations Fundamentals](#) on page 2099.
- For IPv6, see [IPv6 Routing Basics](#) on page 2167.

The following table describes IPv4 IS-IS accept policy filters.

| Filters into               | Filter                                                                                                                                | Description                                                                                                                                                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global Routing Table (GRT) | <code>accept route-map WORD&lt;1-64&gt;</code>                                                                                        | By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.                                                                                                                                                      |
|                            | <code>accept route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                                                    | This is the default accept policy with configuration to accept specific DvR host routes from the DvR backbone.                                                                                                                                                        |
|                            | <code>accept adv-rtr &lt;x.xx.xx&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                            | The device filters based on the specific advertising BEB defined by the SPBM nickname.<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                                   |
|                            | <code>accept i-sid &lt;1-16777215&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                           | The device filters based on the I-SID, which represents a local or remote Layer 3 VSN.<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                                   |
|                            | <code>accept adv-rtr &lt;x.xx.xx&gt; i-sid &lt;1-16777215&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>   | The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN.<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                  |
|                            | <code>accept isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                         | The device filters based on the list of I-SIDs.<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                                                                          |
|                            | <code>accept adv-rtr &lt;x.xx.xx&gt; isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code> | The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter. |

| Filters into                                       | Filter                                                                                                                                     | Description                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual Routing and Forwarding (VRF) routing table | <code>isis accept adv-rtr &lt;x.xx.xx&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                            | The device filters based on the specific advertising BEB defined by the SPBM nickname.<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                                                                          |
|                                                    | <code>isis accept i-sid &lt;0-16777215&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                           | The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                  |
|                                                    | <code>isis accept adv-rtr &lt;x.xx.xx&gt; i-sid &lt;0-16777215&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>   | The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter. |
|                                                    | <code>isis accept isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                         | The device filters based on the list of I-SIDs to which the IS-IS accept policy applies. The number 0 represents the Global Routing Table (GRT).<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                |
|                                                    | <code>isis accept adv-rtr &lt;x.xx.xx&gt; isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code> | The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                        |
|                                                    | <code>isis accept route-map WORD&lt;1-64&gt; route-map WORD&lt;1-64&gt; backbone-route-map WORD&lt;1-64&gt;</code>                         | The device filters based on the route policy.<br>The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.                                                                                                                                   |

The following table describes the IPv6 IS-IS accept policy filters:

| Filters into               | Filter                                                                                           | Description                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Global Routing Table (GRT) | <i>ipv6 accept route-map WORD&lt;1-64&gt;</i>                                                    | By default, the device accepts all routes advertised. This is the default accept policy.                                                 |
|                            | <i>ipv6 accept adv-rtr &lt;x.xx.xx&gt; route-map WORD&lt;1-64&gt;</i>                            | The device filters based on the specific advertising BEB defined by the SPBM nickname.                                                   |
|                            | <i>ipv6 accept i-sid &lt;1-16777215&gt; route-map WORD&lt;1-64&gt;</i>                           | The device filters based on the I-SID, which represents a local or remote Layer 3 VSN.                                                   |
|                            | <i>ipv6 accept adv-rtr &lt;x.xx.xx&gt; i-sid &lt;1-16777215&gt; route-map WORD&lt;1-64&gt;</i>   | The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN.                  |
|                            | <i>ipv6 accept isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt;</i>                         | The device filters based on the list of I-SIDs.                                                                                          |
|                            | <i>ipv6 accept adv-rtr &lt;x.xx.xx&gt; isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt;</i> | The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT). |



| Filters into                                       | Filter                                                                                                      | Description                                                                                                                                                                     |
|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual Routing and Forwarding (VRF) routing table | <code>ipv6 isis accept route-map WORD&lt;1-64&gt; route-map WORD&lt;1-64&gt;</code>                         | The device filters based on the route policy.                                                                                                                                   |
|                                                    | <code>ipv6 isis accept adv-rtr &lt;x.xx.xx&gt; route-map WORD&lt;1-64&gt;</code>                            | The device filters based on the specific advertising BEB defined by the SPBM nickname.                                                                                          |
|                                                    | <code>ipv6 isis accept i-sid &lt;0-16777215&gt; route-map WORD&lt;1-64&gt;</code>                           | The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT).                                  |
|                                                    | <code>ipv6 isis accept adv-rtr &lt;x.xx.xx&gt; i-sid &lt;0-16777215&gt; route-map WORD&lt;1-64&gt;</code>   | The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT). |
|                                                    | <code>ipv6 isis accept isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt;</code>                         | The device filters based on the list of I-SIDs to which the IS-IS accept policy applies. The number 0 represents the Global Routing Table (GRT).                                |
|                                                    | <code>ipv6 isis accept adv-rtr &lt;x.xx.xx&gt; isid-list WORD&lt;1-32&gt; route-map WORD&lt;1-64&gt;</code> | The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT).                                        |

### IS-IS Accept Policies for the GRT and VRFs

You can create an IS-IS accept policy for incoming routes for the Global Routing Table (GRT), which accepts routes into the routing table, or for a Virtual Routing and Forwarding (VRF) instance, which accepts incoming routes to the routing table of the VRF.

If you create an IS-IS accept policy on the switch for either the GRT or a VRF that operates at a global default level, the accept policy applies to all routes for all BEBs in the GRT or VRF.

If you create an IS-IS accept policy on the switch for a specific advertising BEB for either the GRT or a VRF, the IS-IS accept policy instance applies for that specific advertising BEB. If you use a more specific filter, the system gives preference to the specific filter over the global default level.

### IS-IS Accept Policies for Inter-VRF Route Redistribution

You can also use the filter mechanism for IS-IS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT. For inter-VRF route redistribution, you match the filter based on the I-SID, which represents the Layer 3 VSN context.

You can apply the filter at the global default level, where the IS-IS accept policy applies to all routes for that I-SID from all BEBs, or at a specific advertising BEB level, where the filter only applies to a specific advertising BEB. The device gives preference to a specific filter for a specific advertising BEB over the global default filter.

For inter-VRF route redistribution, an I-SID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the I-SID is the source VRF (or remote VRF).

**Note**

If the primary B-VLAN is down either because you did not configure at least one network-to-network interface (NNI) or all configured NNIs are down, the switch does not redistribute inter-VRF routes through IS-IS accept policies.

**IS-IS Accept Policy Considerations**

Consider the following when you configure IS-IS accept policies:

- If a VRF uses a different protocol to redistribute routes from another VRF, the IS-IS accept policy feature cannot be used. You can only use the IS-IS accept policy for inter-VSN route redistribution between VRFs.
- IPv4 and IPv6 IS-IS accept policies can exist on the same VRF and GRT; The I-SID list configuration is shared across both protocol versions.

**Precedence Rules in the Same VSN**

The following precedence rules apply for IS-IS accept policies used in the same VSN:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply either a default filter for all advertising BEBs or a filter for a specific advertising BEB.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device prefers the **accept adv-rtr** filter, which filters based on a specific advertising BEB, over the default filter for all advertising BEBs.
- The device accepts all routes within the same VSN by default. You can apply a route policy to filter or change the characteristics of the route by metric or preference.
- The **i-sid** or **isid-list** filters are not valid for routes within the same VSN.

**Precedence Rules for Inter-VSN Route Redistribution**

The following precedence rules apply for IS-IS accept policies used for inter-VSN route redistribution:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply filters at a global default level for all BEBs for a specific I-SID or I-SID list, or you can apply filters for a specific advertising BEB for a specific I-SID or I-SID list.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device requires a specific filter to redistribute routes between VSNs through the use of the **i-sid** or **isid-list** filters.

- The **i-sid** filter takes precedence over the **isid-list** filter.
- The **adv-rtr** filter for a specific advertising BEB takes precedence over a filter with the same **i-sid** filter without the **adv-rtr** filter.
- The **i-sid** or **isid-list** filters only apply to routes for inter-VSN route redistribution.
- If multiple **isid-list** filters have the same I-SID within the list, the first on the list alphabetically has the higher precedence.

### Route Preference

The relative value of the route preference among different protocols determines which protocol the device prefers. If multiple protocols are in the routing table, the device prefers the route with the lower value. You can change the value at the protocol level, and you can also change the preference of incoming IS-IS routes using the route-map with the IS-IS Accept policy filter for IPv4 only.

### Route Metric

Use a route-map to change the metric of a route when you accept a remote IS-IS route with IS-IS accept policies.

You can use a route-map to change the metric of a route when you redistribute the route from another protocol to IS-IS through the route redistribution mechanism.

You can also configure the route metric with the base **redistribute** command without the use of a route-map.



#### Note

For both IPv4 and IPv6 IS-IS accept policies, if there is a mismatch in the route-map (inbound filtering) configured, all routes are accepted by default. Unlike the redistribute route-map (outbound filtering), where if there is a mismatch, all routes are denied by default. For more information, see [IPv4 Routing Operations Fundamentals](#) on page 2099.

For more information on the configuration of a route-map:

- For IPv4, see [IPv4 Routing Operations Fundamentals](#) on page 2099.
- For IPv6, see [IPv6 Routing Basics](#) on page 2167.

## IP Shortcuts Configuration using the CLI

This section provides procedures to configure IP Shortcuts using the command line interface (CLI).

### Configure SPBM IPv4 Shortcuts

#### Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, create the Customer VLANs, add slots/ports, and add the IP addresses and network masks.

## About This Task

To enable IP shortcuts on the BEBs, you can configure a circuitless IP (CLIP) address (loopback address), and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.



### Note

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address.

## Procedure

1. Enter Loopback Interface Configuration mode  

```
enable
```

```
configure terminal
```

```
interface Loopback <1-256>
```
2. Configure a CLIP interface to use as the source address for SPBM IP shortcuts:  

```
ip address [<1-256>] <A.B.C.D/X>
```
3. Exit the Loopback Interface Configuration mode to Global Configuration mode:  

```
exit
```
4. Log on to IS-IS Router Configuration mode:  

```
router isis
```
5. Specify the CLIP interface as the source address for SPBM IP shortcuts:  

```
ip-source-address <A.B.C.D>
```
6. Configure SPBM IP shortcuts:  

```
spbm <1-100> ip enable
```
7. Display the status of SPBM IP shortcuts on the switch:  

```
show isis spbm
```
8. Identify routes on the local switch to be announced into the SPBM network:  

```
redistribute {bgp | direct | ospf | rip | static}
```
9. (Optional) Tag routes that BGP or OSPF redistribute into IS-IS:  

```
redistribute {bgp | ospf} tag <1-4294967295>
```
10. Enable routes to be announced into the SPBM network  

```
redistribute {bgp | direct | ospf | rip | static} enable
```
11. Exit to Global Configuration mode:  

```
exit
```
12. Apply the configured redistribution:  

```
isis apply redistribute {bgp | direct | ospf | rip | static | vrf
WORD<1-16>}
```

## Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 192.0.2.2/8
Switch:1(config-if)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#ip-source-address 192.0.2.2
Switch:1(config-isis)#spbm 1 ip enable
Switch:1(config-isis)#show isis spbm

=====
ISIS SPBM Info
=====
SPBM B-VID PRIMARY NICK LSDB IP IPV6 MULTICAST
INSTANCE VLAN NAME TRAP

1 4086-4087 4086 3.03.01 disable enable enable disable

=====
ISIS SPBM SMLT Info
=====
SPBM SMLT-SPLIT-BEB SMLT-VIRTUAL-BMAC SMLT-PEER-SYSTEM-ID SMLT-PEER-NICK-NAME SMLT-PEER-HOST-NAME
INSTANCE

1 primary 00:00:03:03:03:03 0000.0303.0302 02.02.01 Switch-3

Total Num of SPBM instances: 1

Switch:1(config-isis)#redistribute rip
Switch:1(config-isis)#redistribute rip enable
Switch:1(config-isis)#exit
Switch:1(config)#isis apply redistribute rip
```

## Variable Definitions

The following table defines parameters for the **ip address** command.

| Variable    | Value                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------|
| <1-256>     | Specifies an interface ID value. This value is optional.                                              |
| <A.B.C.D/X> | Specifies an IP address and subnet mask. Use the <b>no</b> option to delete the specified IP address. |
| <A.B.C.D>   | Specifies an IP address. Use the <b>no</b> option to delete the specified IP address.                 |

The following table defines parameters for the **ip-source-address** command.

| Variable  | Value                                                                            |
|-----------|----------------------------------------------------------------------------------|
| <A.B.C.D> | Specifies the CLIP interface to use as the source address for SPBM IP shortcuts. |

The following table defines parameters for the **spbm** command.

| Variable                             | Value                                                                                                                                                         |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>&lt;1-100&gt; ip enable</code> | Enables or disables SPBM IP shortcut state.<br>The default is disabled. Use the <code>no</code> or <code>default</code> options to disable SPBM IP shortcuts. |

The following table defines parameters for the **redistribute** command.

| Variable                                          | Value                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>{bgp   direct   ospf   rip   static}</code> | Specifies the protocol.                                                                                                                                                                                                                                                                                                                         |
| <code>enable</code>                               | Enables the redistribution of the specified protocol into the SPBM network.<br>The default is disabled. Use the <b>no</b> option to disable the redistribution.                                                                                                                                                                                 |
| <code>metric &lt;0-65535&gt;</code>               | Configures the metric (cost) to apply to redistributed routes. The default is 1.                                                                                                                                                                                                                                                                |
| <code>metric-type {external internal}</code>      | Configures the type of route to import into the protocol.<br>The default is internal.                                                                                                                                                                                                                                                           |
| <code>route-map WORD&lt;0-64&gt;</code>           | Configures the route policy to apply to redistributed routes. Type a name between 0 to 64 characters in length.                                                                                                                                                                                                                                 |
| <code>subnets {allow suppress}</code>             | Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose <code>suppress</code> to advertise subnets aggregated to their classful subnet. Choose <code>allow</code> to advertise the subnets individually with the learned or configured mask of the subnet. The default is <code>allow</code> . |
| <code>tag &lt;1-4294967295&gt;</code>             | Apply a tag to routes that BGP or OSPF redistribute to IS-IS. You can then configure route filtering to prevent tagged IS-IS routes from being redistributed back to either protocol.                                                                                                                                                           |

The following table defines parameters for the **isis apply redistribute** command.

| Variable                                          | Value                   |
|---------------------------------------------------|-------------------------|
| <code>{bgp   direct   ospf   rip   static}</code> | Specifies the protocol. |

### Configure SPBM IPv6 Shortcuts

#### Before You Begin

- Enable IPv4 Shortcuts before you enable IPv6 Shortcuts. IPv6 Shortcuts depend on IPv4 Shortcuts for some functions.
- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, create the Customer VLANs, add slots/ports, and add the IPv6 addresses and network masks.

## About This Task

To enable IPv6 Shortcuts on the BEBs, you must configure a circuitless IPv6 (CLIPv6) address (loopback address), and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 236. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.



### Note

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address, and the CLIPv6 address prefix must be 128.

## Procedure

1. Enter Loopback Interface Configuration mode
 

```
enable

configure terminal

interface Loopback <1-256>
```
2. Configure a CLIPv6 interface to use as the source address for SPBM IPv6 Shortcuts:
 

```
ipv6 interface address WORD<0-255>
```
3. Exit the Loopback Interface Configuration mode to Global Configuration mode:
 

```
exit
```
4. Log on to IS-IS Router Configuration mode:
 

```
router isis
```
5. Specify the CLIPv6 interface as the source address for SPBM IPv6 Shortcuts:
 

```
ipv6-source-address WORD<0-46>
```
6. Enable SPBM IPv6 Shortcuts:
 

```
spbm <1-100> ipv6 enable
```
7. Display the status of SPBM IPv6 Shortcuts on the switch:
 

```
show isis spbm
```
8. Identify IPv6 routes on the local switch to be announced into the SPBM network.
 

```
ipv6 redistribute {bgp | direct | ospf | rip | static}
```
9. Enable the IPv6 routes to be announced into the SPBM network:
 

```
ipv6 redistribute {bgp | direct | ospf | rip | static} enable
```
10. Exit to Global Configuration mode:
 

```
exit
```
11. (Optional) Change route preference values for IPv6 Shortcut routes learned through IS-IS:
 

```
ipv6 route preference protocol spbm-level1 <0-255>
```
12. Apply the configured redistribution:
 

```
ipv6 isis apply redistribute {bgp | direct | ospf | rip | static |}
[vrf WORD<1-16>]
```

## Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 123
Switch:1(config-if)#ipv6 interface address 123::1/128
Switch:1(config-if)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#ipv6 ipv6-source-address <non-link-local ipv6-address>
Switch:1(config-isis)#spbm 1 ipv6 enable

Switch:1(config-isis)#show isis spbm
=====
SPBM B-VID PRIMARY NICK LSDB IP IPV6 MULTICAST SPB-PIM-GW STP-MULTI
INSTANCE VLAN NAME TRAP

1 10 1.11.16 disable disable disable disable disable disable enable
=====

 ISIS SPBM SMLT Info
=====

SPBM SMLT-SPLIT-BEB SMLT-VIRTUAL-EMAC SMLT-PEER-SYSTEM-ID SMLT-PEER-NICK-NAME SMLT-PEER-HOST-NAME
INSTANCE

1 primary 00:00:00:00:00:00 N/A N/A

Total Num of SPBM instances: 1
=====
```

## Variable Definitions

The following table defines parameters for the IPv6 Shortcuts commands.

| Variable                                                                                 | Value                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ipv6-source-address</i><br><i>WORD&lt;0-46&gt;</i>                                    | Specifies the source IPv6 address for locally generated IPv6 packets whose egress port is an SPBM network-to-network interface (NNI) port. The <i>WORD&lt;0-46&gt;</i> value must be a locally configured loopback IPv6 address (CLIPv6). Use the <b>no</b> option to delete the specified IPv6 address. |
| <i>spbm&lt;1-100&gt; ipv6 enable</i>                                                     | Enables or disables SPBM IPv6 Shortcuts. The default is disabled. Use the <b>no</b> or <b>default</b> options to disable SPBM IPv6 Shortcuts.                                                                                                                                                            |
| <i>ipv6 route preference</i><br><i>protocol spbm-level1</i><br><i>&lt;0-255&gt;</i>      | Sets the route preference value for IPv6 Shortcut routes learned through IS-IS. The default preference is 7.                                                                                                                                                                                             |
| <i>ipv6 redistribute {bgp</i><br><i>  direct   static  </i><br><i>ospf   rip} enable</i> | Specifies the GRT IPv6 route that you want to redistribute into the SPBM IS-IS domain. The default is disabled. Use the <b>no</b> option to disable the redistribution.                                                                                                                                  |

## Configure Inter-VRF IPv4 Accept Policies on VRFs

### Before You Begin

- Enable IS-IS globally.
- Ensure that a route policy exists.



- Ensure that the VRFs exist.
- Configure a route-map.

### About This Task

Configure IS-IS accept policies on a VRF to use inter-VRF accept policies in the SPB cloud. You can use IS-IS accept policies to redistribute routes between different VRFs, including the global routing table (GRT). First you apply the filter, and then you match the filter based on the I-SID, which represents the Layer 3 VSN context.



#### Note

- The **isis apply accept [vrf WORD<1-16>]** command can disrupt traffic and cause temporary traffic loss. After you apply **isis apply accept [vrf<1-16>]**, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply **isis apply accept [vrf WORD<1-16>]** at the end.
- If you use the **accept** command for inter-VRF routes based on the remote I-SID, the device only accepts routes coming from remote BEBs. For instance, if a local Layer 3 VSN exists with the same I-SID, the device does not add the local routes. The assumption is that the device uses existent methods, either through use of another protocol or static configuration, to obtain those routes.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

### Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



#### Note

You cannot use the derived name for the untrusted (\*-unt) or unrestricted (\*-unr) segments of a Segmented VRF to enter this mode.

- (Optional) To accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> [<0-16777215>] [list WORD<1-1024>]
```



#### Note

When creating an I-SID list, you can add I-SID entries until the maximum limit for supported Layer 3 I-SIDs is reached. The system truncates any additional I-SID entries. The maximum limit includes the I-SIDs for locally configured Layer 3 VSNs and the I-SIDs specified for IS-IS accept policy filters.

Use the **show ip isid-list vrf WORD<1-16>** command to view the list of truncated I-SIDs.

- Create an IS-IS accept policy instance to apply to routes from all Backbone Edge Bridges (BEBs):

```
isis accept [i-sid <0-16777215>] [isid-list WORD<1-32>]
```

- Create an IS-IS accept policy instance to apply to routes for a specific BEB:

```
isis accept [adv-rtr <x.xx.xx>] [i-sid <0-16777215>] [isid-list WORD<1-32>]
```

- Specify an IS-IS route policy to apply to routes from all BEBs:

```
isis accept route-map WORD<1-64>
```

- Specify an IS-IS route policy to apply for a specific BEB:

```
isis accept adv-rtr <x.xx.xx> route-map WORD<1-64>
```

- Enable a configured IS-IS accept policy instance:

```
isis accept [adv-rtr <x.xx.xx>] [i-sid <0-16777215>] [isid-list WORD<1-32>] [enable]
```

- Exit VRF Router Configuration mode:

```
exit
```

You are in Global Configuration mode.

- Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
isis apply accept [vrf WORD<1-16>]
```

## Example

Configure Inter-VRF accept policies on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router vrf green
Switch:1(router-vrf)#isis accept i-sid 100
Switch:1(router-vrf)#isis accept i-sid 100 enable
Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept vrf green
```

## Variable definitions

The following table defines parameters for the **ip isid-list** command.

| Variable                  | Value                                                                    |
|---------------------------|--------------------------------------------------------------------------|
| <i>WORD</i> <1-32>        | Creates a name for your I-SID list.                                      |
| <0-16777215>              | Specifies an I-SID value.                                                |
| <i>list WORD</i> <1-1024> | Specifies a list of I-SID values. For example, in the format 1,3,5,8-10. |

The following table defines parameters for the **isis accept** command.

| Variable                          | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>adv-rtr</i> < <i>x.xx.xx</i> > | Specifies a specific advertising BEB in which to apply the IS-IS accept policy to routes for a specific advertising BEB. <i>x.xx.xx</i> specifies an SPBM nickname.<br><br>The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.<br><br>The system requires an explicit filter to redistribute routes from a particular VSN. If the default global filter or the filter for a specific advertising BEB does not exist, the system does not redistribute the routes from the remote VSN. |
| <i>enable</i>                     | Enables the IS-IS accept policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <i>i-sid</i> <0-16777215>         | Configures the I-SID to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <i>isid-list WORD</i> <1-32>      | Configures a list of I-SIDs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <i>route-map WORD</i> <1-64>      | Specifies a route policy.<br>You must configure a route policy earlier in a separate procedure.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

The following table defines parameters for the **isis apply accept** command.

| Variable               | Value                              |
|------------------------|------------------------------------|
| <i>vrf WORD</i> <1-16> | Specifies a specific VRF instance. |

### Configure Inter-VRF IPv6 Accept Policies on VRFs

#### Before You Begin

- Enable IS-IS globally.
- Configure IPv6 Shortcuts. For more information, see [Configure SPBM IPv6 Shortcuts](#) on page 1608.
- Configure IPv6 IPVPN.
- Ensure that a route policy exists.
- Ensure that the VRFs exist.
- Configure a route-map.

## About This Task

Configure IPv6 IS-IS accept policies on a VRF to use inter-VRF accept policies in the SPB cloud. You can use IPv6 IS-IS accept policies to redistribute routes between different VRFs, including the global routing table (GRT). First you apply the filter, and then you match the filter based on the I-SID, which represents the Layer 3 VSN context.



### Note

- The **ipv6 isis apply accept [vrf WORD<1-16>]** command can disrupt traffic and cause temporary traffic loss. After you apply **ipv6 isis apply accept [vrf<1-16>]**, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply **ipv6 isis apply accept [vrf WORD<1-16>]** at the end.
- If you use the **ipv6 accept** command for inter-VRF routes based on the remote I-SID, the device accepts routes from other local VRFs to the current VRF, therefore if the accepted I-SID is configured on the local BEB, the device accepts its own IPv6 routes advertised under the accepted I-SID.
- If the route policy changes, you must reapply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

## Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



### Note

You cannot use the derived name for the untrusted (\*-unt) or unrestricted (\*-unr) segments of a Segmented VRF to enter this mode.

2. (Optional) If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IPv6 IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> {<0-16777215> | list WORD<1-1024>}
```



### Note

When creating an I-SID list, you can add I-SID entries until the maximum limit for supported Layer 3 I-SIDs is reached. The system truncates any additional I-SID entries. The maximum limit includes the I-SIDs for locally configured Layer 3 VSNs and the I-SIDs specified for IS-IS accept policy filters.

Use the command **show ip isid-list vrf WORD<1-16>** to view the list of truncated I-SIDs.

### 3. Configure an IPv6 IS-IS accept policy instance with a route policy.

Use one of the following options:

#### a. Configure an IPv6 IS-IS accept policy based on a specific advertising BEB:

```
ipv6 isis accept adv-rtr <x.xx.xx> [enable] [i-sid <0-16777215>]
[isid-list WORD<1-32>] [route-map WORD<1-64>]
```

#### b. Configure an IPv6 IS-IS accept policy based on a particular I-SID:

```
ipv6 isis accept i-sid <0-16777215> [enable] [route-map WORD<1-64>]
```

#### c. Configure an IPv6 IS-IS accept policy based on a particular I-SID list:

```
ipv6 isis accept isid-list WORD<1-32> [enable] [route-map WORD<1-64>]
```

### 4. Enable the configured IPv6 IS-IS accept policies:

```
ipv6 isis accept [adv-rtr <x.xx.xx>] [i-sid <0-16777215>] [isid-list
WORD<1-32>] enable
```

### 5. Exit to Global Configuration mode:

```
exit
```

### 6. Apply the IPv6 IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
ipv6 isis apply accept [vrf WORD<1-16>]
```

## Example

Configure Inter-VRF accept policies on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf red
Switch:1(router-vrf)#ipv6 isis accept i-sid 100 enable
Switch:1(router-vrf)#exit
Switch:1(config)#ipv6 isis apply accept vrf red
```

## Variable Definitions

The following table defines parameters for the **ip isid-list** command.



### Note

The I-SID lists created can be associated with both IPv4 or IPv6 routes.

| Variable                       | Value                                                                    |
|--------------------------------|--------------------------------------------------------------------------|
| <i>WORD&lt;1-32&gt;</i>        | Creates a name for your I-SID list.                                      |
| <i>&lt;0-16777215&gt;</i>      | Specifies an I-SID value.                                                |
| <i>list WORD&lt;1-1024&gt;</i> | Specifies a list of I-SID values. For example, in the format 1,3,5,8-10. |

The following table defines parameters for the **ipv6 isis accept** command.

| Variable                        | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>adv-rtr</i><br><x.xx.xx>     | Specifies the SPBM nickname for each advertising BEB to allow you to apply the IPv6 IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.<br><br><b>Note:</b><br>An IPv6 IS-IS accept policy that specifies the adv-rtr without an I-SID or I-SID list will filter routes coming from the I-SID on which the policy is configured and from the specified BEB. |
| <i>enable</i>                   | Enables an IPv6 IS-IS accept policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <i>i-sid</i><br><0-16777215>    | Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IPv6 IS-IS accept policy applies.<br>Use the parameter to apply a filter for routes from specific I-SIDs that represent the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter.<br>An I-SID value of 0 represents the global routing table (GRT).                                                                                                 |
| <i>isid-list</i><br>WORD<1-32>  | Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IPv6 IS-IS accept policy applies.<br>Use the parameter to apply a default filter for all routes from specific I-SIDs, that represent the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter.<br>An I-SID value of 0 represents the global routing table (GRT).                                                                             |
| <i>route-map</i> WORD<br><1-64> | Specifies a route policy.<br>You must configure a route policy earlier in a separate procedure.                                                                                                                                                                                                                                                                                                                                                                                                                                 |

The following table defines parameters for the **ipv6 isis apply accept** command.

| Variable              | Value                              |
|-----------------------|------------------------------------|
| <i>vrf</i> WORD<1-16> | Specifies a specific VRF instance. |

### Configure IS-IS Accept Policies

#### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply.
- Ensure that DvR is enabled on the switch before you configure an IS-IS accept policy with a backbone route policy, to accept host routes from the DvR backbone.

### About This Task

Use the following procedure to create and enable IS-IS accept policies to apply to routes from all Backbone Edge Bridges (BEBs) or to all routes from a specific BEB. Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

If DvR is enabled on your switch, and the switch is either a DvR Controller or a non-DvR BEB within the domain, you can configure IS-IS accept policies to accept specific host routes from the DvR backbone.

IS-IS accept policies are disabled by default.



#### Note

- The **isis apply accept [vrf WORD<1-16>]** command can disrupt traffic and cause temporary traffic loss. After you apply **isis apply accept [vrf <1-16>]**, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply **isis apply accept [vrf WORD<1-16>]** at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless the IS-IS accept policy was the last sequence in the configuration.
- The **isis apply accept [vrf WORD<1-16>]** command is not saved in the configuration file. If you use a saved configuration file for IS-IS accept policy configuration, you must apply the **isis apply accept [vrf WORD<1-16>]** command at the end.
- The number of unique Layer 3 VSN I-SIDs used on a BEB is limited to the number of VRFs supported on the switch. This includes the I-SID values used for Layer 3 VSNs and the I-SID values specified for the IS-IS accept policy filters, which can be configured using the **ip isid-list [ISID#]**, **accept i-sid <value>**, or **accept adv-rtr <isis nn> i-sid <value>** commands.

The switch supports 24 VRFs by default, so, in a default configuration, you cannot create an ip isid-list or accept policy with more than 24 unique I-SID entries. However, the configured VRFs take up an entry, so the formula to calculate the limit is: [24 VRF Limit – (currently configured VRFs)]. This gives the number of unique I-SIDs that can be used directly in the IS-IS accept policy filters, which you implement with the **ip isid-list** or **accept policy** command. The I-SIDs used for Layer 3 VSNs can be reused in IS-IS accept policy filters without affecting the limit.

If you increase the VRF scaling, you can create more Layer 3 VSNs. For more information about how to increase the number of supported VRFs, see [Configure the Maximum Number of VRFs](#) on page 4089. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Fabric Engine Release Notes*.

#### Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```



2. (Optional) If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
```

**Note**

When creating an I-SID list, you can add I-SID entries until the maximum limit for supported Layer 3 I-SIDs is reached. The system truncates any additional I-SID entries. The maximum limit includes the I-SIDs for locally configured Layer 3 VSNs and the I-SIDs specified for IS-IS accept policy filters.

Use the command **show ip isid-list vrf WORD<1-16>** to view the list of truncated I-SIDs.

3. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

4. Configure IS-IS accept policies with a route policy or a backbone route policy or a combination of both, to determine which routes the IS-IS accept policy applies to. Configure one of the following types of IS-IS accept policies:

- An IS-IS accept policy with only the route policy:

The IS-IS routes are selectively accepted based on the route policy. Since the backbone route policy is not configured, all host routes from the DvR backbone are denied.

If you do not configure a route policy, by default, all IS-IS routes are accepted.

- An IS-IS accept policy with only the backbone route policy:

The DvR host routes from the DvR backbone are selectively accepted based on the backbone route policy. Since the route policy is not configured, all IS-IS host routes are accepted.

If you do not configure a backbone route policy, all host routes from the DvR backbone are denied.

- An IS-IS accept policy with both route policy and backbone route policy:

IS-IS routes are selectively accepted based on the route policy and host routes from the DvR backbone are selectively accepted based on the backbone route policy.

5. Configure an IS-IS accept policy instance with a route policy.

Use one of the following options:

- a. Create an IS-IS accept policy instance to apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>]
```

- b. Create an IS-IS accept policy instance to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> [i-sid <1-16777215>][isid-list WORD <1-32>]
```

- c. Specify an IS-IS route policy to apply to routes from all BEBs:

```
accept route-map WORD<1-64>
```

- d. Specify an IS-IS route policy to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
```

- e. Enable an IS-IS route accept instance:

```
accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list WORD<1-32>]
```

6. Configure an IS-IS accept policy instance with a backbone route policy to accept host routes from the DvR backbone:



**Note**

IS-IS accept policies typically apply to all IS-IS routes. However, to accept DvR host routes from the DvR backbone, you must explicitly configure the IS-IS accept policy with a backbone route policy.

Use one of the following options:

- a. Create the default IS-IS accept policy instance to accept host routes from the DvR backbone:

```
accept backbone-route-map WORD <1-64>
```

- b. Create an IS-IS accept policy instance to accept host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-route-map WORD<1-64>
```

- c. Create an IS-IS accept policy instance to accept host routes from the DvR backbone and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> backbone-route-map WORD <1-64>
```

7. Configure an IS-IS accept policy with both route policy and backbone route policy, to selectively accept IS-IS routes as well as host routes from the DvR backbone.

- a. Create the default IS-IS accept policy instance with a route policy to accept IS-IS routes and a backbone route policy to accept host routes from the DvR backbone:

```
accept route-map WORD<1-32> backbone-route-map WORD <1-64>
```

- b. Create an accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map WORD<1-32> backbone-route-map WORD<1-64>
```

- c. Create an IS-IS accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to a specific advertising BEB:
 

```
accept adv-rtr <x.xx.xx> route-map WORD<1-32> backbone-route-map
WORD <1-64>
```
8. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:
 

```
isis apply accept [vrf WORD <1-16>]
```

### Examples

Configure an I-SID based IS-IS accept policy with the route policy test:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map test 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 101
Switch:1(config-isis)#accept i-sid 101 route-map test
Switch:1(config-isis)#accept i-sid 101 enable
Switch:1#exit
Switch:1(config)#isis apply accept
```

The following examples show the configuration of an IS-IS accept policy to accept host routes from the DvR backbone.

To accept host routes from the DvR backbone, you must configure a backbone route policy and apply it to the IS-IS accept policy.

1. Configure a route policy for DvR:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map dvrmap1 1
Switch:1(route-map)#enable
```

2. Configure an IS-IS accept policy for I-SID 10, and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap1
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
```

### OR

Configure the default accept policy for IS-IS and DvR, and apply the route policy as a backbone route policy:

```
Switch:1(config)#route-map isismap1 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept route-map isismap1 backbone-route-map dvrmap1
```

## 3. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
Switch:1(config)#exit
```

## 4. Verify the configuration:

```
Switch:1#show ip isis accept

=====
 Isis Accept - GlobalRouter
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE

- 10 - TRUE dvrmap1
- - - isismap1 dvrmap1

2 out of 2 Total Num of Isis Accept Policies displayed
```

Configure an IS-IS accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

## 1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

## 2. Create the route policy dvrmap2 to match the IP prefix list:

```
Switch:1(config)#route-map dvrmap2 1
Switch:1(route-map)#match network listPrefix
Switch:1(route-map)#enable
```

## 3. Create an IS-IS accept policy with I-SID 10 and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
```

## 4. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

The previous command causes IS-IS to accept all routes with I-SID 10. To deny IS-IS routes and accept only DvR host routes, you can configure an additional IS-IS route policy as follows:

```
Switch:1(config)#route-map isismap2 1
Switch:1(route-map)#no permit
Switch:1(route-map)#enable

Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 route-map isismap2 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
```

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

##### 5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1#show ip isis accept

=====
Isis Accept - GlobalRouter
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY

- 10 - TRUE isismap2 dvrmap2

1 out of 1 Total Num of Isis Accept Policies displayed
```

The following examples show the configuration of IS-IS accept policies for a specific VRF instance.

Configure IS-IS accept policies to accept host routes from the DVR backbone, for a specific VRF instance.

1. In the VRF green context, configure the route policy dvrmap3 for DVR:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap3 1
Switch:1(router-vrf-routemap)#enable
```

2. Use one of the following options to configure an IS-IS accept policy, and apply the route policy as a backbone route policy:

Configure an IS-IS accept policy for a specific advertising BEB with nickname 1.11.11:

```
Switch:1(router-vrf-routemap)#isis accept adv-rtr 1.11.11 backbone-route-map dvrmap3
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#isis accept adv-rtr 1.11.11 enable
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY

1.11.11 - - TRUE dvrmap3

1 out of 1 Total Num of Isis Accept Policies displayed
Switch:1(config)#show ip isis accept vrfids 2

=====
Isis Accept - VRF green
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY

1.11.11 - - TRUE dvrmap3
```

```
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for I-SID 10:

```
Switch:1(router-vrf)#isis accept i-sid 10 backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
 Isis Accept - VRF green
=====
```

| ADV_RTR | I-SID | ISID-LIST | ENABLE POLICY | BACKBONE POLICY |
|---------|-------|-----------|---------------|-----------------|
| -       | 10    | -         | TRUE          | dvrmap3         |

```
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for the I-SID list listisids:

```
Switch:1(router-vrf)#isis accept isid-list listisids backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
 Isis Accept - VRF green
=====
```

| ADV_RTR | I-SID | ISID-LIST | ENABLE POLICY | BACKBONE POLICY |
|---------|-------|-----------|---------------|-----------------|
| -       | 10    | listisids | TRUE          | dvrmap3         |

```
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure the default accept policy for IS-IS and DvR:

```
Switch:1(router-vrf)#route-map isismap3 1
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#isis accept route-map isismap3 backbone-route-map dvrmap3
Switch:1(router-vrf)#
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
 Isis Accept - VRF green
=====
```

| ADV_RTR | I-SID | ISID-LIST | ENABLE POLICY | BACKBONE POLICY |
|---------|-------|-----------|---------------|-----------------|
| -       | -     | -         | TRUE isismap3 | dvrmap3         |

```
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure the default accept policy for DvR:

```
Switch:1(router-vrf)#isis accept backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
 Isis Accept - VRF green
=====
```

```

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE

- - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

```

Configure an accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24

```

2. For a specific VRF instance, create a route policy to match the IP prefix list:

```

Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap4 1
Switch:1(router-vrf-route-map)#match network listPrefix
Switch:1(router-vrf-route-map)#enable
Switch:1(router-vrf-route-map)#exit
Switch:1(router-vrf)#

```

3. Create an IS-IS accept policy with I-SID 10, and apply the route policy as the backbone route policy:

```

Switch:1(router-vrf)#accept i-sid 10 backbone-route-map dvrmap4
Switch:1(router-vrf)#accept i-sid 10 enable

```

4. Apply the IS-IS accept policy:

```

Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept

```

5. Verify the configuration:

```

Switch:1(config)#exit
Switch:1(router-vrf)#show ip isis accept vrf green

=====
 Isis Accept - VRF green
=====

ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE

- - - TRUE dvrmap4
1 out of 1 Total Num of Isis Accept Policies displayed

```

## Variable definitions

The following table defines parameters for the **ip isid-list** command.

| Variable                  | Value                                                                    |
|---------------------------|--------------------------------------------------------------------------|
| <i>WORD</i> <1-32>        | Creates a name for your I-SID list.                                      |
| <1-16777215>              | Specifies an I-SID number.                                               |
| <i>list WORD</i> <1-1024> | Specifies a list of I-SID values. For example, in the format 1,3,5,8-10. |

The following table defines parameters for the **accept** command.

| Variable                             | Value                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>adv-rtr</i> <x.xx.xx>             | Specifies the SPBM nickname for each advertising BEB to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.                                                                                                                                        |
| <i>backbone-route-map</i> WORD<1-64> | Specifies the DvR backbone route map.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <i>enable</i>                        | Enables an IS-IS accept policy.                                                                                                                                                                                                                                                                                                                                                                                                               |
| <i>i-sid</i><br><1-16777215>         | Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies.<br>Use the parameter to apply a filter for routes from specific I-SIDs that represent the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter.<br>An I-SID value of 0 represents the global routing table (GRT).                    |
| <i>isid-list</i><br>WORD<1-32>       | Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IS-IS accept policy applies.<br>Use the parameter to apply a default filter for all routes from specific I-SIDs that represent the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter.<br>An I-SID value of 0 represents the global routing table (GRT). |
| <i>route-map</i><br>WORD<1-64>       | Specifies a route policy by name.<br>You must configure the route policy earlier in a separate procedure.                                                                                                                                                                                                                                                                                                                                     |

The following table defines parameters for the **isis apply accept** command.

| Variable              | Value                              |
|-----------------------|------------------------------------|
| <i>vrf</i> WORD<1-16> | Specifies a specific VRF instance. |

#### View IS-IS Accept Policy Information

Use the following procedure to view IS-IS accept policy information on the switch.

#### Procedure

1. Display IS-IS accept policy information:

```
show ip isis accept [vrf WORD<1-16>][vrfids WORD<0-512>]
```

2. Display I-SID list information:

```
show ip isid-list [vrf WORD<1-16>][vrfids WORD<0-512>][WORD<1-32>]
```



## 3. Display route information:

```
show ip route [vrf WORD<1-16>]
```

The NH VRF/ISID column displays the I-SID for inter-Virtual Services Network (VSN) routes redistributed with IS-IS accept policies, only if the I-SID redistributed does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays. If the I-SID is 0, the column represents and displays as the GlobalRouter.

The existing IS-IS routes for Layer 3 VSNs continue to display as the VRF name of the IP VSN.

## 4. Display the SPBM IP unicast Forwarding Information Base (FIB):

```
show isis spbm ip-unicast-fib [all] [id <1-16777215>] [spbm-nh-as-mac]
[home|remote]
```

## Example

View IS-IS accept policy information:

```
Switch:1#show ip route vrf test
=====
IP Route - VRF test
=====
NH INTER
DST MASK NEXT NH INTER
VRF/ISID COST FACE PROT AGE TYPE PRF ACCESS

1.1.1.5 255.255.255.255 1.1.1.5 GlobalRouter 0 0 ISIS 0 IB 200 Trusted
1.1.1.13 255.255.255.255 Switch13 GRT 10 1000 ISIS 0 IBSV 7 Trusted
1.1.1.200 255.255.255.255 Switch200 GRT 10 1000 ISIS 0 IBSV 7 Trusted
5.7.1.0 255.255.255.0 5.7.1.1 - 1 7 LOC 0 DB 0 Trusted
13.7.1.0 255.255.255.0 Switch13 GlobalRouter 10 1000 ISIS 0 IBSV 7 Trusted
100.0.0.0 255.255.255.0 100.0.0.1 GlobalRouter 0 100 ISIS 0 IB 200 Trusted
111.1.1.0 255.255.255.0 111.1.1.1 hub 0 111 ISIS 0 IB 200 Trusted

Switch:1(config)#show isis spbm ip-unicast-fib
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF DEST OUTGOING SPBM PREFIX IP ROUTE
VRF ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE AREA AREA-NAME

GRT - 101 1.1.1.13/32 Switch13 1000 1/7 10 44 7 HOME area-9.00.02
GRT - 101 1.1.1.13/32 Switch13 1001 1/7 10 44 7 HOME area-9.00.02

Home : Total number of SPBM IP-UNICAST FIB entries 2
Remote: Total number of SPBM IP-UNICAST FIB entries 0

Switch:1(config)#show ip isid-list test
=====
IP ISID LIST
=====
List Name I-SID VRF

test 1 GlobalRouter
 3 GlobalRouter
 4 GlobalRouter
 5 GlobalRouter
 10 GlobalRouter
 22 GlobalRouter

All 6 out of 6 Total Num of Isid Lists displayed

Switch:1(router-vrf)#show ip isid-list vrf red
=====
IP ISID LIST red
=====
List Name I-SID VRF

test1 11 1
 12 1
 13 1
 14 1
 15 1
```

## Variable Definitions

The following table defines parameters for the **show ip isis accept** command.

| Variable                        | Value                                                         |
|---------------------------------|---------------------------------------------------------------|
| <i>vrf WORD&lt;1-16&gt;</i>     | Displays I-SID list information for a particular VRF by name. |
| <i>vrfids WORD&lt;0-512&gt;</i> | Displays I-SID list information for a particular VRF ID.      |

The following table defines parameters for the **show ip isid-list** command.

| Variable                        | Value                                                             |
|---------------------------------|-------------------------------------------------------------------|
| <i>vrf WORD&lt;1-16&gt;</i>     | Displays I-SID list information for a particular VRF by name.     |
| <i>vrfids WORD&lt;0-512&gt;</i> | Displays I-SID list information for a particular VRF ID.          |
| <i>WORD&lt;1-32&gt;</i>         | Displays I-SID list information for a particular I-SID list name. |

The following table defines parameters for the **show ip route** command.

| Variable                    | Value                                                         |
|-----------------------------|---------------------------------------------------------------|
| <i>vrf WORD&lt;1-16&gt;</i> | Displays I-SID list information for a particular VRF by name. |

The following table defines parameters for the **show isis spbm ip-unicast-fib** command.

| Variable                     | Value                                                                                             |
|------------------------------|---------------------------------------------------------------------------------------------------|
| <i>all</i>                   | Displays all IS-IS SPBM IP unicast Forwarding Information Base (FIB) information.                 |
| <i>home</i>                  | Displays the IS-IS SPBM IP unicast FIB information that the system configures in the home area.   |
| <i>id &lt;1-16777215&gt;</i> | Displays IS-IS SPBM IP unicast FIB information by I-SID ID.                                       |
| <i>remote</i>                | Displays the IS-IS SPBM IP unicast FIB information that the system configures in the remote area. |
| <i>spbm-nh-as-mac</i>        | Displays the next hop BMAC of the IP unicast FIB entry.                                           |

## Configure IPv6 IS-IS Accept Policies

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Configure IPv6 Shortcuts. For more information, see [Configure SPBM IPv6 Shortcuts](#) on page 1608.
- You must configure a route-map.

## About This Task

Perform the following procedure to create and enable IPv6 IS-IS accept policies based on a particular Backbone Edge Bridge (BEB), I-SID, or I-SID list. IPv6 IS-IS accept policies filter incoming IS-IS routes that the device receives over the SPBM cloud. IPv6 IS-IS accept policies apply to incoming traffic and determine whether to add the route to the routing table.

IPv6 IS-IS accept policies are disabled by default.



### Note

- The I-SID lists created can be associated with both IPv4 or IPv6 routes.
- The **ipv6 isis apply accept [vrf WORD<1-16>]** command can disrupt traffic and cause temporary traffic loss. After you apply **ipv6 isis apply accept [vrf <1-16>]**, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply **ipv6 isis apply accept [vrf WORD<1-16>]** at the end.
- If the route policy associated with an accept policy changes, you must reapply the IPv6 IS-IS accept policy, unless the IPv6 IS-IS accept policy was the last sequence in the configuration.
- The **ipv6 isis apply accept [vrf WORD<1-16>]** command is not saved in the configuration file. If you use a saved configuration file for IPv6 IS-IS accept policy configuration, you must apply the **ipv6 isis apply accept [vrf WORD<1-16>]** command at the end.

The number of unique Layer 3 VSN I-SIDs used on a BEB is limited to the number of VRFs supported on the switch. This includes the I-SID values used for Layer 3 VSNs and the I-SID values specified for the IPv6 IS-IS accept policy filters.

The switch supports 24 VRFs by default, so, in a default configuration, you cannot create an I-SID list or accept policy with more than 24 unique I-SID entries. However, the configured VRFs take up an entry, so the formula to calculate the limit is: [24 VRF Limit – (currently configured VRFs)]. This gives the number of unique I-SIDs that can be used directly in the IPv6 IS-IS accept policy filters, which you implement with the **ip isid-list** or **ipv6 accept** command. The I-SIDs used for Layer 3 VSNs can be reused in IPv6 IS-IS accept policy filters without affecting the limit.

If you increase the VRF scaling, you can create more Layer 3 VSNs. For more information about how to increase the number of supported VRFs, see [Configure the Maximum Number of VRFs](#) on page 4089. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Fabric Engine Release Notes*.

## Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. (Optional) If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IPv6 IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> {<1-16777215> | list WORD<1-1024>}
```



### Note

When creating an I-SID list, you can add I-SID entries until the maximum limit for supported Layer 3 I-SIDs is reached. The system truncates any additional I-SID entries. The maximum limit includes the I-SIDs for locally configured Layer 3 VSNs and the I-SIDs specified for IS-IS accept policy filters.

Use the command **show ip isid-list vrf WORD<1-16>** to view the list of truncated I-SIDs.

3. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



### Note

You cannot use the derived name for the untrusted (\*-unt) or unrestricted (\*-unr) segments of a Segmented VRF to enter this mode.

4. Configure an IPv6 IS-IS accept policy instance with a route policy.

Use one of the following options:

- a. Configure an IPv6 IS-IS accept policy based on a specific advertising BEB:

```
ipv6 isis accept adv-rtr <x.xx.xx> [enable] [i-sid <0-16777215>]
[isid-list WORD<1-32>] [[route-map WORD<1-64>]
```

- b. Configure an IPv6 IS-IS accept policy based on a particular I-SID:

```
ipv6 isis accept i-sid <0-16777215> [enable] [route-map WORD<1-64>]
```

- c. Configure an IPv6 IS-IS accept policy based on a particular I-SID list:

```
ipv6 isis accept isid-list WORD<1-32> [enable] [route-map WORD<1-64>]
```

- d. Specify a particular route-map to use for all IS-IS routes from all BEBs unless a more specific filter exists for the advertising BEB.:

```
ipv6 isis accept route-map WORD<1-64>
```

5. Enable the configured IPv6 IS-IS accept policies:

```
ipv6 isis accept [adv-rtr <x.xx.xx>] [i-sid <0-16777215>] [isid-list
WORD<1-32>] enable
```

## 6. Exit to Global Configuration mode:

```
exit
```

## 7. Apply the IPv6 IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
ipv6 isis apply accept [vrf WORD <1-16>]
```

**Example**

Configure an IPv6 IS-IS accept policy based on a particular I-SID:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrftest
Switch:1(config-isis)#ipv6 isis accept i-sid 101 route-map test
Switch:1(config-isis)#ipv6 isis accept i-sid 101 enable
Switch:1#exit
Switch:1(config)#ipv6 isis apply accept
```

**Variable Definitions**

The following table defines parameters for the **ip isid-list** command.

**Note**

The I-SID lists created can be associated with both IPv4 or IPv6 routes.

| Variable                  | Value                                                                    |
|---------------------------|--------------------------------------------------------------------------|
| <i>WORD</i> <1-32>        | Creates a name for your I-SID list.                                      |
| <1-16777215>              | Specifies an I-SID number.                                               |
| <i>list WORD</i> <1-1024> | Specifies a list of I-SID values. For example, in the format 1,3,5,8-10. |

The following table defines parameters for the **ipv6 isis accept** command.

| Variable                    | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>adv-rtr</i><br><x.xx.xx> | Specifies the SPBM nickname for each advertising BEB to allow you to apply the IPv6 IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.<br><br><b>Note:</b><br>An IPv6 IS-IS accept policy that specifies the adv-rtr without an I-SID or I-SID list will filter routes coming from the I-SID on which the policy is configured and from the specified BEB. |
| <i>enable</i>               | Enables an IPv6 IS-IS accept policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

| Variable                       | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>i-sid</i><br><0-16777215>   | Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IPv6 IS-IS accept policy applies.<br><br>Use the parameter to apply a filter for routes from specific I-SIDs that represent the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter.<br><br>An I-SID value of 0 represents the global routing table (GRT).                    |
| <i>isid-list</i><br>WORD<1-32> | Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IPv6 IS-IS accept policy applies.<br><br>Use the parameter to apply a default filter for all routes from specific I-SIDs that represent the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter.<br><br>An I-SID value of 0 represents the global routing table (GRT). |
| <i>route-map</i><br>WORD<1-64> | Specifies a route policy by name.<br><br>You must configure the route policy earlier in a separate procedure.                                                                                                                                                                                                                                                                                                                                              |

The following table defines parameters for the **ipv6 isis apply accept** command.

| Variable              | Value                     |
|-----------------------|---------------------------|
| <i>vrf</i> WORD<1-16> | Specifies a VRF instance. |

### Display IPv6 IS-IS Accept Policy Information

Perform the following procedure to view IPv6 IS-IS accept policy information on the switch.

#### Procedure

1. Enter Privileged EXEC mode:  
enable
2. Display IPv6 IS-IS accept policy information:  
show ipv6 isis accept [vrf WORD<1-16>] [vrffids WORD<0-512>]

#### Example

Display IPv6 IS-IS accept policy information for vrfRED:

```
Switch:1>enable
Switch:1#show ipv6 isis accept vrf vrfRED
=====
 Isis Accept - VRF vrfRED
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY

1.11.11 1001 - TRUE

1 out of 1 Total Num of Isis Accept Policies displayed
```

## IP Shortcuts configuration using EDM

This section provides procedures to configure IP Shortcuts using EDM.

### Configure SPBM IP Shortcuts

#### Before You Begin

- Enable IPv4 Shortcuts before you enable IPv6 Shortcuts. IPv6 Shortcuts depend on IPv4 Shortcuts for some functions.
- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, create the Customer VLANs, add slots/ports, and add the IP addresses and network masks.

#### About This Task

To enable IP shortcuts on the BEBs, you can configure a circuitless IP address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.



#### Note

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address.

#### Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **IS-IS**.
3. From the **Globals** tab, in **IpSourceAddress**, specify the CLIP interface to use as the source address for SBPM IP Shortcuts.



#### Note

For IPv6 Shortcuts, select **ipv6** in **Ipv6SourceAddressType**, and then use **Ipv6SourceAddress** to specify the CLIPv6 interface to use as the source address for SBPM IPv6 Shortcuts.

4. Select **Apply**.
5. In the navigation pane, expand **Configuration > Fabric > SPBM**.
6. Select the **SPBM** tab.
7. In **IpShortcut**, select **enable**.



#### Note

For IPv6 Shortcuts, select **enable** in **Ipv6Shortcut**.

8. Select **Apply**.
9. In the navigation pane, expand **Configuration > IP**.
10. Select **Policy**.

11. Select the **Route Redistribution** tab.
12. Select **Insert** to identify routes on the local switch to be announced into the SPBM network.
13. Specify the source protocols to redistribute into IS-IS.  
In **Protocol**, ensure you specify **isis** as the destination protocol.
14. Select **Insert**.

### Configure IPv4 IS-IS Redistribution

#### Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (\*-unt) or unrestricted (\*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

#### About This Task

Use this procedure to configure IS-IS redistribution. In a non-default VRF, just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute a protocol into IS-IS, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

#### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. In **SrcVrfId**, select the source VRF ID.
6. In **RouteSource**, select the protocol for route redistribution.
7. Complete the remaining optional fields as required.
8. Select **enable**.
9. Select **Insert**.

#### IS-IS Redistribute Field Descriptions

Use the data in the following table to use the **IS-IS Redistribute** tab.

| Name            | Description                                                                                   |
|-----------------|-----------------------------------------------------------------------------------------------|
| <b>DstVrfId</b> | Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution. |
| <b>Protocol</b> | Specifies the protocols that receive the redistributed routes.                                |



| Name               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SrcVrfId</b>    | Specifies the source VRF ID used in the redistribution. For IS-IS, the source VRF ID must be the same as the destination VRF ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>RouteSource</b> | Specifies the source protocol for the route redistribution entry.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Enable</b>      | Enables or disables a redistribution entry. The default is disable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>RoutePolicy</b> | Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Metric</b>      | Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. Use a value that is consistent with the destination protocol.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>MetricType</b>  | Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Subnets</b>     | Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>SetTag</b>      | <p>Tag routes that BGP or OSPF redistribute into IS-IS. You can then configure router filtering to prevent tagged IS-IS routes from being redistributed back to BGP or OSPF.</p> <p>To use IS-IS Route Tagging on GRT IS-IS routes, you must also configure the metric-type as external. If you want to use IS-IS tags on GRT as internal routes, all Fabric nodes must be above a minimum software version. Any switch in the SPB Fabric that runs earlier software versions triggers an exception if you use metric type internal. To ensure this does not occur, if you attempt to configure a tag and the metric-type is not external, the switch reminds you to upgrade the software on all devices. You must ensure all devices in the network run the minimum required software.</p> |

### Configure IPv6 IS-IS Redistribution

Use this procedure to configure IS-IS redistribution for IPv6. In the Virtual Routing and Forwarding (VRF), just like in the Global Router, the IPv6 routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: v6direct, v6static, RIPng, OSPFv3, or BGPv6, within

the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.



#### Note

RIPng is supported only on the Global Router.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IPv6 VPN reachability information, the IPv6 routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

### Before You Begin

Change the VRF instance as required to configure IPv6 IS-IS redistribution on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. Complete the fields as required.
6. Select **Insert**.
7. Select **Apply**.

### Redistribute Field Descriptions

Use the data in the following table to use the **Redistribute** tab.

| Name               | Description                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>DstVrflid</b>   | Specifies the destination Virtual Routing and Forwarding (VRF) ID used in redistribution.                                               |
| <b>Protocol</b>    | Specifies the protocols that receive the redistributed routes.                                                                          |
| <b>SrcVrflid</b>   | Specifies the source Virtual Routing and Forwarding (VRF) ID used in redistribution.                                                    |
| <b>RouteSource</b> | Specifies the source protocol for the route redistribution entry.                                                                       |
| <b>Enable</b>      | Enables or disables a redistribution entry. The default is disabled.                                                                    |
| <b>RoutePolicy</b> | Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain. |

| Name              | Description                                                                                                                                                                                                                                                                                     |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Metric</b>     | Specifies the metric for the redistributed route. The default value is 0. Use a value that is consistent with the destination protocol.                                                                                                                                                         |
| <b>MetricType</b> | Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2. |

### Apply IPv4 IS-IS Accept Policies Globally

#### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Ensure the IP IS-IS filter exists.

#### About This Task

Apply IS-IS accept policies globally. Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

After you apply the IS-IS accept filters, the device removes and re-adds all routes with updated filters.

IS-IS accept policies are disabled by default.



#### Note

- After you apply IS-IS accept policies globally the application can disrupt traffic and cause temporary traffic loss. After you configure the IS-IS accept policies value to **Apply**, the device reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply IS-IS accept policies globally at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

#### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Accept Global** tab.
4. Select a name from the list or enter name in the **DefaultPolicyName** field, to specify the route policy name for the default filter.
5. Select **Apply** to apply the default policy.

## Accept Global Field Descriptions

Use the data in the following table to use the **Accept Global** tab.

| Name                             | Description                                                                                                                                                                                                                                                      |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DefaultPolicyName</b>         | Specifies the route policy name for the default filter.                                                                                                                                                                                                          |
| <b>DefaultBackbonePolicyName</b> | Specifies the backbone host route policy name for the default filter.                                                                                                                                                                                            |
| <b>Apply</b>                     | Applies the default policy when you configure the field to apply. The device only activates the default policy if the route map (the default policy name) has a value. If you do not select apply, the device takes no action. The GRT always returns no action. |

### Configure an IPv4 IS-IS Accept Policy for a Specific Advertising BEB

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB). Specify the SPBM nickname and the IS-IS accept policy name to allow you to apply the IS-IS accept policy.

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.



#### Note

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply.

### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Accept Nick Name** tab.
4. Select **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Select **Insert**.

## Accept Nick Name Field Descriptions

Use the data in the following table to use the **Accept Nick Name** tab.

| Name                      | Description                                                                                                                                                                                                                                                                                                                          |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AdvertisingRtr</b>     | Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.<br><br>The value is 2.5 bytes in the format <x.xx.xx>. |
| <b>Enable</b>             | Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.                                                                                                                                                                                                        |
| <b>PolicyName</b>         | Specifies a route policy.<br>You must configure a policy earlier in a separate procedure.                                                                                                                                                                                                                                            |
| <b>BackbonePolicyName</b> | Specifies the route policy for the backbone routes.<br>You must configure a policy earlier in a separate procedure.                                                                                                                                                                                                                  |

### Configure an IS-IS Accept Policy to Apply for a Specific I-SID

Configure an IS-IS accept policy for a specific I-SID number to represent a local or remote Layer 3 VSN, which allows the system to redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).



#### Note

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply.

### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Accept Isid** tab.
4. Select **Insert**.
5. In the **Isid** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.

7. In the **PolicyName** field, specify the route-map name.
8. Select **Insert**.

### Accept Isid Field Descriptions

Use the data in the following table to use the **Accept Isid** tab.

| Name                      | Description                                                                                                                                                                                                                                                                                                                   |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Isid</b>               | Configures a specific I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies.<br>Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter.<br>An I-SID value of 0 represents the global routing table (GRT). |
| <b>Enable</b>             | Enables or disables the I-SID entry. You must enable the value to filter. The default is disabled.                                                                                                                                                                                                                            |
| <b>PolicyName</b>         | Specifies the route map name. You must configure a policy earlier in a separate procedure.                                                                                                                                                                                                                                    |
| <b>BackbonePolicyName</b> | Specifies the backbone route map name. You must configure a policy earlier in a separate procedure.                                                                                                                                                                                                                           |

### Configure an IPv4 IS-IS Accept Policy for a Specific Advertising BEB and I-SID

Configure a specific advertising Backbone Edge Bridge (BEB) with a specific I-SID to apply the IS-IS accept policy to routes for a specific advertising BEB.



#### Note

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply.

### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Accept Nick-Name Isid** tab.
4. Select **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Isid** field, specify an I-SID number.

7. Select **enable** in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Select **Insert**.

#### Accept Nick-Name Isid Field Descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid** tab.

| Name                      | Description                                                                                                                                                            |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AdvertisingRtr</b>     | Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.<br>The value is 2.5 bytes in the format <x.xx.xx>. |
| <b>Isid</b>               | Specifies an I-SID used to filter. The value 0 is used for the Global Router.                                                                                          |
| <b>Enable</b>             | Enables or disables the I-SID entry. The default is disabled.                                                                                                          |
| <b>PolicyName</b>         | Specifies the route policy name. You must configure a policy earlier in a separate procedure.                                                                          |
| <b>BackBonePolicyName</b> | Specifies the backbone route policy name. You must configure a policy earlier in a separate procedure.                                                                 |

#### Configure an I-SID List for an IPv4 IS-IS Accept Policy

Configures a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. After you create the list of I-SID numbers, you must then create, configure, and enable the IS-IS accept policy.



#### Note

When creating an I-SID list, you can add I-SID entries until the maximum limit for supported Layer 3 I-SIDs is reached. The system truncates any additional I-SID entries. The maximum limit includes the I-SIDs for locally configured Layer 3 VSNs and the I-SIDs specified for IS-IS accept policy filters.  
Refresh the EDM tab to view the actual list of I-SIDs in the I-SID list.

#### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.

#### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Isid-List** tab.
4. Select **Insert**.
5. In the **Name** field, specify a name for the I-SID list.

6. Select **Isid** or **Isid-List**.
7. Specify an I-SID number or a list of I-SID numbers.
8. Select **Insert**.

### Isid-List Field Descriptions

Use the data in the following table to use the **Isid-List** tab.

| Name                     | Description                                                                                                                                                                                                            |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>              | Specifies the name of the I-SID list.                                                                                                                                                                                  |
| <b>Isid or Isid-List</b> | Specifies that you either want to add a particular I-SID or a list of I-SID numbers.                                                                                                                                   |
| <b>Isid</b>              | Specifies a particular I-SID number or a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies.<br>An I-SID value of 0 represents the global routing table (GRT). |

### Configure an IPv4 IS-IS Accept Policy for a Specific I-SID List

Configure an IS-IS accept policy for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.



#### Note

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply.

### Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Accept Isid-List** tab.
4. Select **Insert**.
5. In the **Name** field, specify the I-SID list name.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Select **Insert**.



## Accept Isid-List Field Descriptions

Use the data in the following table to use **Accept Isid-List** tab.

| Name                      | Description                                                                                             |
|---------------------------|---------------------------------------------------------------------------------------------------------|
| <b>Name</b>               | Specifies the name of I-SID list.                                                                       |
| <b>Enable</b>             | Enables or disables the I-SID list entry. The value must be enabled to filter. The default is disabled. |
| <b>PolicyName</b>         | Specifies the route policy name.                                                                        |
| <b>BackBonePolicyName</b> | Specifies the backbone route policy name.                                                               |

### *Configure an IPv4 IS-IS Accept Policy for a Specific Advertising BEB and I-SID-list*

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB) for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.



#### Note

If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

## Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply.

## About This Task

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

## Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IS-IS Policies**.
3. Select the **Accept Nick-Name Isid-List** tab.
4. Select **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Name** field, specify an I-SID list name.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Select **Insert**.

## Accept Nick-Name Isid-List Field Descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid-List** tab.

| Name                      | Description                                                                                                                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AdvertisingRtr</b>     | Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter is present the device applies the specific filter.<br><br>The value is 2.5 bytes in the format <x.xx.xx>. |
| <b>Name</b>               | Specifies the name of the I-SID list used to filter.                                                                                                                                                                                                                                                  |
| <b>Enable</b>             | Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.                                                                                                                                                                         |
| <b>PolicyName</b>         | Specifies a route policy name.                                                                                                                                                                                                                                                                        |
| <b>BackBonePolicyName</b> | Specifies a backbone route policy name.                                                                                                                                                                                                                                                               |

### Apply IPv6 IS-IS Accept Policies Globally

Apply IPv6 IS-IS accept policies globally. Use IPv6 IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud.

After you apply the IPv6 IS-IS accept policy filters, the device removes and re-adds all IPv6 routes with updated filters.

IPv6 IS-IS accept policies are disabled by default.



#### Note

- After you apply IPv6 IS-IS accept policies globally the application can disrupt traffic and cause temporary traffic loss. After you configure the IPv6 IS-IS accept policies value to **Apply**, the device reapplies the accept policies, which deletes all of the IPv6 IS-IS routes, and adds the IPv6 IS-IS routes again. You should make all the relevant accept policy changes, and then apply IPv6 IS-IS accept policies globally at the end.
- If the route policy changes, you must reapply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Ensure the IPv6 IS-IS filter exists.
- Change the VRF instance as required to apply IPv6 IS-IS accept policies on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

## Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Accept Global** tab.
4. (Optional) Select a name from the list or enter name in the **DefaultPolicyName** field, to specify the route policy name for the default filter.
5. Select **Apply** to apply the default policy.
6. Select **Apply**.

## Accept Global Field Descriptions

Use the data in the following table to configure the **Accept Global** tab.

| Name                             | Description                                                                                                                                                                                                                                      |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DefaultPolicyName</b>         | Specifies the route policy name for the default filter.                                                                                                                                                                                          |
| <b>Apply</b>                     | Applies the default policy when you select apply. The device only activates the default policy if the route map (the default policy name) has a value. If you do not select apply, the device takes no action. The GRT always returns no action. |
| <b>NickNameTableSize</b>         | Shows the IPv6 IS-IS In Filter nickname table size.                                                                                                                                                                                              |
| <b>IsidTableSize</b>             | Shows the IPv6 IS-IS In Filter I-SID table size.                                                                                                                                                                                                 |
| <b>NickNameIsidTableSize</b>     | Shows the IPv6 IS-IS In Filter nickname I-SID table size.                                                                                                                                                                                        |
| <b>IsidListTableSize</b>         | Shows the IPv6 IS-IS In Filter I-SID List table size.                                                                                                                                                                                            |
| <b>NickNameIsidListTableSize</b> | Shows the IPv6 IS-IS In Filter nickname I-SID List table size.                                                                                                                                                                                   |

### Configure an IPv6 IS-IS Accept Policy for a specific Advertising BEB

Configure an IPv6 IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB). Specify the SPBM nickname and the IS-IS accept policy name to allow you to apply the IPv6 IS-IS accept policy.

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.



#### Note

If the route policy changes, you must re-apply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map.
- Change the VRF instance as required to configure an IPv6 IS-IS accept policy for a particular advertising BEB on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Accept Nick Name** tab.
4. Select **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. Select **Enable** to apply the filter.
7. (Optional) In the **PolicyName** field, specify the route-map name.
8. Select **Insert**.

### Accept Nick Name Field Descriptions

Use the data in the following table to configure the **Accept Nick Name** tab.

| Name                  | Description                                                                                                                                                                                                                                                      |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AdvertisingRtr</b> | Specifies the SPBM nickname to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter. |
| <b>Enable</b>         | Enables the SPBM nickname advertising router entry. The default is disabled.                                                                                                                                                                                     |
| <b>PolicyName</b>     | Specifies a route policy.                                                                                                                                                                                                                                        |

### Configure an IPv6 IS-IS Accept Policy for a specific I-SID

Configure an IPv6 IS-IS accept policy for a specific I-SID to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.



#### Note

If the route policy changes, you must re-apply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.

- You must configure a route-map.
- Change the VRF instance as required to configure an IPv6 IS-IS accept policy for a particular I-SID on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Accept Isid** tab.
4. Select **Insert**.
5. In the **Isid** field, specify the I-SID value.
6. Select **Enable** to apply the filter.
7. (Optional) In the **PolicyName** field, specify the route-map name.
8. Select **Insert**.

### Accept Isid Field Descriptions

Use the data in the following table to configure the **Accept Isid** tab.

| Name              | Description                                                                                                                                                                                    |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Isid</b>       | Specifies a particular I-SID number that represents local or remote Layer 3 VSNs to which the IPv6 IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT). |
| <b>Enable</b>     | Enables or disables the I-SID entry. The default is disabled.                                                                                                                                  |
| <b>PolicyName</b> | Specifies the route policy name.                                                                                                                                                               |

*Configure an IPv6 IS-IS accept policy for a specific advertising BEB and I-SID*

Configure a specific advertising Backbone Edge Bridge (BEB) with a specific I-SID to allow you to apply the IPv6 IS-IS accept policy to routes for a specific advertising BEB.



#### Note

If the route policy changes, you must re-apply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map.
- Change the VRF instance as required to configure an IPv6 IS-IS accept policy for a particular advertising BED and I-SID on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

## Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Accept Nick-Name Isid** tab.
4. Select **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Isid** field, specify an I-SID number.
7. Select **Enable** to apply the filter.
8. (Optional) In the **PolicyName** field, specify the route-map name.
9. Select **Insert**.

## Accept Nick-Name Isid Field Descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid** tab.

| Name                  | Description                                                                                            |
|-----------------------|--------------------------------------------------------------------------------------------------------|
| <b>AdvertisingRtr</b> | Specifies the SPBM nickname to apply the IS-IS accept policy to routes for a specific advertising BEB. |
| <b>Isid</b>           | Specifies the I-SID value. The value 0 is used for the Global Router.                                  |
| <b>Enable</b>         | Enables or disables the I-SID entry. The default is disabled.                                          |
| <b>PolicyName</b>     | Specifies the route policy name. You must configure a policy earlier in a separate procedure.          |

## Configure an I-SID List for an IPv6 IS-IS Accept Policy

Configures a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IPv6 IS-IS accept policy applies. After you create the list of I-SID numbers, you must then create, configure, and enable the IPv6 IS-IS accept policy.



### Note

When creating an I-SID list, you can add I-SID entries until the maximum limit for supported Layer 3 I-SIDs is reached. The system truncates any additional I-SID entries. The maximum limit includes the I-SIDs for locally configured Layer 3 VSNs and the I-SIDs specified for IS-IS accept policy filters.

Refresh the EDM tab to view the actual list of I-SIDs in the I-SID list.

## Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Change the VRF instance as required to configure an I-SID list for an IPv6 IS-IS accept policy on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Isid-List** tab.
4. Select **Insert**.
5. In the **Name** field, specify a name for the I-SID list.
6. Select **Isid** or **Isid-List**.
7. Specify an I-SID number or a list of I-SID numbers.
8. Select **Insert**.
9. Select **Apply**.

### Isid-List Field Descriptions

Use the data in the following table to configure the **Isid-List** tab.

| Name        | Description                                                                                                                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b> | Specifies the name of the I-SID list.                                                                                                                                                                                       |
| <b>Isid</b> | Specifies a particular I-SID number or a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IPv6 IS-IS accept policy applies.<br>An I-SID value of 0 represents the global routing table (GRT). |

#### *Configure an IPv6 IS-IS Accept Policy for a specific I-SID List*

Configure an IPv6 IS-IS accept policy for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.



#### Note

If the route policy changes, you must re-apply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map.
- Change the VRF instance as required to configure an IPv6 IS-IS accept policy for a particular I-SID list on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Accept Isid-List** tab.
4. Select **Insert**.

5. In the **Name** field, specify the I-SID list name.
6. Select **Enable** to apply the filter.
7. (Optional) In the **PolicyName** field, specify the route-map name.
8. Select **Insert**.

### Accept Isid-List Field Descriptions

Use the data in the following table to configure **Accept Isid-List** tab.

| Name              | Description                                                        |
|-------------------|--------------------------------------------------------------------|
| <b>Name</b>       | Specifies the name of I-SID list.                                  |
| <b>Enable</b>     | Enables or disables the I-SID list entry. The default is disabled. |
| <b>PolicyName</b> | Specifies the route policy name.                                   |

### *Configure an IPv6 IS-IS Accept Policy for a specific Advertising BEB and I-SID List*

Configure an IPv6 IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB) for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.



#### Note

If the route policy changes, you must reapply the IPv6 IS-IS accept policy, unless it was the last sequence in the configuration.

### Before You Begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map.
- Change the VRF instance as required to configure an IPv6 IS-IS accept policy for a particular advertising BEB and I-SID list on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

### About This Task

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IS-IS**.
3. Select the **Accept Nick-Name Isid-List** tab.
4. Select **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Name** field, specify an I-SID list name.
7. Select **Enable** to apply the filter.



8. (Optional) In the **PolicyName** field, specify the route-map name.
9. Select **Insert**.

#### Accept Nick-Name Isid-List Field Descriptions


Use the data in the following table to configure the **Accept Nick-Name Isid-List** tab.

| Name                  | Description                                                                                                                                                                                                                       |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AdvertisingRtr</b> | Specifies the SPBM nickname to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter is present the device applies the specific filter. |
| <b>Name</b>           | Specifies the I-SID list name.                                                                                                                                                                                                    |
| <b>Enable</b>         | Enables or disables the SPBM nickname advertising router entry. The default is disabled.                                                                                                                                          |
| <b>PolicyName</b>     | Specifies a route policy name.                                                                                                                                                                                                    |

## IP Shortcuts SPBM Configuration Example

The following figure shows a sample IP Shortcuts over SPBM deployment.

Examples and network illustrations in this document may not include your specific platform. Unless otherwise noted, the concept illustrated applies to all supported platforms. For information about high-level feature support, see *Fabric Engine and VOSS Feature Matrix*.



**Figure 117: SPBM IP Shortcuts**

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see [SPBM and IS-IS Infrastructure Fundamentals](#) on page 1178.

Note the following:

- IP IS-IS redistribution needs to be configured to inject IP shortcuts routes into IS-IS. The one exception is the circuitless IP address configured as the IS-IS ip-

source-address. This address is automatically advertised without the need for a redistribution rule.

- In the displayed configuration, only direct routes are injected (the same configuration is possible for static routes). To inject IPv6 routes, you must enable route redistribution of IPv6 direct, IPv6 static, and OSPFv3 routes into IS-IS.
- No IP address needs to be configured on SwitchG.

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example.

### SwitchC

#### CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter

```
interface loopback 1
ip address 1 10.0.0.1/255.255.255.255
exit
```

#### ISIS CONFIGURATION

```
router isis
ip-source-address 10.0.0.1
```

#### ISIS SPBM CONFIGURATION

```
spbm 1 ip enable
exit
```

#### VLAN CONFIGURATION

```
vlan create 13 type port-mstprstp 0
vlan members 13 1/2 portmember
interface Vlan 13
ip address 10.0.13.1 255.255.255.0
exit
```

#### IP REDISTRIBUTION CONFIGURATION - GlobalRouter

```
router isis
redistribute direct
redistribute direct metric 1
redistribute direct enable
exit
```

#### IP REDISTRIBUTE APPLY CONFIGURATIONS

```
isis apply redistribute direct
```

### SwitchD

#### CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter

```
interface loopback 1
ip address 1 10.0.0.2/255.255.255.255
exit
```

#### ISIS CONFIGURATION

```
router isis
ip-source-address 10.0.0.2
```

```

ISIS SPBM CONFIGURATION

spbm 1 ip enable
exit

VLAN CONFIGURATION

vlan create 14 type port-mstprstp 0
vlan member add 14 1/2
interface Vlan 14
ip address 10.0.14.1 255.255.255.0
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router isis
redistribute direct
redistribute direct metric 1
redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct

```

### Verifying Operation — SwitchC

```

SwitchC:1# show isis spbm ip-unicast-fib
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX PREFIX IP ROUTE
 INTERFACE COST COST TYPE PREFERENCE AREA AREA-NAME

GRT - 10.0.0.2/32 SwitchD 4000 1/30 20 1 Internal 7 HOME area-9.00.02
GRT - 10.0.14.1/24 SwitchD 4000 1/30 20 1 Internal 7 HOME area-9.00.02

Home : Total number of SPBM IP-UNICAST FIB entries 2
Remote: Total number of SPBM IP-UNICAST FIB entries 0
=====

SwitchC:1# show ip route
=====
IP Route - GlobalRouter
=====
DST MASK NEXT NH INTER
VRF COST FACE PROT AGE TYPE PRF ACCESS

10.0.0.1 255.255.255.255 10.0.0.1 - 1 0 LOC 0 DB 0 Trusted
10.0.0.2 255.255.255.255 SwitchD Glob~ 20 4000 ISIS 0 IBS 7 Trusted
10.0.13.1 255.255.255.0 10.0.13.1 - 1 13 LOC 0 DB 0 Trusted
10.0.14.1 255.255.255.0 SwitchD Glob~ 20 4000 ISIS 0 IBS 7 Trusted

4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

### Verifying operation — SwitchD

```

SwitchD:1# show isis spbm ip-unicast-fib
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX PREFIX IP ROUTE
 INTERFACE COST COST TYPE PREFERENCE AREA AREA-NAME

GRT - 10.0.0.1/32 SwitchC 4000 1/20 20 1 Internal 7 HOME area-9.00.02
GRT - 10.0.13.1/24 SwitchC 4000 1/20 20 1 Internal 7 HOME area-9.00.02

Home : Total number of SPBM IP-UNICAST FIB entries 2
Remote: Total number of SPBM IP-UNICAST FIB entries 0

```

```

SwitchD:1# show ip route
=====
IP Route - GlobalRouter
=====
DST MASK NEXT VRF NH INTER
COST FACE PROT AGE TYPE PRF ACCESS

10.0.0.1 255.255.255.255 SwitchC Glob~ 20 4000 ISIS 0 IBS 7 Trusted
10.0.0.2 255.255.255.255 10.0.0.2 - 1 0 LOC 0 DB 0 Trusted
10.0.13.1 255.255.255.0 SwitchC Glob~ 20 4000 ISIS 0 IBS 7 Trusted
10.0.14.1 255.255.255.0 10.0.14.1 - 1 14 LOC 0 DB 0 Trusted

4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

## Layer 3 VSN Configuration

**Table 143: Layer 3 VSN product support**

| Feature     | Product         | Release introduced                                                                                                                                                                                                                                                                                 |
|-------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Layer 3 VSN | 4220 Series     | Fabric Engine 9.2<br>The 4220 Series switches support only a single active VRF with IP configuration.                                                                                                                                                                                              |
|             | 5320 Series     | Fabric Engine 8.6<br>For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> <li>• 5320-16P-2MXT-2X</li> <li>• 5320-24T-4X-XT</li> <li>• 5320-24T-24S-4XE-XT</li> <li>• 5320-48P-8XE</li> <li>• 5320-48T-8XE</li> </ul> |
|             | 5420 Series     | VOSS 8.4                                                                                                                                                                                                                                                                                           |
|             | 5520 Series     | VOSS 8.2.5                                                                                                                                                                                                                                                                                         |
|             | 5720 Series     | Fabric Engine 8.7                                                                                                                                                                                                                                                                                  |
|             | 7520 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|             | 7720 Series     | Fabric Engine 8.10                                                                                                                                                                                                                                                                                 |
|             | 7830 Series     | Fabric Engine 9.3                                                                                                                                                                                                                                                                                  |
|             | VSP 4900 Series | VOSS 8.1                                                                                                                                                                                                                                                                                           |
|             | VSP 7400 Series | VOSS 8.0                                                                                                                                                                                                                                                                                           |


## Layer 3 VSN Configuration Fundamentals

This section provides fundamental concepts on Layer 3 VSN.

For information about supported service types, see [Fabric Connect Service Types](#) on page 1518.

### SPBM Layer 3 VSN

The SPBM Layer 3 VSN feature is a mechanism to provide IP connectivity over SPBM for VRFs. SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF.



**Figure 118: SPBM Layer 3 VSN**

In the preceding figure, the BEBs connect over the SPBM cloud running IS-IS. VRF red and green are configured on the BEBs. VRF red on BEB A has to send and receive routes from VRF red on BEB D. Similar operations are required for VRF green on BEB A and BEB D.

IS-IS TLV 184 is used to advertise SPBM Layer 3 VSN route information across the SPBM cloud. To associate advertised routes with the appropriate VRF, each VRF is associated with an I-SID. All VRFs in the network that share the same I-SID participate in the same VSN.



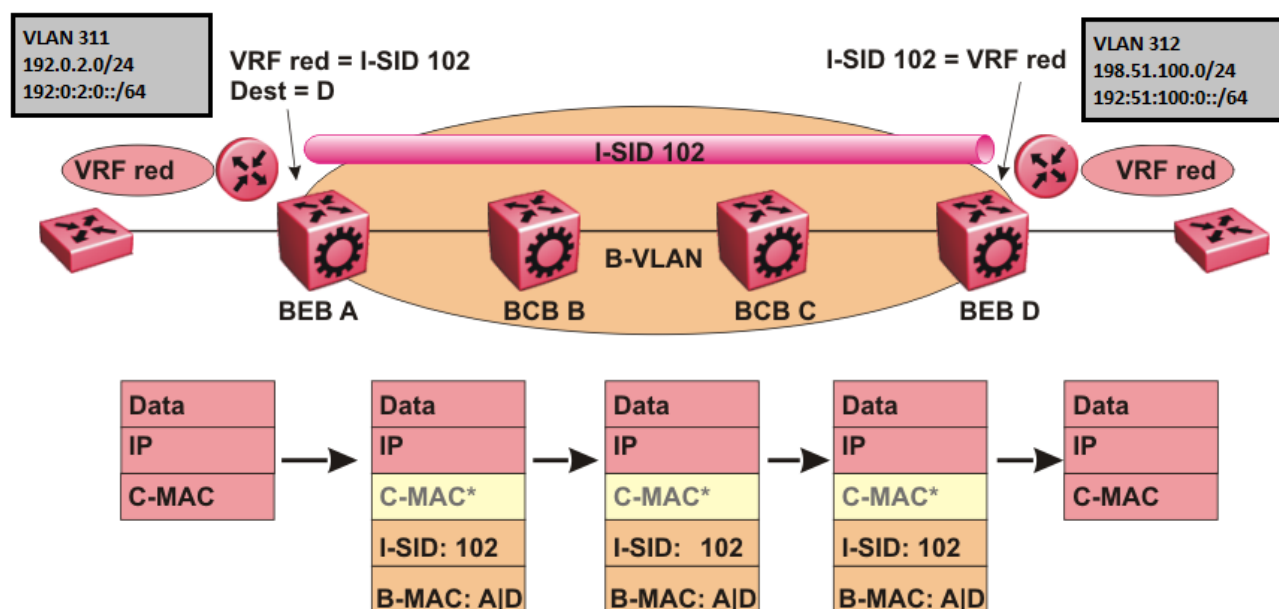
#### Note

IPv4 Layer 3 VSN and IPv6 Layer 3 VSN coexist and share the same I-SID. You need to configure the I-SID only one time. The advantage of having two separate VPNs, one for IPv4 and one for IPv6, is because it gives the user an option to enable them separately.

In the preceding example, I-SID 101 is associated with VRF green and I-SID 102 is associated with VRF red. The I-SID ties the advertised routes to a particular VRF. This identifier must be the same on all edge nodes for a particular VRF, and must be unique across all the VRFs on the same node.

When IS-IS receives an update from an edge node, it looks for the Layer 3 VSN TLV, and if one exists, it looks at the I-SID identifier. If that identifier is mapped to a local VRF, it extracts the IPv4 or IPv6 routes and adds them to the RTM of that VRF.

With SPBM Layer 3 VSN, the packet forwarding works in a similar fashion as the IP Shortcuts on the Global Router, with the difference that the encapsulation includes the I-SID to identify the VRF that the packet belongs to. The following figure shows the packet forwarding for VRF red.



**Figure 119: Packet forwarding in SPBM Layer 3 VSN**

When BEB A receives traffic from VRF red that must be forwarded to the far-end location, it performs a lookup and determines that VRF red is associated with I-SID 102 and that BEB D is the destination for I-SID 102. BEB A then encapsulates the IP data into a new BMAC header, using destination BMAC: D.



**Note**

With SPBM Layer 3 VSN, the CMAC header is all null. This header does not have any significance in the backbone. It is included to maintain the same 802.1ah format for ease of implementation.

At BEB D, the node strips off the BMAC encapsulation, and performs a lookup to determine the destination for traffic with I-SID 102. After identifying the destination as VRF red, the node forwards the packet to the destination VRF.



**Note**

IPv4 Layer 3 VSN and IPv6 Layer 3 VSN coexist and share the same I-SID. The advantage of having two separate VPNs, one for IPv4 and one for IPv6 is because it gives user an option to enable them separately.

#### *Enable/disable ICMP Response on VRFs/Layer 3 VSNs*

This feature supports VRFs/Layer 3 VSNs to operate in stealth mode by disabling ICMP responses on specific VRFs/Layer 3 VSNs.

If the ICMP response is disabled, the switch does not respond to any ICMP requests received on the VRFs/Layer 3 VSNs.

If the ICMP response is enabled, the switch responds to ICMP requests received on the VRF/Layer 3 VSNs.

### *Layer 3 VSN Restrictions*

#### **Layer 3 VSN Route Resolution**

Layer 3 VSN IPv4 BGP and static routes that have their next-hops resolved using IS-IS routes can result in traffic loss.

Choose the following workarounds, based on your deployment and needs:

- Use static routes to reach the loopbacks used as BGP peers. Static routes have better preference than IS-IS. Use static routes with next-hops reachable on the UNI side (Layer 2 VSN).
- Use OSPF to reach the loopbacks used as BGP peers, but ensure that the OSPF route towards the BGP peer is chosen as the best route because IS-IS has a better preference than OSPF. There are several ways to accomplish this: either do not redistribute that route in IS-IS if it is not needed, or control the redistribution with a route-map.
- Have BGP peers reachable directly using a C-VLAN; do not use loopback interfaces as BGP peer addresses.
- If none of the workaround scenarios are suitable for your deployment, do not use internal Border Gateway Protocol (iBGP) peering.

#### **IPv6 Layer 3 VSN Restrictions**

Consider the following when you configure the IPv6 Layer 3 VSN :

- You can enable IPv6 Layer3 VSN only when the **boot config flags spbm-config-mode** command is enabled.
- You must enable IPv4 Shortcuts and IPv6 Shortcuts.

## Layer 3 VSN configuration using the CLI

This section provides a procedure to configure Layer 3 VSNs using the command line interface (CLI).

### *Configure SPBM IPv4 Layer 3 VSN*

#### **Before You Begin**

- Configure the required SPBM IS-IS infrastructure.
- Configure a VRF on the switch. For more information, see [VRF Lite Configuration using CLI](#) on page 4081.
- Create the Customer VLANs and add slots/ports.

#### **About This Task**

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 VSN to advertise IPv4 routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF. In the VRF, just like in the Global Router (VRF 0), the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

### Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



#### Note

You cannot use the derived name for the untrusted (\*-unt) or unrestricted (\*-unr) segments of a Segmented VRF to enter this mode.

2. Create an IPv4 VPN instance on the VRF:

```
ipvpn
```

3. Configure SPBM Layer 3 VSN:

```
i-sid <0-16777215>
```

4. Enable IPv4 VPN on the VRF:

```
ipvpn enable
```

By default, a new IPv4 VPN instance is disabled.

5. Display all IPv4 VPNs:

```
show ip ipvpn [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

6. Identify routes on the local switch to be announced into the SPBM network:

```
isis redistribute {direct | bgp | ospf | rip | static}
```

7. Enable routes on the local switch to be announced into the SPBM network:

```
isis redistribute {direct | bgp | ospf | rip | static} enable
```

8. Identify other routing protocols to which to redistribute IS-IS routes:

```
ip {bgp | ospf | rip} redistribute isis
```

9. (Optional) Tag routes that BGP or OSPF redistribute into IS-IS:

```
isis redistribute {bgp | ospf} tag <1-4294967295>
```

10. Enable IS-IS redistribution to other routing protocols:

```
ip {bgp | ospf | rip} redistribute isis enable
```

11. Return to Privileged EXEC mode:

```
end
```



## 12. Apply the configured redistribution:

```
isis apply redistribute {direct | bgp | ospf | rip | static} vrf
WORD<1-16>
```

```
ip bgp apply redistribute isis vrf WORD<1-16>
```

```
ip ospf apply redistribute isis vrf WORD<1-16>
```

```
ip rip apply redistribute isis vrf WORD<1-16>
```

## 13. Display the redistribution configuration:

```
show ip isis redistribute [vrf WORD<1-16>] [vrffids WORD<0-512>]
```

## Example

Create the IPv4 VPN instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router vrf green
Switch:1(router-vrf)#ipvpn
Switch:1(router-vrf)#i-sid 109
Switch:1(router-vrf)#ipvpn enable
Switch:1(router-vrf)#show ip ipvpn
=====
IPv4 IPVPN
=====
VRF Name VRF ID IPv4 IPVPN IPv6 IPVPN Trusted Unrestricted Untrusted I-SID Name

green 1 enabled disabled 109 0 0 ExtremeServer1

1 out of 1 Total IPv4 L3 VSN, 1 active IPv4 and 0 active IPv6 displayed.
Switch:1(router-vrf)#isis redistribute ospf
Switch:1(router-vrf)#isis redistribute ospf enable
Switch:1(router-vrf)#isis redistribute ospf enable
Switch:1(config)#end
Switch:1#isis apply redistribute ospf vrf vrfred
Switch:1#show ip isis redistribute vrf vrfred
=====
ISIS Redistribute List - VRF vrfred
=====
SOURCE MET MTYPE SUBNET ENABLE LEVEL RPOLICY TAG

LOC 1 internal allow FALSE l1 0
```

## Variable Definitions

The following table defines parameters for the **show ip ipvpn** command.

| Variable                   | Value                   |
|----------------------------|-------------------------|
| <i>vrf</i> WORD<1-16>      | Specifies the VRF name. |
| <i>vrffids</i> WORD<0-512> | Specifies the VRF ID.   |

The following table defines parameters for the **i-sid** command.

| Variable     | Value                                                                                                                                           |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <0-16777215> | Assigns an I-SID to the VRF being configured.<br>Use the <b>no</b> or <b>default</b> option to remove the I-SID to VRF allocation for this VRF. |

The following table defines parameters for the **isis redistribute** command.

| Variable                                    | Value                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>{direct   bgp   ospf   rip   static}</i> | Specifies the protocol.                                                                                                                                                                                                                                                                                 |
| <i>enable</i>                               | Enables the redistribution of the specified protocol into the SPBM network.<br>The default is disabled. Use the <b>no</b> or <b>default</b> options to disable the redistribution.                                                                                                                      |
| <i>metric &lt;0-65535&gt;</i>               | Configures the metric (cost) to apply to redistributed routes. The default is 1.                                                                                                                                                                                                                        |
| <i>metric-type {external internal}</i>      | Configures the type of route to import into the protocol. The default is internal.                                                                                                                                                                                                                      |
| <i>route-map WORD&lt;0-64&gt;</i>           | Configures the route policy to apply to redistributed routes. Specifies a name.                                                                                                                                                                                                                         |
| <i>subnets {allow suppress}</i>             | Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow. |
| <i>tag &lt;1-4294967295&gt;</i>             | Apply a tag to routes that BGP or OSPF redistribute to IS-IS. You can then configure router filtering to prevent tagged IS-IS routes from being redistributed back to either protocol.                                                                                                                  |

The following table defines parameters for the **isis apply redistribute** command.

| Variable                                    | Value                                                                    |
|---------------------------------------------|--------------------------------------------------------------------------|
| <i>{direct   bgp   ospf   rip   static}</i> | Specifies the protocol.                                                  |
| <i>vrf WORD&lt;1-16&gt;</i>                 | Applies IS-IS redistribute for a particular VRF. Specifies the VRF name. |

### Configure SPBM IPv6 Layer 3 VSN using CLI

#### Before You Begin

- You must enable IPv6 Shortcuts.
- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF instance on the switch. For more information, see [VRF Lite Configuration using CLI](#) on page 4081.

#### About This Task

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 VSN to advertise IPv6 routes across the SPBM network using the following procedure.

## Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable

configure terminal

router vrf WORD<1-16>
```



### Note

You cannot use the derived name for the untrusted (\*-unt) or unrestricted (\*-unr) segments of a Segmented VRF to enter this mode.

2. Create an IPv6 VPN instance on the VRF:

```
ipv6 ipvpn
```

3. Configure SPBM Layer 3 VSN:

```
i-sid <0-16777215>
```

4. Enable IPv6 VPN on the VRF:

```
ipv6 ipvpn enable
```

5. Display all IPv6 VPNs:

```
show ipv6 ipvpn [vrf WORD<1-16> | vrfids WORD<0-512>]
```

6. Identify routes on the local switch to be announced into the SPBM network:

```
ipv6 isis redistribute {bgp | direct | ospf | static}
```

7. Enable routes on the local switch to be announced into the SPBM network:

```
ipv6 isis redistribute {direct | bgp | ospf | rip | static} enable
```

8. Identify the routing protocol to which to redistribute IS-IS routes:

```
ipv6 ospf redistribute isis
```

9. Enable IS-IS redistribution to OSPF:

```
ipv6 ospf redistribute isis enable
```

10. Return to Privileged EXEC mode:

```
end
```

11. Apply the configured redistribution to a specific VRF:

```
ipv6 isis apply redistribute {direct | bgp | ospf | rip | static} vrf
WORD<1-16>
```

12. Apply the OSPF configuration to a specific VRF:

```
ipv6 ospf apply redistribute isis vrf WORD<1-16>
```

13. Display the redistribution configuration:

```
show ipv6 isis redistribute [vrf WORD<1-16> | vrfids WORD<0-512>]
```

14. Verify IPv6 IS-IS routes:

```
show ipv6 route vrf WORD<1-16>
```

## Examples

Create the IPv6 VPN instance:

```
Switch:1>enable
Switch:1#configure terminal
```

```

Switch:1(config)#router vrf vrfred
Switch:1(router-vrf)#ipv6 ipvpn
Switch:1(router-vrf)#i-sid 100
Switch:1(router-vrf)#ipv6 ipvpn enable
Switch:1(router-vrf)#show ipv6 ipvpn
=====
 IPv6 IPVPN
=====
VRF Name VRF ID IPv6 IPVPN IPv4 IPVPN I-SID I-SID Name

vrfred 2 enabled disabled 100 ISID-100

1 out of 1 Total IPv6 L3 VSN, 1 active IPv6 and 0 active IPv4 displayed.
Switch:1(router-vrf)#ipv6 isis redistribute direct enable
Switch:1(router-vrf)#ipv6 ospf redistribute isis enable
Switch:1(router-vrf)#ipv6 ospf apply redistribute isis vrf vrfred

Switch:1(router-vrf)#show ipv6 route vrfred
=====
 IPv6 Routing Table Information - VRF vrfred
=====
Destination Address/PrefixLen NEXT HOP VID/BID/TID PROTO COST AGE TYPE PREF

55:0:0:0:0:0:0:0/64 Switch V-2 ISIS 10 0 B 7

1 out of 1 Total Num of Route Entries displayed.

TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

```

## Variable Definitions

The following table defines parameters for the **ipv6 ipvpn** command.

| Variable      | Value                                        |
|---------------|----------------------------------------------|
| <i>enable</i> | Enables IPv6 IPVPN. The default is disabled. |

The following table defines parameters for the **show ipv6 ipvpn** command.

| Variable                        | Value                   |
|---------------------------------|-------------------------|
| <i>vrf WORD&lt;1-16&gt;</i>     | Specifies the VRF name. |
| <i>vrfids WORD&lt;0-512&gt;</i> | Specifies the VRF ID.   |

The following table defines parameters for the **i-sid** command.

| Variable                  | Value                                         |
|---------------------------|-----------------------------------------------|
| <i>&lt;0-16777215&gt;</i> | Assigns an I-SID to the VRF being configured. |

The following table defines parameters for the **isis redistribute** command.

| Variable                              | Value                                                                                                   |
|---------------------------------------|---------------------------------------------------------------------------------------------------------|
| <i>{bgp   direct   ospf   static}</i> | Specifies the protocol.                                                                                 |
| <i>enable</i>                         | Enables the redistribution of the specified protocol into the SPBM network.<br>The default is disabled. |

### Display SPBM IPv6 Unicast Forwarding Information Base

#### About This Task

Perform this procedure to display SPBM IPv6 unicast Forwarding Information Base (FIB).

#### Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display SPBM IPv6 unicast FIB:

```
show isis spbm ipv6-unicast-fib [all] [id <1-16777215>] [spbm-nh-as-
mac] [home|remote]
```

#### Example

```
Switch:1>show isis spbm ipv6-unicast-fib
=====
SPBM IPv6-UNICAST FIB ENTRY INFO
=====
```

| VRF  | ISID | Dest<br>ISID | Destination       | NH BEB | VLAN | INTERFACE | SPBM<br>COST | PREFIX<br>COST | METRIC<br>TYPE | IP ROUTE<br>PREFERENCE | AREA | AREA-NAME    |
|------|------|--------------|-------------------|--------|------|-----------|--------------|----------------|----------------|------------------------|------|--------------|
| GRT  | -    | -            | 00:16:ca:23:73:df | e12    | 10   | 1/22      | 10           | 1              | Internal       | 7                      | HOME | area-9.00.02 |
| GRT  | -    | 11           | 00:16:ca:23:73:df | esp    | 20   | 1/22      | 10           | 1              | Internal       | 7                      | HOME | area-9.00.02 |
| vrfl | 11   | 100          | 00:18:b0:bb:b3:df | e12    | 10   | 1/22      | 10           | 1              | External       | 7                      | HOME | area-9.00.02 |
| vrfl | 11   | 11           | 00:14:c7:e1:33:e0 | ess    | 20   | 1/22      | 10           | 1              | External       | 7                      | HOME | area-9.00.02 |

```

Home: Total number of SPBM IPv6-UNICAST FIB entries 4
Remote: Total number of SPBM IPv6-UNICAST FIB entries 0

```

#### Variable Definitions

The following table defines parameters for the **show isis spbm ipv6-unicast-fib** command.

| Variable                     | Value                                                                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------|
| <i>all</i>                   | Displays all IS-IS SPBM IPv6 unicast Forwarding Information Base (FIB) information for all VRFs.    |
| <i>home</i>                  | Displays the IS-IS SPBM IPv6 unicast FIB information that the system configures in the home area.   |
| <i>id &lt;1-16777215&gt;</i> | Displays IS-IS SPBM IPv6 unicast FIB information by I-SID ID.                                       |
| <i>remote</i>                | Displays the IS-IS SPBM IPv6 unicast FIB information that the system configures in the remote area. |
| <i>spbm-nh-as-mac</i>        | Displays the next hop BMAC of the IPv6 unicast FIB entry.                                           |

*Display IS-IS Link State Database Information*

Perform the following procedure to display the IS-IS link state database related information on the switch.

**Procedure**

1. To enter User EXEC mode, log on to the switch.
2. Display IS-IS link state database information:

```
show isis lsdb ipv6-unicast [i-sid <0-16777215>] [lspid
xxxx.xxxx.xxxx.xx-xx] [sysid xxxx.xxxx.xxxx] [home|remote]
```

**Example**

```
Switch:1>show isis lsdb ipv6-unicast
=====
ISIS IPv6-UNICAST-ROUTE SUMMARY
=====
```

| I-SID | ADDRESS            | PREFIX<br>LENGTH | METRIC | METRIC<br>TYPE | TLV<br>TYPE | LSP<br>FRAG | HOST<br>NAME | AREA |
|-------|--------------------|------------------|--------|----------------|-------------|-------------|--------------|------|
| 4     | 2222:0:0:0:0:0:0:0 | 64               | 1      | Internal       | 184         | 0x2         | 4210         | HOME |
| 4     | 2222:0:0:0:0:0:0:0 | 64               | 1      | Internal       | 184         | 0x2         | 4210         |      |

```
REMOTE
=====
2 out of 2 Total Num of Entries
```

## Layer 3 VSN configuration using EDM

This section provides procedures to configure Layer 3 VSNs using Enterprise Device Manager (EDM).

*Configure SPBM IPv4 Layer 3 VSN***Before You Begin**

- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF and IP VPN instance on the switch. For more information about how to configure a VRF, see [VRF Lite Configuration Using EDM](#) on page 4092.
- You must create the Customer VLANs and add slots/ports.

**About This Task**

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 Virtual Services Network (VSN) to advertise IPv4 routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF. In a non-default VRF, just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute a protocol into IS-IS, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols. This procedure includes a route policy. To configure IS-IS redistribution, see [Configure IPv4 IS-IS Redistribution](#) on page 1634.

## Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP-VPN**.
3. Select the **VPN** tab.
4. To create an IP VPN instance, select **Insert**.
5. Select the ellipsis (...), select a VRF to associate with the IP VPN, and select **Ok**.
6. Select **Insert**.
7. In **Enable**, select **enable** to enable the IP VPN on the VRF.
8. In **Trusted Isid**, specify an I-SID to associate with the VPN.
9. Select **Apply**.
10. In the navigation pane, expand **Configuration > IP**.
11. Select **Policy**.
12. To identify routes on the local switch to be announced into the SPBM network, select the **Route Redistribution** tab.
13. Select **Insert**.
14. In **DstVrflid**, select the ellipsis (...), and then select the destination VRF ID and select **Ok**.
15. In **Protocol**, select **isis** as the route destination.
16. In **SrcVrflid**, select the ellipsis (...) and then select the source VRF ID and select **Ok**.
17. In **RouteSource**, select the source protocol.
18. In **Enable**, select **enable**.
19. In **RoutePolicy**, select the ellipsis (...), choose the route policy to apply to the redistributed routes and select **Ok**.
20. Configure the other parameters as required.
21. Select **Insert**.
22. To apply the redistribution configuration, select the **Applying Policy** tab.
23. Select **RedistributeApply**, and then select **Apply**.

## Configure SPBM IPv6 Layer 3 VSN

### Before You Begin

- You must enable IPv6 Shortcuts.
- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF and IPv6 VPN instance on the switch. For more information, see [VRF Lite Configuration Using EDM](#) on page 4092.

### About This Task

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 Virtual Services Network (VSN) to advertise IPv6 routes across the SPBM network from one VRF to another using the following procedure.


SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF.

### Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6-VPN**.
3. Select the **VPN** tab.
4. Select **Insert**.
5. Select the ellipsis [...], and select a VRF.
6. Select **Ok**.
7. Select **Insert**.
8. In **IsidNumber**, double-click the **0** value, and then enter the service instance identifier (I-SID) to assign to the IPv6-VPN.
9. Select **Apply**.
10. In **Enable**, select **true** or **false**.
11. Select **Apply**.
12. In the navigation pane, expand **Configuration > VRF Context View**.
13. Select **Set VRF Context View**.
14. Select the **VRF** tab.
15. Select a context to view.
16. Select **Launch VRF Context view**.
17. In the navigation pane, expand **Configuration > IPv6**.
18. Select **IS-IS**.
19. Select the **Redistribute** tab.
20. Select **Insert**.
21. Configure the parameters as required.
22. Select **Insert**.
23. Select **Apply**.

### Layer 3 VSN configuration example

The following figure shows a sample Layer 3 VSN deployment.



**Figure 120: Layer 3 VSN**

The following sections show the steps required to configure the Layer 3 VSN parameters in this example.




Note that IP IS-IS redistribution needs to be configured to inject the VRF routes into IS-IS.

You must first configure basic SPBM and IS-IS infrastructure.

### *VRF green configuration*

The following figure shows the green VRF in this Layer 3 VSN example.



**Figure 121: Layer 3 VSN — VRF green**

The following sections show the steps required to configure the green VRF parameters in this example.

### **VRF green – Switch-C**

```
VRF CONFIGURATION

ip vrf green vrfid 1

VLAN CONFIGURATION

vlan create 101 type port-mstprstp 0
vlan mlt 101 1
vlan members 101 1/2 portmember
interface Vlan 101
vrf green
ip address 10.1.101.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 13990001
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf green
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS
```

```
isis apply redistribute direct vrf green
```

### VRF green – Switch-D

```
VRF CONFIGURATION

ip vrf green vrfid 1

VLAN CONFIGURATION

vlan create 102 type port-mstprstp 0
vlan mlt 102 1
vlan members add 102 1/2 portmember
interface vlan 102
vrf green
ip address 10.1.102.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 13990001
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF


router vrf green
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf green
```

### VRF red configuration

The following figure shows the red VRF in this Layer 3 VSN example.



**Figure 122: Layer 3 VSN — VRF red**

The following sections show the steps required to configure the red VRF parameters in this example.

**VRF red – Switch-C**

```
VRF CONFIGURATION

ip vrf red vrfid 2

VLAN CONFIGURATION

vlan create 201 type port-mstprstp 0
vlan mlt 201 1
vlan members 201 1/2 portmember
interface Vlan 201
vrf red
ip address 10.2.201.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf red
ipvpn
i-sid 13990002
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf red
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf red
```

**VRF red – Switch-D**

```
VRF CONFIGURATION

ip vrf red vrfid 2

VLAN CONFIGURATION

vlan create 202 type port-mstprstp 0
vlan mlt 101 1
vlan members 202 1/2 portmember
interface Vlan 202
vrf red
ip address 10.3.202.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf red
ipvpn
i-sid 13990002
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf red
isis redistribute direct
```

```

isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf red

```

### Verifying Layer 3 VSN operation

The following sections show the steps required to verify the Layer 3 VSN configuration in this example.

#### Switch-C

```

Switch-C:1# show isis spbm ip-unicast-fib
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

GRT - - 10.0.0.2/32 Switch-D 4000 1/3 20 1 7
GRT - - 10.0.14.0/24 Switch-D 4000 1/3 20 1 7

Total number of SPBM IP-UNICAST FIB entries 2
=====

Switch-C:1# show isis spbm ip-unicast-fib id 13990001
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

green - 13990001 10.1.101.0/24 Switch-D 4000 1/2 20 1 7

Total number of SPBM IP-UNICAST FIB entries 1
=====

Switch-C:1# show isis spbm ip-unicast-fib id 13990002
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

red - 13990002 10.2.202.0/24 Switch-D 4000 1/3 20 1 7

Total number of SPBM IP-UNICAST FIB entries 1
=====

Switch-C:1# show isis spbm ip-unicast-fib id all
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

GRT - - 10.0.0.2/32 Switch-D 4000 1/3 20 1 7
GRT - - 10.0.14.0/24 Switch-D 4000 1/3 20 1 7
green - 13990001 10.1.102.0/24 Switch-D 4000 1/3 20 1 7
red - 13990002 10.2.202.0/24 Switch-D 4000 1/3 20 1 7

```

```

Total number of SPBM IP-UNICAST FIB entries 4

```

## Switch-D

```
Switch-D:1# show isis spbm ip-unicast-fib
```

```
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

GRT - - 10.0.0.1/32 Switch-C 4000 1/2 20 1 7
GRT - - 10.0.13.0/24 Switch-C 4000 1/2 20 1 7

Total number of SPBM IP-UNICAST FIB entries 2

```

```
Switch-D:1# show isis spbm ip-unicast-fib id 13990001
```

```
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

green - 13990001 10.1.101.0/24 Switch-C 4000 1/2 20 1 7

Total number of SPBM IP-UNICAST FIB entries 1

```

```
Switch-D:1# show isis spbm ip-unicast-fib id 13990002
```

```
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

red - 13990002 10.2.201.0/24 Switch-C 4000 1/2 20 1 7

Total number of SPBM IP-UNICAST FIB entries 1

```

```
Switch-D:1# show isis spbm ip-unicast-fib id all
```

```
=====
 SPBM IP-UNICAST FIB ENTRY INFO
=====
 VRF DEST
VRF ISID ISID Destination NH BEB VLAN OUTGOING SPBM PREFIX IP ROUTE
 ISID ISID Destination NH BEB VLAN INTERFACE COST COST PREFERENCE

GRT - - 10.0.0.1/32 Switch-C 4000 1/2 20 1 7
GRT - - 10.0.13.0/24 Switch-C 4000 1/2 20 1 7
green - 13990001 10.1.101.0/24 Switch-C 4000 1/2 20 1 7
red - 13990002 10.2.201.0/24 Switch-C 4000 1/2 20 1 7

Total number of SPBM IP-UNICAST FIB entries 4

```

## VRF green—Switch-C

```
Switch-C:1# show ip route vrf green
```

```
=====
 IP Route - VRF green
=====
 NH INTER
DST MASK NEXT VRF/ISID COST FACE PROT AGE TYPE PRF ACCESS

10.1.101.0 255.255.255.0 10.1.101.1 - 1 101 LOC 0 DB 0 Trusted
```

```

10.1.102.0 255.255.255.0 Switch-D vrf green 20 4000 ISIS 0 IBSV 7 Trusted

2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.

TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

### VRF green—Switch-D

```

Switch-D:1# show ip route vrf green
=====
IP Route - VRF green
=====
DST MASK NEXT NH INTER
 VRF/ISID COST FACE PROT AGE TYPE PRF ACCESS

10.1.101.0 255.255.255.0 Switch-C vrf green 20 4000 ISIS 0 IBSV 7 Trusted
10.1.102.0 255.255.255.0 10.1.102.1 - 1 102 LOC 0 DB 0 Trusted

2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.

TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

### VRF red—Switch-C

```

Switch-C:1# show ip route vrf red
=====
IP Route - VRF red
=====
DST MASK NEXT NH INTER
 VRF/ISID COST FACE PROT AGE TYPE PRF ACCESS

10.2.201.0 255.255.255.0 10.2.201.1 - 1 201 LOC 0 DB 0 Trusted
10.2.202.0 255.255.255.0 Switch-D vrf red 20 4000 ISIS 0 IBSV 7 Trusted

2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.

TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

### VRF red—Switch-D

```

Switch-D:1# show ip route vrf red
=====
IP Route - VRF red
=====
DST MASK NEXT NH INTER
 VRF/ISID COST FACE PROT AGE TYPE PRF ACCESS

10.2.201.0 255.255.255.0 Switch-C vrf red 20 4000 ISIS 0 IBSV 7 Trusted
10.2.202.0 255.255.255.0 10.2.202.1 - 1 202 LOC 0 DB 0 Trusted
2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.

TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```



# Internet Key Exchange

[IKEv2](#) on page 1676

[IKE Limitations](#) on page 1678

[IKE Configuration using CLI](#) on page 1678

[IKE Configuration using EDM](#) on page 1689

**Table 144: Internet Key Exchange product support**

| Feature                        | Product         | Release introduced |
|--------------------------------|-----------------|--------------------|
| Internet Key Exchange (IKE) v2 | 4220 Series     | Fabric Engine 9.2  |
|                                | 5320 Series     | Fabric Engine 8.6  |
|                                | 5420 Series     | VOSS 8.4           |
|                                | 5520 Series     | VOSS 8.2.5         |
|                                | 5720 Series     | Fabric Engine 8.7  |
|                                | 7520 Series     | Fabric Engine 8.10 |
|                                | 7720 Series     | Fabric Engine 8.10 |
|                                | 7830 Series     | Fabric Engine 9.3  |
|                                | VSP 4900 Series | VOSS 8.1           |
|                                | VSP 7400 Series | VOSS 8.0           |

Internet Key Exchange (IKE) protocol creates a Security Association (SA) in IPsec. The SA is the relationship between two network devices that define attributes such as authentication mechanism, encryption and hash algorithms, exchange mode, and key length for secured communications. The SA should be agreed by both devices.

The IKE protocol is based on Internet Security Association and Key Management Protocol (ISAKMP) which helps in building a secured connection between two or more hosts using the following concepts:

- authentication
- encryption
- key management
- security association (SA)
- policy

IKE uses a key exchange mechanism based on the Diffie-Hellman encryption key exchange protocol. IKE provides periodic automatic key renegotiation, pre-shared and public key infrastructures, and anti-replay defense. It is layered on top of the UDP protocol and uses UDP port 500 to exchange information between peers.

## IKE Phases

---

A switch negotiates with a peer using IKE in two phases.

- In phase 1, the switch negotiates the IKE SA to protect the negotiations that take place in phase 2. The SAs negotiated in phase 1 are bi-directional, and are applicable to traffic originating in both directions.
- In phase 2, the peers negotiate and establish the SAs for IPsec and session keys through quick mode. A Diffie-Hellman key exchange is done to achieve perfect forward secrecy, which ensures that the compromise of a single key does not permit access to data other than that protected by that compromised key. The SAs in phase 2 are uni-directional. They are used according to the direction of the traffic. The quick mode is initiated by either of the peer endpoints irrespective of who initiated phase 1.

## IKE Modes

---

There are two modes of exchanging messages in Phase 1:

- Main mode

This is a secure mode of exchanging messages. It allows protection of the confidentiality of the peers during negotiation. This mode provides more flexibility in proposals compared to aggressive mode. As the main mode requires a total of 6 messages to be exchanged between peers, it is more time consuming.

- Aggressive mode

This mode is less secure than the main mode. It does not protect the confidentiality of the peers. However, it requires only a total of 3 messages to be exchanged for phase 1, which makes this mode faster than the main mode. The number of total message exchange is reduced in this mode because some messages are embedded in other messages.

The mode of message exchange in phase 2 is called quick mode. In this mode a total of 3 messages are exchanged between the peers. This mode is used to establish IPsec SA. The negotiations in the quick mode are protected during the phase 1 negotiations in main mode.

## IKE Policies

---

A combination of security parameters used during the IKE SA negotiation is called a policy. The policies must be configured on both the peers and at least one of the policies should match on both ends to have a successful negotiation for. If a policy is not configured on both peers or if a policy does not match on both ends, an SA cannot be setup and data cannot be exchanged.



The following are the attributes of an IKE policy:

- Encryption — This is the cryptographic algorithm that is sent in the proposal by the initiator or responder during the phase 1 negotiation. This cryptographic algorithm is used to encrypt phase 2 negotiation messages. The supported encryption algorithms are:
  - DES
  - 3DES
  - AES
- Hash function — This function is used as part of the authentication mechanism during the authentication of peers in phase 1. It is always used with the authentication algorithm. The supported values are:
  - MD5
  - SHA1
  - SHA256
- Authentication — This process authenticates the peers. Following are the supported authentication modes:
  - Digital Signatures — The digital signatures use digital certificate which is signed by the certificate authority (CA) for authentication.
  - Pre-shared keys (PSK) — The PSKs are shared out-of-band between the peers before hand. Using PSK in main mode exchange limits identifying the peer to an IP address (and not host name).
- Diffie-Hellman (DH) Group — This is an algorithm used by two peers that are unknown to each other to establish a shared secret key. This key that is decided during phase 1 is used to encrypt subsequent message exchanges during phase 2 to establish security associations (SA) and security policies (SP) for IPsec sessions. The supported DH Groups are as follows:
  - Group 1 (MODP768)
  - Group 2 (MODP1024)
  - Group 14 (MODP2048)
- Lifetime — This is a time and data limit agreed by peers to protect an SA from getting compromised. It ensures that the peers renegotiate the SAs just before the lifetime value expires, that is, when the time limit is reached.
- Dead-peer detection – This is a process in which the switch waits for a response from peer for a limited number of seconds before declaring the peer as dead. It is a keep-alive mechanism required to perform IKE peer fail-over and to reclaim lost resources by freeing up SAs that are no longer in use.

## IKE Authentication

---

The security gateway of a peer must authenticate the security gateway of the peer it intends to communicate with. This ensures that IKE SAs are established between the peers. The switch supports the following two authentication methods:

- Digital certificates (using RSA algorithms)

For digital certificate authentication, the initiator signs the message interchange data using the private key. The responder uses the public key of the initiator to

verify the signature. The public key is exchanged by messages containing an X.509v3 certificate. This certificate provides an assurance that the identity of a peer, as represented in the certificate, is associated with a particular public key.

- Pre-shared keys

Pre-shared key authentication, the same secret must be configured on both security gateways before the gateways can authenticate each other.

## Signature Authentication

---


The switch receives the digital signature of its peer in a message exchange. The switch verifies the digital signature by using the public key of the peer. The certificate of the peer, received during the IKE negotiation, contains the public key. To ensure that the peer certificate is valid, the switch verifies its digital signature by using the certificate authority (CA) public key contained in the root CA certificate. The switch and its IKE peer require at least one common trusted root CA for authentication to work.

When IKE is configured to use digital certificates for authentication, the certificates are retrieved from the trusted certificate store in the switch, based on the provided distinguished name. The certificates received from the peer are verified with the public key.

## IKEv2

---

The software supports IKEv2, which is an enhancement of the IKEv1 protocol. All IKEv2 communications consist of pairs of messages: a request and a response. The IKEv2 protocol uses a non-reliable transport protocol (UDP using ports 500). The pairs of exchanges allows ensuring of reliability to the IKEv2 protocol, as there is an expected response for each request.



IKEv2 provides a number of improvements over IKEv1, including the following:

- A simplified initial exchange of messages that reduces latency and increases connection establishment speed.
  - IKEv2 makes use of a single four-message exchange instead of the eight different initial exchanges of IKEv1.
  - It improves upon IKEv1's latency by making the initial exchange to be of two round trips of four messages, and allows the ability to add setup of a child SA on that exchange.
- Improved reliability through the use of sequence numbers, and acknowledgments.
  - IKEv2 reduces the number of possible error states by making the protocol reliable as all messages are acknowledged and sequenced.
- IKE SA integrity algorithms are supported only in IKEv2.
- Traffic Selectors are specified in IKEv2 by their own payloads type and not by overloading ID payloads. This makes the Traffic Selectors more flexible.
- No lifetime negotiations for IKEv2, but in IKEv1 SA lifetimes are negotiated.

## IKEv2 OCSP Validation

Confirmation of certificate reliability is essential to achieve the security assurances public key cryptography provides. One fundamental element of such confirmation is reference to certificate revocation status. IKEv2 enables the use of Online Certificate Status Protocol (OCSP) for in-band signaling of certificate revocation status. The IKEv2 supports the authentication methods as pre shared key and digital certificate. It allows the verification of the digital certificate sent by the peer whether it is revoked or not.

This is done through a method by sending the digital certificate to the OCSP server. The OCSP server in turn verifies the certificate status and sends the response back. Based on the response from OCSP server, the device validates the certificate.

## IKE Limitations

---

This section describes the restrictions associated with this feature.

- XAUTH ( 2-factor authentication ) is not supported.
- Domain of Interpretation is not supported other than for IPsec.
- Custom IKE messages and vendor ID for the messages are not supported.
- IKE fragmentation is not supported.
- IKE and IPsec are not supported on the Segmented Management Instance interfaces, or with management applications such as RADIUS and TACACS+. You can configure RADIUS security with RADsec on supported devices.

## IKE Configuration using CLI

---

The topics in this section provide the IKE CLI configuration.

### Configure an IKE Phase 1 Profile

#### About This Task

Use the following procedure to configure an IKE Phase 1 profile.

#### Procedure

1. Enter Global Configuration mode:  
`enable`  
`configure terminal`
2. Create an IKE phase 1 profile:  
`ike profile WORD<1-32>`
3. Configure the IKE phase 1 profile hash algorithm:  
`ike profile WORD<1-32> hash-algo <md5|sha|sha256|any>`
4. Configure the IKE phase 1 profile encryption algorithm:  
`ike profile WORD<1-32> encrypt-algo <desCbc|3DesCbc|aesCbc|any>`
5. Configure the IKE phase 1 profile Diffie-Hellman group:  
`ike profile WORD<1-32> dh-group <modp768|modp1024|modp2048|any>`
6. Configure the IKE phase 1 encryption key length:  
`ike profile WORD<1-32> encrypt-key-len <128|192|256>`
7. Configure the IKE phase 1 lifetime, in seconds:  
`ike profile WORD<1-32> lifetime-sec <0-4294967295>`
8. (Optional) Delete the IKE Phase 1 profile:  
`no ike profile WORD<1-32>`

### Variable Definitions

The following table defines parameters for the **ike profile** commands.

| Variable                                                    | Value                                                                                                                                                                                                                                        |
|-------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>profile WORD&lt;1-32&gt;</code>                       | Specifies the IKE profile name.                                                                                                                                                                                                              |
| <code>hash-algo &lt;md5 sha sha256 any&gt;</code>           | Specifies the type of hash algorithm. The default value is sha256.                                                                                                                                                                           |
| <code>encrypt-algo&lt;desCbc 3DesCbc aesCbc any&gt;</code>  | Specifies the type of encryption algorithm. The default value is aesCbc.                                                                                                                                                                     |
| <code>dh-group &lt;modp768 modp1024 modp2048 any&gt;</code> | Specifies the Diffie-Hellman (DH) group. DH groups categorize the key used in the key exchange process, by its strength. The key from a higher group number is more secure. The default value is modp2048.                                   |
| <code>encrypt-key-len &lt;128 192 256&gt;</code>            | Specifies the length of the encryption key. The default is 256.                                                                                                                                                                              |
| <code>lifetime-sec &lt;0-4294967295&gt;</code>              | Specifies the lifetime value in seconds. The lifetime ensures that the peers renegotiate the SAs just before the expiry of the lifetime value, to ensure that Security Associations are not compromised. The default value is 86400 seconds. |

## Create an IKE Phase 1 Policy

IKE policy establishes Security Associations (SA) and message exchanges with IKE peers to successfully set up secured channels.

### About This Task

Use the following procedure to create the IKE Phase 1 policy.

### Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```
2. Create an IKE Phase 1 profile:  

```
ike policy WORD<1-32> laddr WORD<1-256> raddr WORD<1-256>
```
3. (Optional) Delete the IKE Phase 1 profile:  

```
no ike policy WORD<1-32>
```

### Variable Definition

The following table defines parameters for the **ike policy <1-320> laddr** command.

| Variable                             | Value                                         |
|--------------------------------------|-----------------------------------------------|
| <code>policy WORD&lt;1-32&gt;</code> | Specifies the name of the IKE Phase 1 policy. |
| <code>laddr WORD&lt;1-256&gt;</code> | Specifies the local IPv4 or IPv6 address.     |
| <code>raddr WORD&lt;1-256&gt;</code> | Specifies the remote IPv4 or IPv6 address.    |

## Configure Profile for IKE Phase 1 Policy

Use the following procedure to configure the IKE Phase1 profile to be used for the IKE Phase 1 policy.

### Procedure

1. Enter Global Configuration mode:  
`enable`  
  
`configure terminal`
2. Configure the profile name to be used for IKE Phase 1 policy:  
`ike policy WORD<1-32> profile WORD<1-32>`

### Variable Definition

The following table defines parameters for the **ike policy WORD<1-32> profile WORD<1-32>** command.

| Variable                              | Value                                                                    |
|---------------------------------------|--------------------------------------------------------------------------|
| <code>policy WORD&lt;1-32&gt;</code>  | Specifies the name of the IKE Phase 1 policy.                            |
| <code>profile WORD&lt;1-32&gt;</code> | Specifies the name of the IKE Phase 1 profile to be used for the policy. |

## Configure IKE Phase 2 Perfect Forward Secrecy

Use the following procedure to configure IKE Phase 2 perfect forward secrecy (PFS).

### About This Task

A Diffie-Hellman key exchange is done to achieve perfect forward secrecy. This ensures that the compromise of even a single key does not permit access to data other than that protected by that key.

### Procedure

1. Enter Global Configuration mode:  
`enable`  
  
`configure terminal`
2. Configure the IKE Phase 2 perfect forward secrecy:  
`ike policy WORD<1-32> p2-pfs <enable|disable> [use-ike-group <enable|disable>] [dh-group <modp768|modp1024|modp2048|any]`

### 3. (Optional) Disable Phase 2 perfect forward secrecy:

```
no ike policy <1-32> p2-pfs
```

#### Variable Definition

The following table defines parameters for the **ike policy WORD<1-32> p2-pfs** command.

| Variable                                              | Value                                                                                                                                                                                                                                                              |
|-------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>policy WORD&lt;1-32&gt;</i>                        | Specifies the name of the IKE Phase 1 policy.                                                                                                                                                                                                                      |
| <i>p2-pfs</i>                                         | Enables the Phase 2 perfect forward secrecy.                                                                                                                                                                                                                       |
| <i>dh-group &lt;modp768 modp1024 modp2048 any&gt;</i> | Configures the Diffie-Hellman (DH) group to be used for Phase 2 perfect forward secrecy (PFS). The default value is modp2048.<br><br><b>Note:</b><br>For Federal Information Processing Standards (FIPS) compliance, only the default value modp2048 is supported. |
| <i>use-ike-group &lt;enable disable&gt;</i>           | Specifies whether to use the IKE Phase 1 DH group for Phase 2 PFS or not to use it. The default is enable.                                                                                                                                                         |

## Configure the IKE Authentication Method

Use the following procedure to configure the IKE authentication method. The default is pre-shared key.

### About This Task

As part of the IKE protocol, one security gateway must authenticate another security gateway to make sure that IKE SAs are established with the intended party. The switch supports two authentication methods:

- Digital certificates

Configure peer identity name for IKE phase 1 and revocation check method.

- Pre-shared keys

Configure the same secret on both security gateways before the gateways can authenticate each other.

### Procedure

#### 1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

## 2. Configure the IKE authentication method using one of the following:

- To use a digital certificate:

```
ike policy WORD<1-32> auth-method digital-certificate [peer-name
WORD <1-64> | revocation-check-method <crl|none|ocsp>]
```

- To use a pre-shared key:

```
ike policy WORD<1-32> auth-method pre-shared-key
```

```
ike policy WORD<1-32> pre-shared-key WORD<0-32>
```

### Variable Definitions

The following table defines parameters for the **ike policy WORD<1-32> auth-method** command.

| Variable                                                                | Value                                                  |
|-------------------------------------------------------------------------|--------------------------------------------------------|
| <i>pre-shared-key</i>                                                   | Specifies the authentication method as pre-shared key. |
| <i>digital-certificate peer-name WORD &lt;1-64&gt;</i>                  | Specifies peer identity name for IKE phase 1.          |
| <i>digital-certificate revocation-check-method&lt;crl none ocsp&gt;</i> | Specifies the revocation check method.                 |

The following table defines parameters for the **ike policy WORD<1-32> pre-shared-key** command.

| Variable                               | Value                                                                                                                              |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <i>pre-shared-key WORD&lt;0-32&gt;</i> | Specifies the pre-shared key. For Federal Information Processing Standards (FIPS) compliance, the minimum length is 14 characters. |

## Configure Dead-Peer Detection Timeout

Use the following procedure to configure the dead-peer detection (DPD) timeout for the IKE Phase 1 policy.

### About This Task

Dead Peer Detection (DPD) timeout is the interval for which the system sends messages to a peer to confirm its availability.



### Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```
2. Configure the DPD timeout:  

```
ike policy WORD<1-32> dpd-timeout <1-4294967295>
```

### Variable Definition

The following table defines parameters for the **ike policy WORD<1-32> dpd-timeout** command.

| Variable                                | Value                                                                                                        |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <i>policy WORD&lt;1-32&gt;</i>          | Specifies the name of the IKE Phase 1 policy.                                                                |
| <i>dpd-timeout &lt;1-4294967295&gt;</i> | Specifies the dead peer detection timeout in seconds for the IKE Phase 1 policy. The default is 300 seconds. |

## Enable the Admin State of IKE Phase 1 Policy

Use the following procedure to enable admin state of IKE Phase 1 policy.

### Procedure

1. Enter Global Configuration mode:  

```
enable
```

```
configure terminal
```
2. Enable admin state of IKE Phase 1 policy:  

```
ike policy WORD<1-32> enable
```
3. (Optional) Disable IKE Phase 1 policy:  

```
no ike policy WORD<1-32> enable
```

## Display IKE Profiles

Use the following procedure to display the configured IKE profiles:

### Procedure

1. Enter Privileged EXEC mode:  

```
enable
```
2. Display all IKE profiles:  

```
show ike profile
```
3. Display a specific ike profile:  

```
show ike profile WORD<1-32>
```

### Example

```
Switch:1#show ike profile
```

```
=====
```

| IKE Profile      |              |                 |                    |             |                  |                     |
|------------------|--------------|-----------------|--------------------|-------------|------------------|---------------------|
| Name             | Hash<br>Algo | Encrypt<br>Algo | Encrypt<br>Key Len | DH<br>Group | Exchange<br>Mode | Lifetime<br>seconds |
| DFLT_IKE_PROFILE | sha256       | aesCbc          | 256                | modp2048    | main             | 86400               |
| ikePRO           | sha256       | aesCbc          | 256                | modp2048    | main             | 180                 |
| test             | sha256       | aesCbc          | 256                | modp2048    | main             | 86400               |

### Variable Definition

The following table defines parameters for the **show ike profile** command.

| Variable                        | Value                                              |
|---------------------------------|----------------------------------------------------|
| <i>profile WORD&lt;1-32&gt;</i> | Specifies the name of the profile to be displayed. |

## Display IKE Policies

Use the following procedure to display the configured IKE policies.

### Procedure

1. Enter Privileged EXEC mode:  
`enable`
2. Display all IKE policies:  
`show ike policy`
3. Display a specific IKE policy:  
`show ike policy WORD<1-32>`
4. Display a specific IKE policy at local address.  
`show ike policy WORD<1-32> laddr WORD<1-256>`
5. Display a specific IKE policy at remote address.  
`show ike policy WORD<1-32> laddr WORD<1-256> raddr WORD<1-256>`

### Example

Switch:1#show ike policy

| IKE Policy     |              |                 |                 |                 |
|----------------|--------------|-----------------|-----------------|-----------------|
| Policy<br>Name | Addr<br>Type | Local Address   | Remote Address  | Profile<br>Name |
| iketest3       | IPv4         | 192.168.152.104 | 192.168.149.207 | test            |
| vlpol          | IPv4         | 192.168.152.104 | 192.168.152.152 | ikepro          |

| IKE Policy                         |                    |              |                |                                     |
|------------------------------------|--------------------|--------------|----------------|-------------------------------------|
| Policy<br>identity<br>Name<br>name | Profile<br>Version | Auth-Method  | Pre-Shared Key | Revocation-Check<br>Method<br>peer- |
| iketest3                           | 2                  | digital-cert |                | ocsp                                |
| vlpol                              | 1                  | digital-cert |                | ocsp                                |

| IKE Policy  |             |             |            |         |                |          |        |
|-------------|-------------|-------------|------------|---------|----------------|----------|--------|
| Policy Name | DPD Timeout | Admin State | Oper State | P2 PFS  | Use IKE DH Grp | DH Group | IntfId |
| iketest3    | 300         | enable      | up         | disable | enable         | modp1024 | 3047   |
| vlpol       | 300         | enable      | up         | disable | enable         | modp1024 | 3047   |

### Variable Definition

The following table defines parameters for the **show ike policy** command.

| Variable                       | Value                                             |
|--------------------------------|---------------------------------------------------|
| <i>policy WORD&lt;1-32&gt;</i> | Specifies the name of the policy to be displayed. |
| <i>laddr WORD&lt;1-256&gt;</i> | Specifies the local IPv4 or IPv6 address.         |
| <i>raddr WORD&lt;1-256&gt;</i> | Specifies the remote IPv4 or IPv6 address.        |

## Display IKE Security Association

Use the following procedure to display the configured IKE Phase 1 for version 1 and 2 security associations (SA).

### Procedure

1. Enter Privileged EXEC mode:  
enable
2. Display all the security associations:  
show ike sa
3. Display security associations for IKE Phase 1 for version 1:  
show ike sa version v1 WORD<1-32> laddr WORD<1-256> raddr WORD<1-256>
4. Display security associations for IKE Phase 1 for version 2:  
show ike sa version v2 WORD<1-32> laddr WORD<1-256> raddr WORD<1-256>

### Example

```
Switch:1(config)#show ike sa version v1
```

```
=====
IKE V1 Phase 1 Security Association
=====
```

| Policy Name | Addr Type | Local Address | Remote Address | Initiator/Responder |
|-------------|-----------|---------------|----------------|---------------------|
| ikepsk      | IPv4      | 192.0.2.5     | 198.51.100.15  | Initiator           |

```
=====
IKE V1 Phase 1 Security Association
=====
```

| Name   | DPD Timeout | Hash Algo | Encrypt Algo | DH Group | Lifetime seconds | Status |
|--------|-------------|-----------|--------------|----------|------------------|--------|
| ikepsk | 300         | sha       | aesCbc       | modp2048 | 3600             | active |

```
Switch:1(config)#show ike sa version v2
```

| IKE V2 Phase 1 Security Association |             |               |              |                |          |                     |        |
|-------------------------------------|-------------|---------------|--------------|----------------|----------|---------------------|--------|
| Policy Name                         | Addr Type   | Local Address |              | Remote Address |          | Initiator/Responder |        |
| v2policy                            | IPv4        | 203.0.113.6   |              | 198.51.100.20  |          | Responder           |        |
| IKE V2 Phase 1 Security Association |             |               |              |                |          |                     |        |
| Name                                | DPD Timeout | Hash Algo     | Encrypt Algo | Integrity Algo | DH Group | Lifetime seconds    | Status |
| v2policy                            | 300         | sha256        | aesCbc       |                | modp2048 | 86400               | active |

### Variable Definition

The following table defines parameters for the **show ike sa** command.

| Variable                                                                                                         | Value                                                                   |
|------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| <i>sa</i>                                                                                                        | Specifies the IKE security association identifier.                      |
| <i>version v1</i><br><i>WORD&lt;1-32&gt; laddr</i><br><i>WORD&lt;1-256&gt; raddr</i><br><i>WORD&lt;1-256&gt;</i> | Specifies the local IPv4 or IPv6 address for IKE Phase 1, version 1 SA. |
| <i>version v2</i><br><i>WORD&lt;1-32&gt; laddr</i><br><i>WORD&lt;1-256&gt; raddr</i><br><i>WORD&lt;1-256&gt;</i> | Specifies the local IPv4 or IPv6 address for IKE Phase 1, version 2 SA. |

## Configure an IKEv2 Profile

### About This Task

Use the following procedure to configure an IKEv2 profile.

#### Procedure

1. Enter Global Configuration mode:  
enable  
  
configure terminal
2. Create an IKEv2-profile:  
ike v2-profile WORD<1-32>
3. Configure the IKEv2 profile hash algorithm:  
ike v2-profile WORD<1-32> hash-algo <md5|sha|sha256|any>
4. Configure the IKEv2 profile encryption algorithm:  
ike v2-profile WORD<1-32> encrypt-algo <desCbc|3DesCbc|aesCbc|any>

## 5. Configure the IKEv2 profile integrity algorithm

```
ike v2-profile WORD<1-32> integrity-algo <hmac-md5|hmac-sha|hmac-sha256|aes-xcbc|any>
```

## 6. Configure the IKEv2 profile dh group

```
ike v2-profile WORD<1-32> dh-group <modp768|modp1024|modp2048|any>
```

## 7. Configure the IKEv2 profile encryption key length:

```
ike v2-profile WORD<1-32> encrypt-key-len <128|192|256>
```

## 8. Configure the IKEv2 profile lifetime, in seconds:

```
ike v2-profile WORD<1-32> lifetime-sec <0-4294967295>
```

## 9. (Optional) Delete the IKEv2 profile:

```
no ike v2-profile WORD<1-32>
```

## Variable Definition

The following table defines parameters for the **ike v2-profile** commands.

| Variable                                           | Value                                                                                                                                                                                                                                        |
|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>profile</i> WORD<1-32>                          | Specifies the IKE v2-profile name.                                                                                                                                                                                                           |
| <i>hash-algo</i> <md5 sha sha256 any>              | Specifies the type of hash algorithm. The default value is sha256.                                                                                                                                                                           |
| <i>encrypt-algo</i> <desCbc 3DesCbc aesCbc any>    | Specifies the type of encryption algorithm. The default value is aesCbc.                                                                                                                                                                     |
| <i>integrity-algo</i> <md5 sha-1 sha-256 aes-xcbc> | Specifies the type of integrity algorithm. The default is sha256.                                                                                                                                                                            |
| <i>dh-group</i> <modp768 modp1024 modp2048 any>    | Specifies the Diffie-Hellman (DH) group. DH groups categorize the key used in the key exchange process, by its strength. The key from a higher group number is more secure. The default value is modp2048.                                   |
| <i>encrypt-key-len</i> <128 192 256>               | Specifies the length of the encryption key. The default is 256.                                                                                                                                                                              |
| <i>lifetime-sec</i> <0-4294967295>                 | Specifies the lifetime value in seconds. The lifetime ensures that the peers renegotiate the SAs just before the expiry of the lifetime value, to ensure that Security Associations are not compromised. The default value is 86400 seconds. |

## Display IKEv2 Profiles

Use the following procedure to display the configured IKEv2 profiles.

## Procedure

## 1. Enter Privileged EXEC mode:

```
enable
```

## 2. Display all IKEv2 profiles:

```
show ike v2-profile
```

3. Display a specific IKEv2 profile:
- ```
show ike v2-profile WORD<1-32>
```

Example

```
Switch:1#show ike v2-profile test
=====
                                IKE2 Profile
=====
Encrypt      Exchange      Hash      Encrypt
Name          Mode          Algo          Algo          Key
Length
-----
test          main          sha256          aesCbc
256
=====
                                IKE2 Profile
=====
Lifetime      DH      Integrity
Name          Group          Algorithm
seconds
-----
test          modp2048          sha256
180
```

Variable Definitions

The following table defines parameters for the **show ike v-2profile** command.

Variable	Value
WORD<1-32>	Specifies the name of the policy.

Configure x509 Certificate Identity

About This Task

Use the following procedure to bind a certificate identity to the IKE certificate store.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Configure the certificate subject name:

```
ike certificate-identity cert-subject-name WORD<1-45>
```

Variable Definitions

The following table defines parameters for the **ike certificate-identity** command.

Variable	Value
<i>cert-subject-name WORD<1-45></i>	Specifies the digital certificate subject name to be used as the identity certificate. If a subject name is not specified, the default subject name is Global.

IKE Configuration using EDM

The topics in this section provide the IKE EDM configuration.

Configure Digital Certificate Subject Name

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IKE**.
3. Select **Globals** tab.
4. In **CertIdentitySubjectName**, type the digital certificate subject name use as identity certificate.
5. Select **Apply**.

Globals Field Definitions

Use the data in the following table to use the **Globals** tab.

Name	Description
CertIdentitySubjectName	Specifies the digital certificate subject name use as identity certificate in IKE.

Configure IKE Phase 1 Profile

Use the following procedure to create and configure an IKE Phase 1 profile.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IKE**.
3. Select the **Profile** tab.
4. Select **Insert**.
5. In **Name**, type a profile name.
6. Complete the remaining optional configuration to customize the policy.
7. Select **Insert**.

Profile *Field Descriptions*

Use the data in the following table to use the **Profile** tab.

Name	Description
Name	Specifies the name of the profile.
HashAlgorithm	Specifies the hash algorithms that can be used during IKE Phase 1 SA negotiation. The default value is sha256.
EncryptionAlgorithm	Specifies the encryption algorithms that can be used during IKE Phase 1 SA negotiation. The default value is aesCbc.
EncryptKeyLen	Specifies the key length that should be used during IKE Phase 1 SA negotiation. The default value is 256.
DHGroup	Specifies the Diffie-Hellman groups that can be used during IKE Phase 1 SA negotiation. The default value is modp2048.
ExchangeMode	Specifies the IKE Phase 1 negotiation mode. The default value is main.
LifetimeSeconds	Specifies the amount of time for which an IKE Phase 1 SA can remain valid during IKE Phase 1 negotiation. A value of 0 means no the SA always remains valid. The default value is 86400 seconds.

Configure IKEv2 Profile

Use the following procedure to create and configure an IKEv2 profile.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IKE**.
3. Select the **V2 Profile** tab.
4. Select **Insert**.
5. In **Name**, type a profile name.
6. Complete the remaining optional configuration to customize the policy.
7. Select **Insert**.

V2 Profile *Field Descriptions*

Use the data in the following table to use the **V2 Profile** tab.

Name	Description
Name	Specifies the IKE v2 profile name.
HashAlgorithm	Specifies the type of hash algorithm that can be used during IKE version 2 SA version 2 negotiation. The default value is sha256.
EncryptionAlgorithm	Specifies the encryption algorithms that can be used during IKE version 2 SA version 2 negotiation. The default value is aesCbc.
EncryptKeyLen	Specifies the type of encryption algorithm. The default value is keylen=256.
DHGroup	Specifies the Diffie-Hellman (DH) group. DH groups categorize the key used in the key exchange process, by its strength. The key from a higher group number is more secure. The default value is modp2048.
ExchangeMode	Specifies the IKE v2 profile negotiation mode. The default value is main.
LifetimeSeconds	Specifies the lifetime value in seconds. The lifetime ensures that the peers renegotiate the SAs just before the expiry of the lifetime value, to ensure that Security Associations are not compromised. The default value is 86400 seconds.
IntegrityAlgorithm	Specifies the type of integrity algorithm. The default value is sha256.

Configure IKE Phase 1 Policy

Use the following procedure to create and configure an IKE Phase 1 policy.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IKE**.
3. Select the **Policy** tab.
4. Select **Insert**.
5. In **LocalIfIndex**, select either **Port** or **Vlan**, and then select an interface.
6. In **LocalAddrType**, select the type of the local address.
7. In **LocalAddr**, type the address of the local peer.
8. In **RemoteAddrType**, select the type of the remote address.

9. In **RemoteAddr**, type the address of the remote peer.
10. In **Name**, type the name for the policy.
Name must be assigned when creating the policy. When the policy is created, the name cannot be changed.
11. Complete the remaining optional configuration to customize the policy.
12. Select **Insert**.

Policy Field Descriptions

Use the data in the following table to use the **Policy** tab.

Name	Description
LocalIfIndex	Specifies the Interface Index of the local address. Only port and vlan interfaces are supported.
LocalAddrType	Specifies whether the local address is an IPv4 or IPv6 address.
LocalAddr	Specifies the address of the local peer.
RemoteAddrType	Specifies whether the remote address is an IPv4 or IPv6 address.
RemoteAddr	Specifies the address of the remote peer.
Name	Specifies the name given to the policy. The name should be assigned while creating the policy. You cannot change the name after the policy is created.
ProfileName	Specifies the name of the profile that should be used for this policy.
ProfileVersion	Specifies the profile version used for the policy.
PeerName	Specifies the peer name.
AuthenticationMethod	Specifies the proposed authentication method for the Phase 1 security association. The default authentication method is pre-shared key.
PSKValue	Specifies the value of the Pre-Shared Key if the authentication method is set to PSK.
DPDTimeout	Specifies the Dead Peer Detection timeout in seconds. Default value is 300 seconds.
P2PFS	Specifies whether or not the perfect forward secrecy (PFS) is used when refreshing keys. To use PFS, select enable. The default value is disable.

Name	Description
P2PfsUselkeGroup	Specifies whether or not to use the same GroupId (Diffie-Hellman Group) for phase 2 as was used in phase 1. Ignore this entry if P2PFS is disabled. The default value is enable.
P2PfsDHGroup	Specifies the Diffie-Hellman group to use for phase 2 when P2PFS is enabled and P2PfsUselkeGroup is disabled. The default value is mod1024.
AdminState	Specifies whether the policy is administratively enabled or disabled. The default value is disable.
OperStatus	Shows is the policy is operationally up or down.
RevocationCheckMethod	Specifies the revocation check method as OCSP, CRL or none.

Display IKE Phase 1 Security Association

Use the following procedure to view the IKE Phase 1 security association.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IKE**.
3. Select the **SA** tab.

SA Field Descriptions

Use the data in the following table to use the **SA** tab.

Name	Description
Id	Specifies the profile ID.
LocalIfIndex	Specifies the Interface Index of the local address. Only port and vlan interfaces are supported.
LocalAddrType	Specifies whether the local address is an IPv4 or IPv6 address.
LocalAddr	Specifies the address of the local peer.
RemoteAddrType	Specifies whether the remote address is an IPv4 or IPv6 address.
RemoteAddr	Specifies the address of the remote peer.
Name	Specifies the name given to the SA.

Name	Description
AuthenticationMethod	Specifies the proposed authentication method for the Phase 1 security association. The default authentication method is pre-shared key.
DPDTimeout	Specifies the Dead Peer Detection timeout in seconds.
HashAlgorithm	Specifies the hash algorithm negotiated for this IKE Phase 1 SA.
EncryptionAlgorithm	Specifies the encryption algorithm negotiated for this IKE Phase 1 SA.
EncryptKeyLen	Specifies the encryption key length negotiated for this IKE Phase 1 SA.
DHGroup	Specifies the Diffie-Hellman group negotiated for this IKE Phase 1 SA.
ExchangeMode	Specifies the IKE Phase 1 SA mode.
LifetimeSeconds	Specifies the amount of time for which an IKE Phase 1 SA can remain valid during IKE Phase 1 negotiation. A value of 0 means no the SA always remains valid.
Status	Specifies whether the SA is active or inactive.
Initiator	Specifies whether specifies the whether the SA is created by an initiator or a responder.

Display IKE V2 Security Association

Use the following procedure to view the IKE version 2 security association.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IKE**.
3. Select the **V2 SA** tab.

V2 SA Field Descriptions

Use the data in the following table to use the **V2 SA** tab.

Name	Description
Id	Specifies the profile ID.
LocalIfIndex	Specifies the Interface Index of the local address. Only port and vlan interfaces are supported.

Name	Description
LocalAddrType	Specifies whether the local address is an IPv4 or IPv6 address.
LocalAddr	Specifies the address of the local peer.
RemoteAddrType	Specifies whether the remote address is an IPv4 or IPv6 address.
RemoteAddr	Specifies the address of the remote peer.
Name	Specifies the name given to the SA.
AuthenticationMethod	Specifies the proposed authentication method for theVersion 2 security association. The default authentication method is pre-shared key.
DPDTimeout	Specifies the Dead Peer Detection timeout in seconds.
HashAlgorithm	Specifies the hash algorithm negotiated for this IKE Version 2 SA.
EncryptionAlgorithm	Specifies the encryption algorithm negotiated for this IKE Version 2 SA.
EncryptKeyLen	Specifies the encryption key length negotiated for this IKE Version 2 SA.
DHGroup	Specifies the Diffie-Hellman group negotiated for this IKE Version 2 SA.
ExchangeMode	Specifies the IKE Version 2 SA mode.
LifetimeSeconds	Specifies the amount of time for which an IKE Version 2 SA can remain valid during IKE Version 2 negotiation. A value of 0 means no the SA always remains valid.
Status	Specifies whether the SA is active or inactive.
Initiator	Specifies whether specifies the whether the SA is created by an initiator or a responder.
IntegrityAlgorithm	Specifies the type of integrity algorithm.



IP Multicast

- [IP Multicast Fundamentals](#) on page 1696
- [IP Multicast Basic Configuration Using CLI](#) on page 1772
- [IP Multicast Basic Configuration Using EDM](#) on page 1796
- [Multicast Listener Discovery](#) on page 1814
- [PIM Configuration Using the CLI](#) on page 1838
- [PIM Configuration Using EDM](#) on page 1849
- [IGMP Configuration Using the CLI](#) on page 1867
- [IGMP Configuration Using EDM](#) on page 1895
- [Route Management Using the CLI](#) on page 1916
- [Route Management Using EDM](#) on page 1924
- [Multicast route statistics configuration using the CLI](#) on page 1933
- [Multicast route statistics configuration using EDM](#) on page 1942

This section describes how to administer and configure IP Multicast Routing protocols.

The topics in this section provide conceptual background, as well as CLI and EDM configuration procedures.

IP Multicast Fundamentals

IP multicast extends the benefits of Layer 2 multicasting on LANs to WANs. Use multicasting techniques on LANs to help clients and servers find each other. With IP multicast, a source can send information to multiple destinations in a WAN with a single transmission. IP multicast results in efficiency at the source and saves a significant amount of bandwidth.

Enabling Multicast on the Switch

Shortest Path Bridging (SPB) and Protocol Independent Multicast (PIM) cannot interoperate with each other on the switch at the same time. To ensure that SPB and PIM stay mutually exclusive, use the boot flag called **spbm-config-mode**:

- The **spbm-config-mode** boot flag is enabled by default. This configuration enables you to configure SPB and IS-IS, but you cannot configure PIM either globally or on an interface.

- If you disable the boot flag, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.

**Important**

- Any change to the **spbm-config-mode** boot flag requires a reboot for the change to take effect.
- If you plan to disable the boot flag, remove all SPB configurations first.
- If you plan to use the default (enabled) setting, remove all PIM configurations first.

Simplified Virtual-IST

Simplified Virtual-IST (vIST) is for conventional network deployments that use SMLT and not SPB. The Simplified vIST feature provides a single CLI command to enable the virtual IST for SMLT deployments.

- Simplified vIST is available **ONLY** for conventional multicast deployments with PIM and IGMP when the boot flag (**spbm-config-mode**) is disabled.
- When the boot flag is enabled (default setting), Simplified vIST is not available. This means that you continue to configure SPB/IS-IS for vIST.
- Simplified VIST requires that the two vIST devices be directly connected.

**Note**

- PIM is supported with Simplified vIST only, not SPB vIST. However, you do not have to configure Simplified vIST to run PIM or IGMP Snooping in a **non-SMLT** topology.
- Do not configure LACP on SPB network-to-network interface (NNI) MLT links or on the Simplified Virtual IST.
- Do not configure ECMP in PIM Simplified vIST scenarios. Running PIM in a Simplified vIST environment with ECMP enabled can lead to incorrect behavior since there are multiple options in terms of choosing the upstream node towards a host or source. For example, since the path chosen cannot be predicted (it is determined by the downstream PIM neighbor), we can end up not adding the Virtual IST MLT port in the PIM mroute's outgoing port list on the joined interface if the PIM Join Prune Message was received on an alternative path, different from the interface the local router considers to be the correct upstream to the source.

Traffic loss can occur in such an environment. Do not enable ECMP in PIM vIST scenarios.

After you disable the **spbm-config-mode** boot flag, you can configure PIM or IGMP Snooping on any VLAN including the vIST VLAN.

To configure the boot flag and Simplified vIST, see [Configure IP Multicast in SMLT Topologies](#) on page 1772 or [Configure Multicast on the Switch](#) on page 1796.

vIST VLAN IP Addresses

Do not configure an RP or BSR on the vIST VLAN because you cannot ping them outside of the vIST VLAN subnet. When you enter the **ip pim enable** command on the vIST VLAN, the following message displays:

```
WARNING: Please do not use virtual IST VLAN IP address for BSR and
RP related configurations, as unicast packets to virtual IST vlan
IP address from outside of virtual IST vlan subnet will be dropped.
Use Loopback or CLIP interface IP address for BSR and RP related
configurations.
```

Overview of IP Multicast

IP multicast transmits messages to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting, except that multicasting transmits to specific groups and broadcasting transmits to all receivers on a network. Because IP multicast transmits only one stream of data to the network where it replicates to many receivers, multicasting saves a considerable amount of bandwidth.

IP multicast services benefit applications such as video conferencing, dissemination of datagram information, and dissemination of mail or news to a large number of recipients.

Multicast protocols use different techniques to discover delivery paths.

A distribution tree is a set of multicast routers and subnetworks that permit the members of a group to receive traffic from a source. The source of the tree depends on the algorithm used by the multicast protocol. The following diagram is an example of a simple distribution tree where S is the multicast source and the arrows indicate the multicast broadcast procedure.




Figure 123: Multicast distribution tree and broadcasting

Broadcast and prune methods use multicast traffic to build the distribution tree. Periodically, the source sends or broadcasts data to the extremities of the internetwork to search for active group members. If no local members of the group exist, the router sends a message to the host, removing itself from the distribution tree, and thus pruning the router.

The following diagram illustrates how the host prunes routers from the distribution tree. First, the router sends a message to the source, after which the pruned routers do not receive multicast data.




Figure 124: Pruning routers from a distribution tree

Reverse path multicast is based on the concept that a multicast distribution tree is built on the shortest path from the source to each subnetwork that contains active receivers. After a datagram arrives on an interface, the router determines the reverse path to the source of the datagram by examining the routing table of known network sources. If the datagram is not on the optimal delivery tree, the router discards it.

Multicast host groups and their group members enable the IP multicast router to transmit just to those groups interested in receiving the traffic. The switch uses the Internet Group Membership Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. A router communicates with the hosts on a local network by sending IGMP queries. Hosts respond by issuing IGMP reports. For more information about host groups, see [Multicast Host Groups](#) and [Multicast Addresses](#) on page 1701. For more information about IGMP, see [Internet Group Management Protocol](#).

Multicast traffic forwarding transmits frames to all interfaces or subnets for which it receives IGMP reports for the multicast group indicated in the destination IP address. Multicast packets forwarded within the same virtual LAN (VLAN) remain unchanged. The switch does not forward packets to networks that do not use members of the multicast group indicated in the destination IP address.

Multicast Host Groups

IP multicast is a method for addressing, routing, and delivering a datagram to a collection of receivers called a host group.

Host groups are permanent or transient, with the following characteristics:

- A permanent host group uses a well-known, administratively assigned IP multicast group address. This address is permanent and defines the group. A permanent host group can consist of zero or more members.

- A transient host group exists only as long as members need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

A host system on an IP network sends a message to a multicast group by using the IP multicast address for the group. To receive a message addressed to a multicast group, however, the host must be a member of the group and must reside on a network where that group is registered with a local multicast router.

An IP multicast host group can consist of zero or more members and places no restrictions on its membership. Host members can reside anywhere, they can join and leave the group at any time, and they can be members of more than one group at the same time.

In general, hosts that are members of the same group reside on different networks. However, a range of multicast addresses (224.0.0.x) is reserved for locally-scoped groups. All message traffic for these hosts typically remains on the local network. Hosts that belong to a group in this address range and that reside in different networks do not receive message traffic for each other.

**Important**

You can apply a special set of filters (global filters) to multicast packets. You can also create, deny, or accept filters to configure the sources that can receive and send data. For more information about how to configure filters, see [Traffic Filtering Fundamentals](#) on page 3681.

Multicast Addresses

Each host group uses a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP version 4 multicast address is a Class D address (the high-order bits are 1110) from 224.0.0.0 to 239.255.255.255. These addresses are assigned statically for use by permanent groups and dynamically for use by transient groups.

The block of addresses from 224.0.0.1 to 224.0.0.255 is reserved for routing protocols and other low-level protocols. Multicast routers do not forward datagrams with addresses in this range because the time-to-live (TTL) value for the packet is usually 1.

Multicast Protocols

You can use the following protocols to enable multicast routing on a switch:

- Internet Group Management Protocol (IGMP)—learns the existence of host group members on directly attached subnets.
- Multicast Router Discovery (MRDISC) protocol—discovers multicast routers in a Layer 2 bridged domain configured for IGMP snoop.
- Protocol Independent Multicast (PIM)
 - Sparse Mode (PIM-SM) protocol—suitable for implementation on networks sparsely populated by receivers.

- Source Specific Multicast (PIM-SSM) protocol—uses a one-to-many model where members can receive traffic from one or more specific sources. This protocol is suitable for television channels and other content-distribution applications.

Static Source Groups

Use static source groups to configure static source-group entries in the PIM-SM, or PIM-SSM multicast routing table. PIM cannot prune these entries from the distribution tree. In other words, even if no receivers for the group exist, the multicast stream for a static source-group entry stays active. PIM never prunes static forwarding entries. If you no longer need the entries, you must manually delete them.

To configure static source groups, you must first globally enable PIM. If you disable PIM, the switch saves all of the configured static source-group entries and deactivates them. After you re-enable PIM, the switch reactivates the static source groups.

Static source groups ensure that the multicast route (mroute) records remain in the distribution tree. After receivers join the group, they do not experience a delay in receiving multicast data because they do not need to graft onto the group, or start a join process in the case of PIM. This timing is essential for applications where the multicast data must send to a receiver as soon as the receiver joins the group, for example, when a switch delivers television channels to receivers. After the receiver turns the channel, which is equivalent to joining a group, the receiver can view the channel immediately.

Static entries result in continuous traffic if the source is active, even if no receivers exist. However, the system does not forward traffic with a static entry if no receivers exist, but forwards it continuously to the switch where the entry is programmed and crosses intermediate switches on the path.

You can configure static source-group entries for a specific source or subnet. If several sources on the same subnet send traffic to the same group, traffic for all these sources flows continuously when using the subnet configuration.

After you configure static source groups, keep the following points in mind:

- If you disable PIM, the switch deactivates all of the static source groups. After you re-enable PIM, the switch activates the static source groups.
- In PIM-SM configuration, the static source-group feature works for both specific source addresses and subnet addresses by using the SrcSubnetMask field.

When the network mask is 255.255.255.255, the full source address is used to match the (S,G) which is the specific source case. When the network mask field is a subnet mask for the source, only the source subnet is used to match (S,G)s.

- In PIM-SSM configurations, static source groups have the following limitations:
 - Subnets: SSM static source groups work only with specific IP addresses. Static source groups cannot work with source subnets, so the mask must use a full 32-bit mask, 255.255.255.255, and the source must use a host address.

IP Multicast over Fabric Connect

IP Multicast over Fabric Connect introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. These extensions, combined with the use of IGMP snooping and querier functions at the edge of the SPBM cloud, efficiently transport IP multicast data by using sub-trees of the VSN shortest path tree per IP multicast group.

Internet Group Management Protocol

Table 145: Internet Group Management Protocol product support

Feature	Product	Release introduced
Internet Group Management Protocol (IGMP), including virtualization	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

A host uses IGMP to register group memberships with the local querier router to receive datagrams sent to this router targeted to a group with a specific IP multicast address.

A router uses IGMP to learn the existence of group members on networks to which it directly attaches. The router periodically sends a general query message to each of its local networks. A host that is a member of a multicasting group identifies itself by sending a response.

IGMP Queries

When multiple IGMP routers operate on a network, one router is elected to send queries. This elected querier periodically sends host membership queries (also known as general queries) to the attached local subnets. The switch supports queries from all three versions of IGMP.

IGMP Host Reports

A host that receives a membership query from a local router can respond with a host membership report, one for each multicast group that joins. A host that receives a query delays its reply by a random interval and listens for a reply from other hosts in the same host group. For example, consider a network that includes two host members—host A and host B—of the same multicast group. The router sends out a host membership query on the local network. Both host A and host B receive the query and listen on the network for a host membership report. The delay timer for host B expires first, so it responds to the query with a membership report. Hearing the response, host A does not send a report of its own for the same group.

Each query from a router to a host includes a maximum response time field. IGMP inserts a value *n* into this field specifying the maximum time in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and *n* tenths of a second for the period that it waits before sending a response. This calculation is true for IGMP versions 2 and 3. For IGMP version 1, this field is 0 but defaults to a value of 100, that is, 10 seconds.

If at least one host on the local network specifies that it is a member of a group, the router forwards to that network all datagrams that bear the multicast address for the group.

Upon initialization, the host can immediately issue a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same as requested reports.

After hosts and routers are in a steady state, they communicate in a way that minimizes the exchange of queries and reports. The designated routers establish a path between the IP multicast stream source and the end stations and periodically query the end stations about whether to continue participation. As long as a client continues to participate, all clients, including nonparticipating end stations on the switch port, receive the IP multicast stream.

Host Leave Messages

If an IGMPv2 host leaves a group and it is the host that issues the most recent report, it also issues a leave group message. The multicast router on the network issues a group-specific query to determine whether other group members exist on the network. If no host responds to the query, the router assumes that no members belonging to that group exist on that interface.

Fast Leave Feature

The switch supports a fast leave feature that is useful for multicast-based television distribution applications. Fast leave relies on an alternative leave process where the

switch stops sending traffic for the group immediately after it receives a leave message, without issuing a query to check if other group members exist on the network. Fast leave alleviates the network from additional bandwidth demand after a customer changes television channels.

The switch provides several fast leave processes for IP multicast:

- immediate leave with one user for each interface
- immediate leave with several users for each interface
- standard IGMP leave based on a Last Member Query Interval (LMQI), which you can configure in tenths of seconds

Fast leave modifies the IGMP leave processing mechanism on an IGMP interface. After the system receives an IGMP leave message on a fast leave enabled interface, the switch does not send a group-specific query and immediately stops sending traffic to the leaving member (IGMP host) port. Without fast leave, traffic continues to forward until the group times out. This situation wastes bandwidth if no receiver that requires the group traffic exists.

Fast leave mode provides two options of the fast leave mechanism—single-user mode and multiple-users mode:

- Single-user mode: In this mode, the port stops receiving traffic immediately after a group member on that port sends a leave message. Use the single-user mode if each interface port connects to only one IGMP host.
- Multiple-users mode: Use this mode if the interface port connects to multiple IGMP hosts. In this case, the port stops receiving traffic after all members leave the IGMP group. The switch removes the leaving IGMP member and, if more group members exist on that port, the switch continues sending traffic to the port.

When operating in multiple-users mode, the switch must use the correct membership information. To support multiple-users mode, multicast receivers on the same interface cannot use IGMP report suppression. If you must use IGMP report suppression, do not use this mode. Instead, use the LMQI (configurable in units of 1/10ths of seconds) to provide a faster leave process while still sending group-specific queries after the interface receives a leave message.

Fast leave mode applies to all fast-leave enabled IGMP interfaces.

IGMP Snoop

The switch provides IP multicast capability and can support all three versions of IGMP to prune group membership for each port within a VLAN. This feature is IGMP snoop.



Important

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

Use the IGMP snoop feature to optimize the multicast data flow, for a group within a VLAN, to only those ports that are members of the group. The switch builds a database

of group members by listening to IGMP reports from each port. The switch suppresses the reports heard by not forwarding them to ports other than the one receiving the report, thus forcing the members to continuously send their own reports. The switch relays group membership from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN. Furthermore, the switch forwards multicast data only to the participating group members and to the multicast routers within the VLAN.

The multicast routing functionality can coexist with IGMP snoop on the same switch, but you can configure only one of IGMP snoop or an IP multicast routing protocol, excluding IGMP, on the same VLAN.

Multicast Group Trace for IGMP Snoop

Use this feature to monitor the multicast group trace for an IGMP snoop-enabled switch. You can view the multicast group trace from CLI.

Multicast group trace tracks the data flow path of the multicast streams. Group trace tracks information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port.

IGMP Proxy

If a switch receives multiple reports for the same multicast group, it does not transmit each report to the multicast upstream router. Instead, the switch consolidates the reports into a single report and forwards the one report. If you add another multicast group or the system receives a query since it last transmitted the report upstream, the system forwards the report onto the multicast router ports. This feature is IGMP proxy.

IGMP Versions

The switch supports IGMPv1, IGMPv2, and IGMPv3. IGMPv1 and IGMPv2 are backward compatible and can exist together on a multicast network. The following list describes the purpose for each version:

- IGMPv1 provides the support for IP multicast routing. IGMPv1 specifies the mechanism to communicate IP multicast group membership requests from a host to its locally attached routers. For more information, see RFC1112.
- IGMPv2 extends the features in IGMPv1 by quickly reporting group membership termination to the routing protocol. This feature is important for multicast groups with highly volatile group membership. For more information, see RFC2236.
- IGMPv3 supports the PIM Source Specific Multicast (SSM) protocol, PIM-SM, and snooping. A host can selectively request or filter traffic from individual sources within a multicast group or from specific source addresses sent to a particular multicast group. Multicast routing protocols use this information to avoid delivering multicast packets from specific sources to networks where there are no interested receivers. For more information, see RFC3376.

For the switch implementation of PIM-SSM, each group can use multiple sources.

The following list identifies group records that a report message includes:

- current-state record

- source-list-change record
- filter-mode-change record

A current-state record is sent by a system in response to a query received on an interface. It reports the current reception state of that interface, with respect to a single multicast address.

The Record Type of a current-state record has one of the following two values:

- **MODE_IS_INCLUDE** — Indicates that the interface has a filter mode of include for the specified multicast address. The source address fields in this group record contain the source list of the interface for the specified multicast address.
- **MODE_IS_EXCLUDE** — Indicates that the interface has a filter mode of exclude for the specified multicast address. The source address fields in this group record contain the source list of the interface for the specified multicast address.

Source-List Change Record — The system sends a source-list-change record after a change of source list occurs that does not coincide with a filter-mode change on the interface for a particular multicast address. The interface on which the change occurs sends a report that includes the record. The record type of a source-list-change record can be one of the following two values:

- **ALLOW_NEW_SOURCES** — Indicates that the source address [i] fields in this group record contain a list of the additional sources that the system wishes to hear from, for packets sent to the specified multicast address. If the change was to an include source list, these are the addresses that were added to the list. If the change was to an exclude source list, these are the addresses that were deleted from the list.
- **BLOCK_OLD_SOURCES** — Indicates that the source address [i] fields in this group record contain a list of the sources that the system no longer wishes to hear from, for packets sent to the specified multicast address. If the change was to an include source list, these are the addresses that were deleted from the list; if the change was to an exclude source list, these are the addresses that were added to the list.

If a change of source list results in both allowing new sources and blocking old sources, then two group records are sent for the same multicast address, one of type **ALLOW_NEW_SOURCES** and one of type **BLOCK_OLD_SOURCES**.

Filter Mode — The switch implements the filter-mode-change record. The system sends a filter-mode-change record whenever the filter mode changes (during a change from include to exclude, or from exclude to include) for a particular multicast address. The interface on which the change occurs sends a report that includes the record. The record type of a filter-mode-change record can be one of the following two values:

- **CHANGE_TO_INCLUDE_MODE** — Indicates that the interface has changed to include filter mode for the specified multicast address. The source address [i] fields in this group record contain the new source list of the interface for the specified multicast address.
- **CHANGE_TO_EXCLUDE_MODE** — Indicates that the interface has changed to exclude filter mode for the specified multicast address. The source address [i] fields in this group record contain the new source list of the interface for the specified multicast address.

After you enable IGMPv3, the following actions occur:

- After you change the version on an interface to or from IGMPv3, the switch experiences a disruption to existing multicast traffic on that interface but traffic does recover. Do not make this change when the system passes multicast traffic.

IGMP States

Multicast routers implementing IGMPv3 keep one state for each group for every port in every attached network. This group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network. This state consists of a set of records of the following form:

- multicast address
- group timer
- filter mode (source records)

Each source record is of the form source address or source timer. If all sources within a given group are desired, an empty source record list is kept with filter-mode set to EXCLUDE. This means hosts on this network want all sources for this group to be forwarded. This is the IGMPv3 equivalent to a IGMPv1 or IGMPv2 group join.

Group Timer

A group timer represents the time for the filter-mode to expire and switch to INCLUDE mode and is used only when a group is in EXCLUDE mode.

Group timers are updated according to the types of group records received. If a group timer is expiring when a router filter-mode for the group is EXCLUDE means, there are no listeners on the attached network in EXCLUDE mode. At this point, a router will transition to INCLUDE filter-mode.

Source Timer

A source timer is maintained for every source record. Source timers are updated according to:

- the type and filter-mode of the group record received
- whenever the source is present in a received record for that group.

If a source timer expires with a router filter-mode for the group of INCLUDE, the router concludes that traffic from this particular source is no longer desired on the attached network, and deletes the associated source record.

If a source record has a running timer with a router filter-mode for the group of EXCLUDE, it means that at least one system desires the source. It should therefore be forwarded by a router on the network. If a source timer expires with a router filter-mode for the group of EXCLUDE, the router informs the routing protocol that there is no receiver on the network interested in traffic from this source. The records are deleted when the group timer expires in the EXCLUDE router filter-mode.

Processing IGMP Messages for Groups in SSM Range

IGMP messages are processed for groups in SSM range in the following scenarios:

1. IGMPv3 interface enabled; PIM-sparse or snooping enabled
 - IGMPv3 reports that contain group records with groups within SSM range are processed with no restrictions.
 - IGMPv2 reports for groups within SSM range translate to IGMPv3 reports with one group record and type IS_EXCLUDE{NULL}. These reports are processed with no restriction as an IGMPv3 report.
 - IGMPv2 leave for groups within SSM range translate to IGMPv3 reports with one group record and type TO_INCLUDE{NULL}. These reports are processed with no restriction as an IGMPv3 report.
2. IGMPv3 interface enabled; PIM-SSM or ssm-snooping enabled
 - IGMPv3 reports that contain group records with groups within SSM range received from members in the EXCLUDE mode are discarded (eg. IS_EXCLUDE and TO_EXCLUDE messages).
 - IGMPv2 reports for groups within SSM range translate to IGMPv3 reports with one group record and type ALLOW{S1,S2,...}. The source list is obtained from the global ssm-map. If there are no sources in the global ssm-map, the message is discarded. These reports are processed with no restriction as an IGMPv3 report.
 - IGMPv2 leave for groups within SSM range translate to IGMPv3 reports with one group record and type BLOCK{S1,S2,...}. The source list is obtained from the global ssm-map. If there are no sources in the global ssm-map, the message is discarded. These reports are processed with no restriction as an IGMPv3 report.



Note

In order to accept v2 messages, you must enable the compatibility mode on the IGMPv3 interface.

IGMPv3 Source-Specific Forwarding Rules

After a multicast router receives a datagram from a source destined to a particular group, the router must decide to forward the datagram to the attached network. The multicast routing protocol uses IGMPv3 information to forward datagrams to all required sources or groups on a subnetwork.

The following table describes the forwarding suggestions that IGMPv3 makes to the routing protocol. The table also identifies the action taken after the source timer expires, based on the filter mode of the group.

Group filter-mode	Source-timer value	Action
INCLUDE	TIMER > 0	Forward the traffic from the source.
INCLUDE	TIMER = 0	Stop forwarding the traffic from the source, and remove the source record. If no more source records exist for the group, delete the group record.
INCLUDE	No source elements	Do not forward the source.

Group filter-mode	Source-timer value	Action
EXCLUDE	TIMER > 0	Forward the traffic from the source.
EXCLUDE	TIMER = 0	Do not forward the traffic from the source. If no more source records exist for the group, delete the group record.
EXCLUDE	No source elements	Forward the traffic from the source.

IGMPv3 Explicit Host Tracking

IGMPv3 explicit host tracking enables the IGMP to track all the source and group members. To track all the source and group members, the sources that are in the include mode hold a list of members who want to receive traffic from that source.

The members that are in the exclude mode are on hold on the reporter list under the port data. By default, IGMPv3 explicit host tracking is disabled.



Important

If explicit host tracking is enabled, you cannot downgrade the IGMPv3 interface to IGMPv1 or IGMPv2.

IGMPv3 Fast Leave

When a BLOCK message is received for a source, you must check if the member that sent this message is the last reporter for the source. If it is the last reporter, delete the source. Else, delete the member. No group and source specific queries are sent.

When a LEAVE message is received, you must check if the member that sent this message is the last reporter for the group. If it is the last reporter, switch to INCLUDE mode if sources are available (if no sources are available the port is deleted). Else, delete the member. No group and source specific queries or group specific queries are sent.



Important

To use the IGMPv3 fast leave feature, you must first enable the explicit host tracking feature.

Synchronization of IGMPv3 over SMLT

The implementation of IGMPv3 offers support for IGMPv3 over SMLT. The Virtual-IST (vIST) peers must be in sync with the IGMPv3 reports received over SMLT links to ensure effective performance. The vIST protocol ensures the infrastructure to send such information from one vIST peer to the other.

The synchronization of IGMPv3 members and their advertised sources is different from IGMPv1 and IGMPv2. Because of IGMPv3 compatibility mode, you must consider the IGMP member version. If you have version 1 or 2 members, you must synchronize the IGMP information as IGMPv1 or IGMPv2 reports, so the peer can build an accurate database. In particular, if members with version 1 or 2 exist, the group filter mode is exclude and the exclude source list is empty. Also no v1 or v2 member will be present on any source from include list.

Each member sends IGMP reports in the same manner for all IGMP versions. The sending mechanism depends on the SMLT state.

After a vIST peer receives an IGMPv3 report over an SMLT link, it must pass the message to its peer. If the SMLT state is up, the vIST peer sends the message encapsulated in an vIST IGMPv3 message. If the SMLT state is down, the vIST peer sends the message as a plain IGMPv3 report.

In both cases the IGMPv3 message is not altered and the receiving vIST peer processes it as expected in SMLT conditions (translating the receiving port to SMLT port if applicable).

**Note**

If you enable compatibility mode and the member sends an IGMPv1 or IGMPv2 report, the message is either a vIST IGMPv1 or v2 encapsulated Message or a plain IGMPv1 or IGMPv2 report.

After SMLT up or down events occur, the vIST peer must synchronize its IGMPv3 database to its peer, taking into account the new state of the SMLT link.

If you enable IGMP explicit host tracking, each include source stores information for each member that advertises that particular source in an include list. This information is synchronized with the vIST peer.

If you do not enable explicit host tracking, each source from include list contains only information related to the last member that sent an IGMPv3 report. Only this information is synchronized with the vIST peer.

Backward Compatibility

IGMPv3 for PIM-SSM is backward compatible with IGMPv2. You can configure the switch to operate in v3-only mode or in v2-v3 compatibility mode. If you configure the switch to use v3-only mode, it ignores all v2 and v1 messages except the query message.

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv1, v2, and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message; if it is a v2 message, PIM-SM or IGMP snoop processes handle the message.

After the switch receives an IGMPv2 leave message and the group address in it is within SSM range, the switch sends the group-and-source specific query. If the group address is not within the SSM range, the switch sends the group specific query.

According to RFC3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.
- If a router that is not explicitly configured to use IGMPv1 or IGMPv2 hears an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure if the switch dynamically downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

In v2-v3 compatibility mode, an IGMPv2 host can only join if you configure a static entry in SSM map and if the interface operates in PIM-SSM mode or IGMP SSM-Snoop mode.

You can use the compatibility mode with Split MultiLink Trunking (SMLT). One core switch sends an SMLT message to the other core switch after it receives an IGMPv3 message. This action synchronizes the IGMP host information.

Implementation of IGMP

You can enable and disable multicast routing on an interface basis. If you disable multicast routing on an interface, the interface does not generate IGMP queries. If the switch or interface is in IGMP router behavior mode, for example, PIM enabled, you cannot configure IGMP snoop. The switch still learns the group membership and snoops multicast receivers on the switch VLAN or ports.

IGMP Layer 2 Querier

In a Layer 2 multicast network, you can enable Layer 2 querier on one of the switches in the VLAN. IGMP Layer 2 querier provides the IGMP querier function so that the switch can provide the recurring queries that maintain IGMP groups when you do not use multicast routing for multicast traffic.

Overview

In a multicast network, if you only need to use Layer 2 switching for the multicast traffic, you do not need multicast routing. However, you must have an IGMP querier on the network for multicast traffic to flow from sources to receivers. A multicast router provides the IGMP querier function. You can also use the IGMP Layer 2 Querier feature to provide a querier on a Layer 2 network without a multicast router.

The Layer 2 querier function originates queries for multicast receivers, and processes the responses accordingly. On the connected Layer 2 VLANs, IGMP snoop continues to provide services as normal. IGMP snoop responds to queries and identifies receivers for the multicast traffic.

You must enable Layer 2 querier and configure an IP address for the querier before it can originate IGMP query messages. If a multicast router exists on the network, the switch automatically disables the Layer 2 querier.

In a Layer 2 multicast network, enable Layer 2 querier on only one of the switches in the VLAN. A Layer 2 multicast domain supports only one Layer 2 querier. No querier election exists.

IGMP Snooping

IGMP Snooping enables Layer 2 switches in the network to examine IGMP control protocol packets exchanged between downstream hosts and upstream routers.

When Layer 2 switches examine the IGMP control protocol packets, they:

- Generate the Layer 2 MAC forwarding tables used for further switching sessions
- Regulate the multicast traffic to prevent it from flooding the Layer 2 segment of the network

IGMP Layer 2 Querier and IGMP Interaction

IGMP Layer 2 Querier uses IGMP to learn which groups have members on each of the attached physical networks, and it maintains a list of multicast group memberships for each attached network and a timer for each membership. In this case, multicast group memberships means the presence of at least one member of a multicast group on a given attached network, not a list of all of the members.

IGMP Layer 2 Querier can assume one of two roles for each of the attached networks:

- Querier
- Non-Querier

After you enable IGMP Layer 2 Querier, the system assumes it is a multicast router, so it sends the General Query, Group Specific/Group, and Source Specific Query when Leave/BLOCK messages are received. IGMP queries are required to maintain an IGMP group.



Note

Group Specific When Leave does not apply to IGMPv1.

IGMP Layer 2 Querier Limitations

The following limitations apply to IGMP Layer 2 Querier.

- IGMP Layer 2 Querier is based on IGMP Snoop. If you disable IGMP Snoop, IGMP Layer 2 Querier does not work until you enable IGMP Snoop and IGMP Layer 2 Querier.
- After you enable IGMP Snoop and IGMP Layer 2 Querier on an interface, if the system receives no IGMP query messages, it becomes the querier.

IGMP Layer 2 Querier Limitations and DvR

The following limitations apply when you configure IGMP Layer 2 Querier on DvR enabled nodes.

- You can configure IGMP Layer 2 Querier only on the DvR Controllers in a DvR domain. When you configure the following parameters on the Controllers, the configuration is automatically pushed to the DvR Leaf nodes within the domain.
 - IGMP version
 - IGMP query interval
 - IGMP query maximum response time
 - IGMP robustness value
 - IGMP last member query interval
 - IGMP compatibility mode
- You cannot configure IGMP snooping on DvR enabled Layer 2 VSNs.

For more information on DvR, see [Distributed Virtual Routing Fundamentals](#) on page 919 .

Multicast Access Control

Multicast access control is a set of features that operate with standard existing multicast protocols. You can configure multicast access control for an IP multicast-enabled port or VLAN with an access control policy that consists of several IP multicast groups.

You can use this feature to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams). For example, in a television distribution application, instead of applying a filter to each channel (multicast group), you can apply a multicast access policy to a range of channels (groups), thereby reducing the total number of filters and providing a more efficient and scalable configuration. Also, if you want to add or remove television channels from a package, you can modify the multicast access policy; you do not need to change filters for individual VLANs or ports. Multicast access policies contain an ID and a name (for example, PremiumChannels), the list of IP multicast addresses, and the subnet mask.

Multicast access control is not a regular filtering configuration. Multicast access control is for multicast streams and relies on handling multicast control and initial data to prevent hosts from sending or receiving specified multicast streams; it does not use filters. Also, multicast access control provides a list of multicast groups in one configuration using the same routing policy prefix list configuration. For information about prefix lists, see [Configuring prefix lists](#) on page 3259. You can configure multicast access control and change it dynamically to support changes in the configuration without restarting the protocol. You can change the access capabilities of a user or service subscriber without loss of service.

The following paragraph describes a typical application.

The local cable television company offers three packages; each one includes 35 channels (35 multicast groups). The company configures each package in an access control policy. This policy applies to a set of VLANs or ports to prevent users from viewing the channels on those VLANs. Use the same policy to prevent users from sending traffic to those groups (also known as spoofing) by specifying the deny-tx option for that port. After you define the packages, you can use them for access policy configuration. You can easily change the package by changing the group range, without changing all the port configurations.

The multicast access control functionality applies to an IP multicast application where you must control user access. You can use it in financial-type applications and other enterprise applications, such as multicast-based video conferencing.

Six types of multicast access control policies exist:

- deny-tx
- deny-rx
- deny-both

- allow-only-tx
- allow-only-rx
- allow-only-both

The tx policies control the sender and ingress interface for a group; the rx policies control the receivers and egress interface for a group.

deny-tx

Use the deny-tx access policy to prevent a matching source from sending multicast traffic to the matching group on the interface where you configure the deny-tx access policy. Configure this policy on the ingress interface to the multicast source. The deny-tx access policy performs the opposite function of the allow-only-tx access policy. Therefore, the deny-tx access policy and the allow-only-tx access policy cannot exist on the same interface at the same time.

For example, in [Figure 125](#), the ingress VLAN, uses a deny-tx access policy. This policy prevents multicast traffic sent by Sender from forwarding from VLAN 1 to a receiver, consequently preventing Receiver 1 and Receiver 2 from receiving data from the multicast group. You can create receive-only VLANs, such as VLAN 1, with the deny-tx policy.




Figure 125: Data flow using deny-tx policy

deny-rx

Use the deny-rx access policy to prevent a matching group from receiving IGMP reports from the matching receiver on the interface where you configure the deny-rx access policy. The deny-rx access policy performs the opposite function of the allow-only-rx access policy. Therefore, the deny-rx access policy and the allow-only-rx access policy cannot exist on the same interface at the same time.

For example, in [Figure 126](#), a VLAN 2 uses a deny-rx access policy, preventing IGMP reports sent by Receiver 1 from receiving on VLAN 2. You can deny a multicast group access to a specific VLAN or receiver using the deny-rx policy.




Figure 126: Data flow using deny-rx policy

deny-both

Use the deny-both access policy to prevent a matching IP address from both sending multicast traffic to, and receiving IGMP reports from, a matching receiver on an interface where you configure the deny-both policy. You can use this policy to eliminate all multicast activity for a receiver or source in a specific multicast group. The deny-both access policy performs the opposite function of the allow-only-both access policy. Therefore, the deny-both access policy and the allow-only-both access policy cannot exist on the same interface at the same time.

For example, in [Figure 127](#), a VLAN 2 uses a deny-both access policy, preventing VLAN 2 from receiving IGMP reports sent by Receiver 2, and preventing multicast traffic sent by Sender 2 from forwarding from VLAN 2. You can prevent certain VLANs from participating in an activity involving the specified multicast groups with the deny-both policy.




Figure 127: Data flow using deny-both policy

allow-only-tx

Use the allow-only-tx policy to allow only the matching source to send multicast traffic to the matching group on the interface where you configure the allow-only-tx policy. The interface discards all other multicast data it receives. The allow-only-tx access policy performs the opposite function of the deny-tx access policy. Therefore, the allow-only-tx access policy and the deny-tx access policy cannot exist on the same interface at the same time.

allow-only-rx

Use the allow-only-rx policy to allow only the matching group to receive IGMP reports from the matching receiver on the interface where you configure the allow-only-rx access policy. The interface discards all other multicast data it receives. The allow-only-rx access policy performs the opposite function of the deny-rx access policy. Therefore,

the allow-only-rx access policy and the deny-rx access policy cannot exist on the same interface at the same time.

allow-only-both

Use the allow-only-both policy to allow only the matching IP address to both send multicast traffic to, and receive IGMP reports from, the matching receiver on the interface where you configure the allow-only-both access policy. The interface discards all other multicast data and IGMP reports. The allow-only-both access policy performs the opposite function of the deny-both access policy. Therefore, the allow-only-both access policy and the deny-both access policy cannot exist on the same interface at the same time.

Host addresses and masks

When you configure multicast access policies, you must specify the host (IP) address and host (subnet) mask of the host to filter (the host that sends multicast traffic).

You can use the host subnet mask to restrict access to a portion of the host network. For example, if you configure the host subnet mask as 255.255.255.255, you use the full host address. To restrict access to a portion of the network of a host, use a subnet mask such as 255.255.255.0. Access control applies to the specified subnet only.

Multicast Stream Limitation Feature

You can configure the multicast stream limitation feature to limit the number of multicast groups that can join a VLAN. By limiting the number of concurrent multicast streams, a service provider can, for example, protect the bandwidth on a specific interface and control access to multicast streams.

Use multicast stream limitation in an environment where you want to limit users to a certain number of multicast streams simultaneously. For example, a television service provider can limit the number of television channels a user can watch at a time. (To a television service provider, a multicast stream is synonymous with a television channel.) If a user purchases a service contract for two single-tuner television receivers, they can use two channels flowing at the same time, but not a third. The service provider can control the bandwidth usage in addition to preventing users from watching more than the allowed number of channels at a point in time.

You can enable the multicast stream limitation feature on the switch by using one of the following methods:

- for each interface—This limitation controls the total number of streams for all clients on this router port.
- for each VLAN—This limitation controls the total number of streams for all clients on this VLAN. This method is equivalent to the interface stream limitation.
- for each VLAN port—This limitation controls the number of streams for all clients on this VLAN port. This method is equivalent to the interface port stream limitation.

You can configure the maximum number of streams for each limit independently. After the number of streams meets the limit, the interface drops additional join reports for

new streams. The maximum number of streams for each limit is 65535 and the default is 4.

Multicast Router Discovery Protocol

The Multicast Router Discovery (MRDISC) protocol can automatically discover multicast-capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and IGMP host membership reports. This feature is useful in a Layer 2 bridging domain that you configure for IGMP snoop.

IGMP multicast router discovery consists of three message types that discover multicast routers on the network:

- Multicast router advertisements: routers advertise that IP multicast forwarding is enabled on an interface.
- Multicast router solicitations: routers solicit a response of multicast router advertisements from all multicast routers on a subnet.
- Multicast router termination messages: a router terminates its multicast routing functions.

Multicast routers send multicast router advertisements periodically on all interfaces where you enable multicast forwarding. Multicast routers also send advertisements in response to multicast router solicitations.

Multicast router solicitations transmit to the IGMP-MRDISC all-routers multicast group that uses a multicast address of 224.0.0.2. Multicast router solicitations do not transmit if a router needs to discover multicast routers on a directly attached subnet.

Multicast router termination messages transmit after a router terminates its multicast routing functions. Other non-IP forwarding devices, such as Layer 2 switches, can send multicast router solicitations to solicit multicast router advertisements.

To function MRDISC on IGMP snoop interface, you must explicitly enable MRDISC. The Solicitation messages are sent only if IGMP snoop and MRDISC are enabled on the switch.

Multicast Flow Distribution over MLT

MultiLink Trunking (MLT) is a mechanism to distribute multicast streams over a multilink trunk and achieve an even distribution of the streams. The distribution is based on source-subnet and group addresses. In applications like television distribution, multicast traffic distribution is particularly important because the bandwidth requirements are substantial when you use a large number of television streams.

The switch enables this feature by default and you can not change the configuration.

Traffic Distribution

Traffic distribution distributes the streams on the multilink trunk links if an MLT configuration change occurs. For example, you can add or delete ports.

This feature distributes active streams according to the distribution algorithm on the multilink trunk links. This distribution can cause minor traffic interruptions. To minimize the effect of distribution of multicast traffic on the multilink trunks, the implementation does not move the streams to the appropriate links at the same time. Instead, it distributes a few streams at every time tick of the system.

To that end, after a multilink trunk port becomes inactive, this feature distributes all the streams on the multilink trunk ports based on the assignment provided by the distribution algorithm.

By default, distribution is enabled and you can not change the configuration.

For more information about MLT, see [MultiLink Trunking](#) on page 2660 .

Multicast Virtualization

Multicast provides simplified extension of internal video and data delivery to remote locations.

Virtualized multicast enables multiple VPN routing instances on devices and supports various unicast routing protocols so that you can provide the services of many virtual routers from one physical device.

You can configure multicast routing support with the Virtual Routing and Forwarding (VRF) Lite feature and you can use VRF Lite to emulate many virtual routers with one router.

Multicast virtualization support includes:

- IGMP snooping
- IGMP in Layer 2 virtual services networks (VSN)
- IGMP in Layer 3 VSNs

To implement multicast virtualization, you must perform the following tasks:

1. Create a VRF. For more information about how to create and configure a VRF, see [Create a VRF Instance](#) on page 4081.
2. Create a VLAN and associate it with the VRF.
3. Enable one of the following: IGMP snooping on the VLAN, Layer 2 VSN, or Layer 3 VSN.

If you use IGMP snooping on the VLAN, ensure the IGMP version on the multicast hosts or other network devices is either the same as the version on the VLAN, or enable compatibility mode.

Multicast virtualization does not support PIM. The switch supports IGMP with PIM only in the Global Router.

VRF Lite Background

VRF Lite provides independent IPv4 forwarding instances and independent routing instances (contexts), which can reside on the same or different VLANs and ports.

While forwarding and routing instances are mapped to IP interfaces, incoming traffic is classified into a VLAN and IP interface and, depending on the IP interface, routed context traffic is forwarded.

Protocol Independent Multicast-Sparse Mode

Table 146: Protocol Independent Multicast - Sparse Mode product support

Feature	Product	Release introduced
Protocol Independent Multicast-Sparse Mode (PIM-SM) for IPv4	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
PIM Infinite Threshold for IPv4 and IPv6	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

**Note**

PIM is supported in Global Routing Table (GRT) only.

PIM-SM, as defined in RFC2362, supports multicast groups spread out across large areas of a company or the Internet. PIM-SM sends multicast traffic only to routers that specifically join a multicast group. This technique reduces traffic flow over WAN links and overhead costs for processing unwanted multicast packets.

Dense-mode protocols use a flood-and-prune technique, which is efficient with densely-populated receivers. However, for sparsely populated networks, PIM-SM is more efficient because it sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic.

PIM-SM is independent of a specific unicast routing protocol, but it does require the presence of a unicast routing protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). PIM-SM uses the information from the unicast routing table to create and maintain multicast trees that enable PIM-enabled routers to communicate.

Typically, a PIM-SM network consists of several multipoint data streams, each targeted to a small number of LANs in the internetwork. For example, customers whose networks consist of multiple hosts on different LANs can use PIM-SM to simultaneously access a video data stream, such as video conferencing, on a different subnet.

**Important**

In some cases, PIM stream initialization can take several seconds.

Hosts

A host is a source, a receiver, or both:

- A source, also known as a sender, sends multicast data to a multicast group.
- A receiver receives multicast data from one or several sources that sends data to a multicast group.

PIM-SM Domain

PIM-SM operates in a domain of contiguous routers on which PIM-SM is enabled.

Each PIM-SM domain requires the following routers:

- designated router (DR)
- rendezvous point (RP) router
- bootstrap router (BSR)

Although a PIM-SM domain can use only one active RP router and one active BSR, you can configure additional routers as a candidate RP (C-RP) router and as a candidate BSR (C-BSR). Candidate routers provide backup protection in case the primary RP router or BSR fails.

As a redundancy option, you can configure several RPs for the same group in a PIM domain. As a load sharing option, you can have several RPs in a PIM-SM domain map to different groups. The switch devices use the hash function defined in the PIM-SM standard to elect the active RP.

Designated Router

The designated router (DR), the router with the highest IP address on a LAN, performs the following tasks:

- sends register messages to the RP router on behalf of directly connected sources

- sends join and prune messages to the RP router on behalf of directly connected receivers
- maintains information about the status of the active RP router for local sources in each multicast group

**Important**

The DR is not a required configuration. Switches act automatically as the DR for directly attached sources and receivers.

Rendezvous Point Router

PIM-SM builds a shared multicast distribution tree within each domain, and the RP router is at the root of this shared tree. Although you can physically locate the RP anywhere on the network, it must be as close to the source as possible. Only one active RP router exists for a multicast group.

At the RP router, receivers meet new sources. Sources use the RP to identify themselves to other routers on the network; receivers use the RP to learn about new sources.

The RP performs the following tasks:

- registers a source that wants to announce itself and send data to group members
- joins a receiver that wants to receive data for the group
- forwards data to group

Candidate Rendezvous Point Router

You can configure a set of routers as C-RP routers that serve as backup to the RP router. If an RP fails, all the routers in the domain apply the same algorithm to elect a new RP from the group of C-RP routers. To make sure that the routers use a complete list of C-RP routers, the C-RP router periodically sends unicast advertisement messages to the BSR. The most common implementation is to configure a PIM-SM router as both a C-RP router and a C-BSR.

The switch devices use the hash function defined in the PIM-SM standard to elect the active RP.

Static Rendezvous Point Router

You can configure a static entry for an RP router with static RP. This feature avoids the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. Static RP-enabled switches cannot learn about RPs through the BSR because the switch loses all dynamically learned BSR information and ignores BSR messages. After you configure static RP entries, the switch adds them to the RP set as if they were learned through the BSR.

**Important**

In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interface configured as an RP.

When you configure a PIM static RP in a switch, the next hop of the unicast route toward the PIM static RP must be a PIM neighbor. The PIM protocol fails to work, due to a route change, if the next hop toward an already configured static RP becomes a non-PIM neighbor. If a PIM neighbor cannot reach the configured RP, the RP does not activate and its state remains invalid.

A static RP-enabled switch can communicate with switches from other vendors that do not use the BSR mechanism. Some vendors use either early implementations of PIM-SM v1 that do not support the BSR or proprietary mechanisms like the Cisco Auto-RP. For a network to work properly with static RP, you must map all the switches in the network (including switches from other vendors) to the same RP or RPs, if several RPs exist in the network.

To avoid a single point of failure, you can also configure redundant static RPs.

Use the static RP feature when you do not need dynamic learning mode, typically in small networks, or for security reasons, where RPs are forced to devices in the network so that they do not learn other RPs.

Static RP Configuration Considerations

Before you can configure a static RP, you must enable PIM-SM and enable static RP.

After you meet these prerequisites, keep in mind the following configuration considerations:

- You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.
- All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.
- Static RPs do not age, that is, they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, you must configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP.
- In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interfaces configured as an RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of vendor switches across the network, you must ensure that all switches and routers use the same active RP because other vendors can use different algorithms to elect the active RP. The switch devices use the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.



Important

To reduce convergence times, create only one static RP for each group. The more static RPs you configure for redundancy, the more time PIM requires to rebuild the mroute table and associate RPs.

- Static RP configured on the switch is active as long as the switch uses a unicast route to the static RP network. If the switch loses this route, the static RP is invalidated and the hash algorithm remaps all affected groups. If the switch regains

this route, the static RP is validated and the hash algorithm remaps the affected groups.

Bootstrap Router

The BSR receives RP router advertisement messages from the candidate RPs. The BSR adds the RP router with its group prefix to the RP set. Only one BSR exists for each PIM-SM domain.

The BSR periodically sends bootstrap messages containing the complete RP set to all routers in the domain. The BSR ensures that all PIM-SM routers send join, prune, and register packets.

Within a PIM-SM domain, you can configure a small set of routers as C-BSRs. The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.



Important

Configure C-BSRs on routers that are central to all candidate RPs.

Shared Trees and Shortest-Path Trees

A PIM-SM domain uses shared trees and shortest-path trees to deliver data packets to group members. This section describes both trees.

Shared Trees

Group members in a PIM-SM domain receive the first packet of data from sources across a shared tree. A shared tree consists of a set of paths that connect all members of a multicast group to the RP. PIM creates a shared tree when sources and receivers send messages toward the RP.

Shortest-Path Trees

After receiving a certain number of packets from the RP, the DR changes from a shared tree to an SPT. Switching to an SPT creates a direct route between the receiver and the source. The switch changes to the SPT after it receives the first packet from the RP.

[Figure 128](#) shows a shared tree and an SPT.




Figure 128: Shared tree and shortest-path tree

Receiver Joining a Group

The following steps describe how a receiver joins a multicast group:

1. A receiver multicasts an IGMP host membership message to the group that it wants to join.
2. After the last-hop router (the DR), normally the PIM router with the highest IP address for that VLAN, receives the IGMP message for a new group join, the router looks up the associated elected RP with responsibility for the group.
3. After it determines the RP router for the group, the last-hop router creates a (*,G) route entry in the multicast forwarding table and sends a (*,G) join message to the RP. After the last-hop router receives data packets from the RP, if the multicast packet arrival rate exceeds the DR threshold, the last-hop router switches to the SPT by sending an (S,G) join message to the source. (S denotes the source unicast IP address, and G denotes the multicast group address.)
4. If the last-hop router switches to the SPT, the following actions occur:
 - All intermediate PIM routers along the path to the source create the (S,G) entry.
 - To trim the shared tree, the router sends an (S,G) prune message to the RP.

You can enable the PIM Infinite Threshold Policy feature to prevent the SPT switchover. Multicast traffic follows the shared tree path through a Rendezvous Point (RP) instead of switching over to SPT.

Receiver Leaving a Group

Before it leaves a multicast group, a receiver sends an IGMP leave message to the DR. If all directly connected members of a multicast group leave or time out, and no downstream members remain, the DR sends a prune message upstream and PIM-SM deletes the route entry after that entry times out.

When the system ages PIM mroutes, it does not clear the (S,G) entry for an inactive route immediately after the expiration period. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.

Source Sending Packets to a Group

The following steps describe how a source sends multicast packets to a group:

1. A source directly attached to a VLAN bridges the multicast data to the DR. The DR for the VLAN (the router with the highest IP address) encapsulates each packet in a register message and sends a unicast message directly to the RP router to distribute to the multicast group.
2. If a downstream group member chooses to receive multicast traffic, the RP router sends a join or prune message toward the source DR and forwards the data down the RP tree after it obtains the data natively.
3. After the receiver DR obtains the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.
4. If no downstream members want to receive multicast traffic, the RP router sends a register-stop message (for the source) to the DR.

The DR starts the register suppression timer after it receives the first register-stop message. During the register suppression timeout period (the default is 60 seconds), the following events occur:

- The DR for the source sends a probe packet to the RP router before the register suppression timer expires. The probe packet prompts the RP router to determine whether new downstream receivers joined the group.
- If no new receivers joined the group, the RP router sends another register-stop message to the DR for the source, and its register suppression timer restarts.
- After the RP router no longer responds with a register-stop message to the source DR probe message, the register suppression timer expires and the DR sends encapsulated multicast packets to the RP router. The RP router uses this method to tell the DR that new members joined the group.

The RP sends a register-stop message to the DR immediately after it receives the first multicast data packet.

Required Elements for PIM-SM Operation

For PIM-SM to operate, the following elements must exist in the PIM-SM domain:

- You must enable an underlying unicast routing protocol for the switch to provide routing table information to PIM-SM.
- You must configure an active BSR to send bootstrap messages to all PIM-v2 configured switches and routers to enable them to learn group-to-RP mapping. If

you configure several BSRs in a network, an active BSR is elected based on priority and IP address (if priority is equal, the BSR with the higher IP address is elected).

- You must include an RP to perform the following tasks:
 - manage one or several IP multicast groups
 - become the root for the shared tree to these groups
 - accept join messages from receiver switches for groups that it manages
 - elect an active RP based on priority and IP address (if priority is equal, the RP with the higher IP address is elected)

PIM-SM Simplified Example

Figure 129 shows a simplified example of a PIM-SM configuration.




Figure 129: PIM-SM simplified example

In the sample configuration, the following events occur:

1. The BSR distributes RP information to all switches in the network.
2. R sends an IGMP membership report to S4.
3. Acting on this report, S4 sends a (*,G) join message to RP.
4. S sends data to G.
5. The DR (S1 in this example) encapsulates the data that it unicasts to RP (S2) in register messages.
6. S2 decapsulates the data, which it forwards to S4.
7. S4 forwards the data to R.

8. If the packet rate exceeds the DR threshold, S4 sends S1 an (S,G) join message.
9. S1 forwards data to S4. After S4 receives data from S1, it prunes the stream from the RP.

**Important**

[Figure 129](#) on page 1727 is a simplified example and is not the best design for a network if you locate the source and receiver as shown. In general, place RPs as close as possible to sources.

PIM-SM Static Source Groups

You can configure static source groups as static source-group entries in the PIM-SM multicast routing table. PIM-SM cannot prune these entries from the distribution tree. For more information about static source groups, see [Static Source Groups](#) on page 1702.

Join and Prune Messages

The DR sends join and prune messages from a receiver toward an RP for the group to either join the shared tree or remove (prune) a branch from it. A single message contains both a join and a prune list. This list includes a set of source addresses that indicate the shortest-path trees or the shared trees that the host wants to join. The DR sends join and prune messages hop-by-hop to each PIM router on the path to the source or the RP.

Register and Register-Stop Messages

The DR sends register messages to the RP for a directly connected source. The register message informs the RP of a new source, causing the RP to send join or prune messages back toward the DR of the source, which forwards the data down the RP tree after it obtains the data natively. After the receiver DR obtains the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.

The DR stops sending encapsulated packets to the RP after it receives a register-stop message. This traffic stops without delay because the RP sends a register-stop message immediately after it receives the first multicast data packet, and joins the shortest-path tree.

PIM-SMLT

IP multicast routing support with Split MultiLink Trunking (SMLT) builds a virtual switch that represents the two switches of the split multilink trunk core.

When switches use PIM in the core, they need to exchange protocol-related updates as part of the interswitch trunking (IST) protocol. IST hides the fact that the edge switch attaches to two physical switches.

PIM-SMLT can work in triangular, square, and full mesh configurations with Layer 3 IP multicast. However, PIM-SSM in square or full mesh SMLT topologies is not supported.

The following rules apply:

- If a VLAN receives traffic from the IST link, it cannot forward on the split multilink trunk link or the edge for the same VLAN.
- If one side of the SMLT link toward the receiver is down, such that the traffic cannot be forwarded directly down the SMLT link from the router on which traffic is ingressing, the IST Peer **MUST** forward that traffic it receives over the IST link down its side of the SMLT toward the receiver. The decision of whether the IST Peer needs to forward traffic received over the IST to SMLT receivers is made in the datapath, which has full knowledge of the remote SMLT link state.
- Traffic can use the IST to route between VLANs if the forwarding decision for the multicast protocol requires that the other side of the core forwards the multicast traffic (follow the IP multicast routing and forwarding rules for routed traffic). Other VLANs that are not part of SMLT continue to function in the same way.
- To create a temporary default route pointing to a peer IST, you must enable PIM on the IST VLAN.
- In a scaled multicast environment, if you must reconfigure the members of an MLT link, either SMLT or IST, by removing the ports from the MLT membership list, you must first shutdown the port by using the **shutdown** command at the port configuration level. Let the unicast and multicast traffic subside, and then remove the port from the MLT membership list. If you reconfigure the MLT without first shutting down the port, it can lead to excessive hardware updates to multicast forwarding records and can result in high utilization of the CPU.



Note

In a scaled PIM over Simplified vIST deployment, disabling all the PIM interfaces (**no ip routing**) causes the VLACP ports to bounce. With no user intervention, the packets start getting processed again in approximately 10 seconds. VLACP enables the ports and full functionality is restored.

SMLT provides for fast fail over in all cases, but does not provide a functionality similar to Routed SMLT (RSMLT).



Important

You must enable square SMLT globally before you configure square or full-mesh configurations.

Traffic Delay with PIM While Restarting Peer SMLT Switches

If you restart peer SMLT switches, you can lose, or experience a delay in, PIM traffic. The local and remote SMLT links must be up to forward traffic. If a remote SMLT link is down, you can experience a traffic delay.

PIM uses a DR to forward data to receivers on a VLAN. If you restart the DR in an SMLT VLAN, you can lose data because of the following actions:

- If the DR is down, the non-DR switch assumes the role and starts forwarding data.
- After the DR comes back up, it takes priority (higher IP address) to forward data so the non-DR switch stops forwarding data.

- The DR is not ready to forward traffic due to protocol convergence and because it takes time to learn the RP set and create the forwarding path. This situation can result in a traffic delay of 2 to 3 minutes because the DR learns the RP set after Open Shortest Path First (OSPF) converges.

A workaround to this delay is to configure the static RP router on the peer SMLT switches. This feature avoids the process of selecting an active RP router from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. After the DR comes back up, traffic resumes as soon as OSPF converges. This workaround reduces the traffic delay to approximately 15 to 65 seconds.

Protocol Independent Multicast-Source Specific Multicast

Table 147: Protocol Independent Multicast-Source Specific Mode

Feature	Product	Release introduced
PIM-Source Specific Mode (PIM-SSM) for IPv4	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0



Note

PIM is supported in Global Routing Table (GRT) only.

Source Specific Multicast optimizes PIM-SM by simplifying the many-to-many model. Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM only builds source-based SPTs. Whereas PIM-SM always joins a shared tree first, and then switches to the source tree, SSM eliminates the need to start with a shared tree by immediately joining a source through the SPT. SSM avoids using an RP and RP-based shared trees, which can be a potential problem.

You can configure multiple channels for the SSM group members in the SSM map table.

This configuration is ideal for applications like television channel distribution and other content-distribution businesses. Banking and trade applications can also use SSM as it provides more control over the hosts receiving and sending data over their networks.

When a v2 report in SSM range is received it is translated to an igmpv3 report message with one group record with type ALLOW and the source lists copied from the igmp ssm map static entries and passed to igmpv3 module. When a v2 leave in SSM range is received it is translated to an igmpv3 report message with one group record with type BLOCK and the source lists copied from the igmp ssm map static entries and passed to igmpv3 module. This behavior is displayed only when PIM-SSM mode is enabled.

SSM applications use IP addresses reserved by the Internet Assigned Numbers Authority (IANA) in the 232/8 range (232.0.0.0 to 232.255.255.255). SSM recognizes packets in this range and controls the behavior of multicast routing devices and hosts that use these addresses. When a source (S) transmits IP datagrams to an SSM destination address (G), a receiver can receive these datagrams by subscribing to the (S,G) channel.

A channel is a source-group (S,G) pair where S is the source that sends to the multicast group and G is an SSM group address. SSM defines channels on an individual or multiple source basis, which enforces the one-to-many concept of SSM applications. In an SSM channel, each group is associated with multiple sources.

SSM Features

PIM-SM requires a unicast protocol to forward multicast traffic within the network to perform the Reverse Path Forwarding (RPF) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that PIM-enabled routers use to communicate. The unicast routing table must contain a route to every multicast source in the network as well as routes to PIM entities like the RPs and BSR.

SSM uses only a subset of the PIM-SM features such as the SPT, DR, and some messages (hello, join, prune, and assert). However, some features are unique to SSM. These features, described in the following sections, are extensions of the IGMP and PIM protocols.

PIM-SSM Architecture

The following diagram illustrates how the PIM-SSM architecture requires routers to perform the following actions:

- support IGMPv3 source-specific host membership reports and queries at the edge routers
- initiate PIM-SSM (S,G) joins directly and immediately after receiving an IGMPv3 join report from the designated router
- restrict forwarding to SPTs within the SSM address range by all PIM-SSM routers




Figure 130: PIM-SSM architecture

The following rules apply to Layer 3 devices with SSM enabled:

- Receive IGMPv3 membership join reports in the SSM range and, if no entry (S,G) exists in the SSM channel table, create one.
- Receive IGMPv2 membership join reports, but only for groups that already use a static (S,G) entry in the SSM channel table.
- Send periodic join messages to maintain a steady SSM tree state.
- Use standard PIM-SM SPT procedures for unicast routing changes, but ignore rules associated with the SPT for the (S,G) route entry.
- Receive prune messages and use standard PIM-SM procedures to remove interfaces from the source tree.
- Forward data packets to interfaces from the downstream neighbors that sent an SSM join, or to interfaces with locally attached SSM group members.
- drop data packets that do not use an exact-match lookup (S,G) in their forwarding database for S and G.

PIM-SSM Static Source Groups

You can configure static source group entries in the PIM-SSM multicast routing table with static source groups. PIM-SSM cannot prune these entries from the distribution tree. For more information about static source groups, see [Static Source Groups](#) on page 1702.

Implementation of SSM and IGMP

The following sections describe how the switch implements PIM-SSM and IGMP.

SSM Range

The standard SSM range is 232/8, but you can extend the range to include an IP multicast address. Although you can configure the SSM range, you cannot configure it for all multicast groups (224/4 or 224.0.0.0/240.0.0.0 or 224.0.0.0/255.0.0.0).

You can extend the SSM range to configure existing applications without changing their group configurations.

SSM Channel Table

You can use the SSM channel to manually configure (S,G) entries that map existing groups to their sending source. These table entries apply to the whole switch, not for each interface, and both IGMPv2 and IGMPv3 hosts use the SSM channel table.

The following rule applies to an SSM channel table for an individual switch:

- You can map one source to multiple groups.
- You can allow multiple sources to the same group.



Important

Different switches can use different mappings for groups to sources, for example, different channels map differently even if they are on the same network.

SSM and IGMPv2

SSM-configured switches can accept reports from IGMPv2 hosts on IGMPv2 interfaces if the group uses an SSM channel table entry. However, the IGMPv2 host groups must exist in the SSM range defined on the switch, which is 232/8 by default.

- After the SSM switch receives an IGMPv2 report for a group that is in the SSM channel table, it joins the specified source immediately.
- After the SSM switch receives an IGMPv2 report for a group that uses an enabled static SSM channel table entry, it triggers PIM-SSM processing as if it received an equivalent IGMPv3 report.
- After the SSM switch receives an IGMPv2 report for a group out of the SSM range, it processes the report as if it is in PIM-SM mode.

Deleting or Disabling an ssm-map with IGMPv1 or IGMPv2

Before you disable or delete an ssm-map, always send IGMPv1 or IGMPv2 leave messages from hosts that operate using IGMPv1 or IGMPv2. If you do not perform this

action, receiving and processing reports in SSM range on an IGMP interface enabled with IGMPv1 or IGMPv2 can lead to unexpected behavior.

Consider the following configuration scenario:

- A device is PIM-enabled, running in SSM mode, with IGMPv1 or IGMPv2 configured on the interface.
- IGMPv1 and IGMPv2 hosts send IGMPv1 or IGMPv2 reports for groups in SSM range.

The following table identifies the expected behaviors in this scenario.

Table 148: Expected behaviors for ssm-map configuration

Action	Expected behavior
You do not configure an ssm-map for the group in SSM range.	IGMPv1 and IGMPv2 reports are not processed.
You do configure an ssm-map for the group in SSM range.	IGMPv1 and IGMPv2 reports are processed and the group in SSM range is learned.

SSM and IGMPv3

The switch supports IGMPv3 for SSM. With IGMPv3, a host can selectively request or filter traffic from sources within the multicast group. IGMPv3 is an interface-level configuration.



Important

IGMPv3 works without PIM-SSM or SSM-snoop enabled on the interface.

The following rules apply to IGMPv3-enabled interfaces:

- Send only IGMPv3 (source-specific) reports for addresses in the SSM range.
- Accept IGMPv3 reports.
- Drop IGMPv2 reports.

The IGMPv2 report mentioned in [SSM and IGMPv2](#) on page 1733 is processed because it is an IGMPv2 report received on an IGMPv2 interface. If an IGMPv2 interface receives an IGMPv3 report, it drops the report even if PIM-SSM is enabled and the entry is in the SSM channel table. The IGMP versions must match.

- Discard IGMP packets with a group address out of the SSM range.

The switch implements IGMPv3 in one of two modes: dynamic and static.

In dynamic mode, the switch learns about new (S,G) pairs from IGMPv3 reports and adds them to the SSM channel table. If you do not enable dynamic mode and an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report.

In static mode, you can statically configure (S,G) entries in the SSM channel table. If an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM

channel table, it ignores the report. The interface also ignores the report if the group is in the table, but the source or mask does not match what is in the table.



Important

After you enable IGMPv3, changes to the query interval and robustness values on the querier switch propagate to other switches on the same VLAN through IGMP query.

Both IGMPv2 and IGMPv3 hosts use the SSM channel table:

- An IGMPv2 host (with an IGMPv2 VLAN) must use an existing SSM channel entry if the group is in the SSM range.
- If you enable dynamic learning for an IGMPv3 host, the SSM channel automatically learns the group. Otherwise, the SSM channel also needs a static entry.

The following table summarizes how a switch in PIM-SSM mode works with IGMP if you disable IGMPv3 compatibility. In the following table, references to matching a static SSM channel entry assumes that the entry is enabled. If an entry is disabled, it is treated as though it is disallowed.

Table 149: PIM-SSM interaction with IGMPv2 and v3 with IGMPv3 compatibility disabled

Host	VLAN	SSM range	Action
IGMPv2 host	IGMPv3 VLAN	In or out of range	Drop report.
IGMPv3 host	IGMPv2 VLAN	In or out of range	Drop report.
IGMPv2 host	IGMPv2 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 host	IGMPv2 VLAN	Out of range	Ignore the SSM channel table and process the report as if it is in PIM-SM mode.
IGMPv3 host	IGMPv3 VLAN	Out of range	Process the report.
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic enabled. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and matches an existing SSM channel entry. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and does not match an existing SSM channel entry. Drop report.

The following table summarizes how a switch in PIM-SSM mode works with IGMP if you enable IGMPv3 compatibility.

Table 150: PIM-SSM interaction with IGMPv2 and v3 with IGMPv3 compatibility enabled

Host	VLAN	SSM range	Action
IGMPv2 Host	IGMPv3 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 Host	IGMPv3 VLAN	Out of range	Process the report as in PIM-SM mode.

If an IGMPv3 group report enters the VLAN port and the port must discard one or more of the groups in that packet after the application of IGMP access controls, the port drops the entire packet and does not forward it on to other ports of the VLAN.

If an IGMPv3 interface receives an IGMPv2 or v1 query, the interface backs down to IGMPv2 or v1. As a result, the interface flushes all senders and receivers on the interface.

Configuration Limitations

Run PIM-SSM on either all switches in the domain or only on the edge routers. If you use a mix of PIM-SSM and PIM-SM switches in the domain, run PIM-SSM on all the edge routers and run PIM-SM on all the core routers.



Important

A PIM domain with edge routers running PIM-SM and core routers running PIM-SSM does not operate properly. If you prefer or require a mixed PIM-SM and PIM-SSM topology, run PIM-SSM on the edge switches and PIM-SM in the core. Ensure a valid RP configuration exists for groups that exist outside of the SSM range. If a valid RP configuration exists, the SSM switches process the joins in SM mode. If no RP exists, the SSM switches drop the reports.

Static source groups cannot conflict with SSM channels. If you configure a static source group or an SSM channel, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) to different sources or multiple groups to a single source for both static source group and an SSM channel.

PIM Passive Interfaces

You can configure the PIM interface as active or passive. The default is active. With an active interface, you can configure transmit and receive PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This

feature is useful when you use a high number of PIM interfaces and these interfaces connect to end users, not to other switches.

A PIM passive interface does not transmit and drops messages of the following type:

- hello
- join
- prune
- register
- register-stop
- assert
- candidate-RP-advertisement
- bootstrap

If a PIM passive interface receives these types of messages, it drops them and the switch logs a message, detailing the type of protocol message and the IP address of the sending device. These log messages help to identify the device that performs routing on the interface, which is useful if you must disable a device that does not operate correctly.

**Important**

A device can send register and register-stop messages to a PIM passive interface, but these messages cannot be sent out of that interface.

The PIM passive interface maintains information about hosts, through IGMP, that are related to senders and receivers, but the interface does not maintain information about PIM neighbors. You can configure a BSR or an RP on a PIM passive interface.

You can also use the PIM passive interface feature as a security measure to prevent routing devices from becoming attached and participating in the multicast routing of the network.

**Important**

Before you change the state (active or passive) of a PIM interface, disable PIM on that interface. This action prevents instability in the PIM operations, especially when neighbors exist or the interface receives streams. After you disable PIM, the switch loses traffic for approximately 80 seconds.

Multicast Route Statistics

Table 151: Mroute statistics product support

Feature	Product	Release introduced
Multicast route (mroute) statistics for IPv4 and IPv6	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6 Not Supported on 5320-16P-2MXT-2X
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The multicast route statistics feature provides statistics for multicast streams through the switch. Using the Command Line Interface (CLI), Simple Network Management Protocol (SNMP) or Enterprise Device Manager (EDM), you can track the number of senders sending multicast streams to a particular group address. You can also obtain a count of the packets or bytes being received for a particular multicast group address and the average size of the frames. Multicast route statistics are supported for both IPv4 and IPv6 group addresses.

Determining the route statistics is especially useful when debugging a multicast network and also when administering the network.

Multicast Route Statistics and DvR

When you enable or clear IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For more information on DvR, see [Distributed Virtual Routing Fundamentals](#) on page 919.

IP Multicast Network Design

Use multicast routing protocols to efficiently distribute a single data source among multiple users in the network. This section provides information about how to design networks that support IP multicast routing.

For more design guidelines, conceptual, and configuration information about IP Multicast over Fabric Connect, see [IP Multicast over Fabric Connect](#) on page 1703.

Multicast Scalability Design Rules

The following section lists the design rules to increase multicast route scaling.

**Important**

The switch software supports the following:

- Protocol-Independent Multicast (PIM)
- Split MultiLink Trunking (SMLT) and Routed-SMLT (RSMLT)

1. Whenever possible, use simple network designs that do not use VLANs that span several switches. Instead, use routed links to connect switches.
2. Whenever possible, group sources sending to the same group in the same subnet. The switch uses a single egress forwarding pointer for all sources in the same subnet sending to the same group. Be aware that these streams have separate hardware forwarding records on the ingress side.
3. Do not configure multicast routing on edge switch interfaces that do not contain multicast senders or receivers. By following this rule, you:
 - Provide secure control over multicast traffic that enters or exits the interface.
 - Reduce the load on the switch, as well as the number of routes. This improves overall performance and scalability.
4. Avoid initializing many (several hundred) multicast streams simultaneously. Initial stream setup is a resource-intensive task, and initializing a large number can increase the setup time. In some cases, this delay can result in stream loss.
5. Whenever possible, do not connect IP multicast sources and receivers by using VLANs that interconnect switches (see the following figure). In some cases, this can result in excessive hardware record use. By placing the source on the interconnected VLAN, traffic takes two paths to the destination, depending on the reverse path forwarding (RPF) checks and the shortest path to the source.

For example, if a receiver is on VLAN 1 on switch S1 and another receiver is on VLAN 2 on switch S1, traffic can be received from two different paths to the two receivers, which results in the use of two forwarding records. If the source on switch S2 is on a different VLAN than VLAN 3, traffic takes a single path to switch S1 where the receivers are located.




Figure 131: IP multicast sources and receivers on interconnected VLANs

IP Multicast Address Range Restrictions

IP multicast routers use D class addresses, which range from 224.0.0.0 to 239.255.255.255. Although you can use subnet masks to configure IP multicast address ranges, the concept of subnets does not exist for multicast group addresses. Consequently, the usual unicast conventions—where you reserve the all 0s subnets, all 1s subnets, all 0s host addresses, and all 1s host addresses—do not apply.

Internet Assigned Numbers Authority (IANA) reserves addresses from 224.0.0.0 through 224.0.0.255 for link-local network applications. Multicast-capable routers do not forward packets with an address in this range. For example, Open Shortest Path First (OSPF) uses 224.0.0.5 and 224.0.0.6, and Virtual Router Redundancy Protocol (VRRP) uses 224.0.0.18 to communicate across local broadcast network segments.

IANA also reserves the range of 224.0.1.0 through 224.0.1.255 for well-known applications. IANA assigns these addresses to specific network applications. For example, the Network Time Protocol (NTP) uses 224.0.1.1, and Mtrace uses 224.0.1.32. RFC1700 contains a complete list of these reserved addresses.

Multicast addresses in the 232.0.0.0/8 (232.0.0.0 to 232.255.255.255) range are reserved only for source-specific multicast (SSM) applications, such as one-to-many applications. While this range is the publicly reserved range for SSM applications, private networks can use other address ranges for SSM.

Finally, addresses in the range 239.0.0.0/8 (239.0.0.0 to 239.255.255.255) are administratively scoped addresses; they are reserved for use in private domains. Do not advertise these addresses outside the private domain. This multicast range is analogous to the 10.0.0.0/8, 172.16.0.0/20, and 192.168.0.0/16 private address ranges in the unicast IP space.

In a private network, only assign multicast addresses from 224.0.2.0 through 238.255.255.255 to applications that are publicly accessible on the Internet. Assign addresses in the 239.0.0.0/8 range to multicast applications that are not publicly accessible.

Although you can use a multicast address you choose on your own private network, it is generally not good design practice to allocate public addresses to private network entities. Do not use public addresses for unicast host or multicast group addresses on private networks.

Multicast MAC Address Mapping Considerations

Like IP, Ethernet has a range of multicast MAC addresses that natively support Layer 2 multicast capabilities. While IP has a total of 28 addressing bits available for multicast addresses, Ethernet has only 23 addressing bits assigned to IP multicast. The Ethernet multicast MAC address space is much larger than 23 bits, but only a subrange of that larger space is allocated to IP multicast. Because of this difference, 32 IP multicast addresses map to one Ethernet multicast MAC address.

IP multicast addresses map to Ethernet multicast MAC addresses by placing the low-order 23 bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01:00:5E:00:00:00. Thus, more than one multicast address maps to the same Ethernet address (see the following figure). For example, all 32 addresses 224.1.1.1, 224.129.1.1, 225.1.1.1, 225.129.1.1, 239.1.1.1, 239.129.1.1 map to the same 01:00:5E:01:01:01 multicast MAC address.




Figure 132: Multicast IP address to MAC address mapping

Most Ethernet switches handle Ethernet multicast by mapping a multicast MAC address to multiple switch ports in the MAC address table. Therefore, when you design the group addresses for multicast applications, take care to efficiently distribute streams only to hosts that are receivers.

As an example, consider two active multicast streams using addresses 239.1.1.1 and 239.129.1.1. Suppose that two Ethernet hosts, receiver A and receiver B, connect to ports on the same switch and only want the stream addressed to 239.1.1.1. Suppose also that two other Ethernet hosts, receiver C and receiver D, also connect to the ports on the same switch as receiver A and B, and want to receive the stream addressed to

239.129.1.1. If the switch uses the Ethernet multicast MAC address to make forwarding decisions, then all four receivers receive both streams—even though each host only wants one stream. This transmission increases the load on both the hosts and the switch. To avoid this extra load, ensure that you manage the IP multicast group addresses used on the network.

When an IP multicast packet is received, the lookup is based on the IP group address, regardless of whether the VLAN is bridged or routed. This problem is particularly true of pure Layer 2 switches.

In a network that includes multiple hardware platforms, the easiest way to ensure that this issue does not arise is to use only a consecutive range of IP multicast addresses that correspond to the lower-order 23 bits of that range. For example, use an address range from 239.0.0.0 through 239.127.255.255. A group address range of this size can still easily accommodate the needs of even the largest private enterprise.

Dynamic Multicast Configuration Changes

You must not perform dynamic multicast configuration changes when multicast streams flow in a network. For example, do not change the routing protocol that runs on an interface, or the IP address, or the subnet mask for an interface until multicast traffic ceases.

For such changes, ensure that you temporarily stop all multicast traffic. If the changes are necessary and you have no control over the applications that send multicast data, you can disable the multicast routing protocols before you perform the change. For example, consider disabling multicast routing before making interface address changes. In all cases, these changes result in traffic interruptions because they affect neighbor-state machines and stream-state machines.

In addition, when removing port members of an MLT group you must first disable the ports. Changing the group set without first shutting the ports down can result in high-CPU utilization and processing in a scaled multicast environment due to the necessary hardware reprogramming on the multicast records.

IGMPv3 Backward Compatibility

IGMPv3 for PIM is backward compatible with IGMPv1/v2. According to RFC3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.
- If a router that is not explicitly configured to use IGMPv1 or IGMPv2, detects an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure the IGMP version of an interface to version 3 regardless of the PIM or snooping mode.

You can configure whether the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does

not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

**Note**

If you enable the explicit host tracking option on an IGMPv3 interface, you cannot downgrade to IGMPv1 or IGMPv2. You must disable explicit host tracking to downgrade the IGMP version.

TTL in IP Multicast Packets

The switch treats multicast data packets with a time-to-live (TTL) of 1 as expired packets and sends them to the CPU before dropping them. To avoid this issue, ensure that the originating application uses a hop count large enough to enable the multicast stream to traverse the network and reach all destinations without reaching a TTL of 1. Ensure that you use a TTL value of 33 or 34 to minimize the effect of looping in an unstable network.

Multicast MAC Filtering

Certain network applications require multiple hosts to share a multicast MAC address. Instead of flooding all ports in the VLAN with this multicast traffic, you can use Multicast MAC filtering to forward traffic to a configured subset of the ports in the VLAN. This multicast MAC address is not an IP multicast MAC address.

At a minimum, map the multicast MAC address to a set of ports within the VLAN. In addition, if traffic is routed on the local host, you must configure an Address Resolution Protocol (ARP) entry to map the shared unicast IP address to the shared multicast MAC address. You must configure an ARP entry because the hosts can also share a virtual IP address, and packets addressed to the virtual IP address need to reach each host.

Ensure that you limit the number of such configured multicast MAC addresses to a maximum of 100. This number is related to the maximum number of possible VLANs you can configure, because for every multicast MAC filter that you configure the maximum number of configurable VLANs reduces by one. Similarly, configuring large numbers of VLANs reduces the maximum number of configurable multicast MAC filters downward from 100.

Although you can configure addresses starting with 01:00:5E, which are reserved for IP multicast address mapping, do not enable IP multicast with streams that match the configured addresses. This configuration can result in incorrect IP multicast forwarding and incorrect multicast MAC filtering.

Guidelines for Multicast Access Policies

Use the following guidelines when you configure multicast access policies:

- Use masks to specify a range of hosts. For example, 10.177.10.8 with a mask of 255.255.255.248 matches hosts addresses 10.177.10.8 through 10.177.10.15. The host subnet address and the host mask must be equal to the host subnet address. An easy way to determine this is to ensure that the mask has an equal or fewer number

of trailing zeros than the host subnet address. For example, 3.3.0.0/255.255.0.0 and 3.3.0.0/255.255.255.0 are valid. However, 3.3.0.0/255.0.0.0 is not.

- Apply receive-access policies to all eligible receivers on a segment. Otherwise, one host joining a group makes that multicast stream available to all.
- Receive access policies are initiated after the switch receives reports with addresses that match the filter criteria.
- Transmit access policies apply after the switch receives the first packet of a multicast stream.

Multicast access policies can apply to a routed PIM interface if Internet Group Management Protocol (IGMP) reports the reception of multicast traffic.

The following rules and limitations apply to IGMP access policy parameters when you use them with IGMP instead of PIM:

- The static member parameter applies to IGMP snooping and PIM on both interconnected links and edge ports.
- The Static Not Allowed to Join parameter applies to IGMP snooping and PIM on both interconnected links and edge ports.
- For multicast access control, the denyRx parameter applies to IGMP snooping and PIM. The DenyTx and DenyBoth parameters apply only to IGMP snooping.

Split-Subnet and Multicast

The split-subnet issue arises when you divide a subnet into two unconnected sections in a network. This division results in the production of erroneous routing information about how to reach the hosts on that subnet. The split-subnet problem applies to all types of traffic, but it has a larger impact on a PIM-SM network.

To avoid the split-subnet problem in PIM networks, ensure that the RP router is not in a subnet that can become a split subnet. Also, avoid having receivers on this subnet. Because the RP is an entity that must be reached by all PIM-enabled switches with receivers in a network, placing the RP on a split-subnet can impact the whole multicast traffic flow. Traffic can be affected even for receivers and senders that are not part of the split-subnet.

Protocol Independent Multicast-Sparse Mode Guidelines

Protocol Independent Multicast-Sparse Mode (PIM-SM) uses an underlying unicast routing information base to perform multicast routing. PIM-SM builds unidirectional shared trees rooted at a RP router for each group and can also create shortest-path trees for each source.

PIM-SM and PIM-SSM Scalability

For more information on interface scaling, see the *Fabric Engine Release Notes*.

The software does not support virtualized PIM. PIM is supported in the Global Routing Table only.

Interfaces that run PIM must also use a unicast routing protocol (PIM uses the unicast routing table), which puts stringent requirements on the system. With a high number of interfaces, take special care to reduce the load on the system.

Use few active IP routed interfaces. You can use IP forwarding without a routing protocol enabled on the interfaces, and enable only one or two with a routing protocol. You can configure proper routing by using IP routing policies to announce and accept routes on the switch. Use PIM passive interfaces on the majority of interfaces.

**Important**

For information on the maximum values for total PIM interfaces and active interfaces, see the *Fabric Engine Release Notes*. If you configure the maximum number of active interfaces, all remaining interfaces must be passive.

When you use PIM-SM, the number of routes can scale up to the unicast route limit because PIM uses the unicast routing table to make forwarding decisions. For higher route scaling, use OSPF instead of Routing Information Protocol (RIP).

As a general rule, a well-designed network does not have many routes in the routing table. For PIM to work properly, ensure that all subnets configured with PIM are reachable and that PIM uses the information in the unicast routing table. For the RPF check, to correctly reach the source of any multicast traffic, PIM requires the unicast routing table.

PIM General Requirements

Design simple PIM networks where VLANs do not span several switches.

PIM relies on unicast routing protocols to perform its multicast forwarding. As a result, include in your PIM network design, a unicast design where the unicast routing table has a route to every source and receiver of multicast traffic, as well as a route to the RP router and Bootstrap router (BSR) in the network. Ensure that the path between a sender and receiver contains PIM-enabled interfaces. Receiver subnets are not always required in the routing table.

Use the following guidelines:


- Ensure that every PIM-SM domain is configured with an RP, either by static definition or via BSR.
- Ensure that every group address used in multicast applications has an RP in the network.
- As a redundancy option, you can configure several RPs for the same group in a PIM domain.
- As a load sharing option, you can have several RPs in a PIM-SM domain map to different groups.
- In order to configure an RP to cover the entire multicast range, configure an RP to use the IP address of 224.0.0.0 and the mask of 240.0.0.0.
- Configure an RP to handle a range of multicast groups by using the mask parameter. For example, an entry for group value of 224.1.1.0 with a mask of 255.255.255.192 covers groups 224.1.1.0 to 224.1.1.63.


- In a PIM domain with both static and dynamic RP switches, you cannot configure one of the (local) interfaces for the static RP switches as the RP. For example, in the following scenario:

(static RP switch) Sw1 ----- Sw2 (BSR/Cand-RP1) -----Sw3

You cannot configure one of the interfaces on switch Sw1 as static RP because the BSR cannot learn this information and propagate it to Sw2 and Sw3. PIM requires that you consistently configure RP on all the routers of the PIM domain, so you can only add the remote interface Candidate-RP1 (Cand-RP) to the static RP table on Sw1.

- If a switch needs to learn an RP-set, and has a unicast route to reach the BSR through this switch, you cannot enable or configure static RP on a switch in a mixed mode of candidate RP and static RP switches. For examples, see the following two figures.





PIM and Shared Tree to Shortest Path Tree Switchover

When an IGMP receiver joins a multicast group, PIM on the leaf router first joins the shared tree. After the first packet is received on the shared tree, the router uses the source address information in the packet to immediately switch over to the shortest path tree (SPT). If you enable PIM Infinite Threshold Policy for IPv4 and IPv6, multicast traffic follows the shared tree path through a Rendezvous Point (RP) instead of switching immediately over to the SPT.

PIM Traffic Delay and SMLT Peer Reboot

PIM uses a designated router (DR) to forward data to receivers on the DR VLAN. The DR is the router with the highest IP address on a LAN. If this router is down, the router with the next highest IP address becomes the DR. However, if the VLAN is an SMLT VLAN, the DR is not a factor in determining which switch forwards the data down to the receiver. Either aggregate switch can forward data to the receiver, because the switches act as one. The switch that forwards depends on where the source is located (on another SMLT/vIST link or on a non-SMLT/non-vIST link) and whether either side of the receiver SMLT link is up or down. If the forwarder switch is rebooted, traffic loss occurs until protocol convergence is completed.

Consider the following cases:

- If the source is on an SMLT link that is not the receiver SMLT, the switch that directly received the data on its side of the source SMLT link forwards it down to the receiver on the receiver SMLT regardless of which switch is the DR for the receiver VLAN. The forwarding switch sends a copy of the data over the vIST link to the peer switch, which drops the data because it knows that the remote SMLT is up and therefore the remote peer has already forwarded the data. If the forwarding switch goes down, the other switch receives the data directly over its source SMLT link and takes over forwarding to the receivers. After the original switch comes back up, the original switch again receives the data directly over its source SMLT. The original switch may not be ready to forward the data because of the protocol reconvergence, so the original switch loses traffic until reconvergence is complete.

- If the source is not learned on another SMLT link or the vIST link on each aggregate switch; they have a route to the source which is not on an SMLT or across the vIST. The switches must choose which one forwards the data down the receiver SMLT link; which one is the designated forwarder, so that duplicate data does not occur. The highest IP address is the designated forwarder. If the designated forwarder becomes disabled, the other takes over. When it is reenabled, the other switch sees that it is no longer the highest IP address and it sees that the remote SMLT link comes up. The other switch then assumes that the vIST peer is capable of being the designated forwarder and it stops forwarding down to the receivers. If the original switch is not ready to forward the data due to reconvergence, traffic loss occurs.

In either case, configuring a static RP helps the situation. To avoid this traffic delay, a workaround is to configure a static RP on the peer SMLT switches. This configuration avoids the process of selecting an active RP router from the list of candidate RPs, and also of dynamically learning about RPs through the BSR mechanism. Then, when the DR comes back, traffic resumes as soon as OSPF converges. This workaround reduces the traffic delay.


Circuitless IP for PIM-SM

Use CLIP to configure a resilient RP and BSR for a PIM network. When you configure an RP or BSR on a regular interface, if it becomes nonoperational, the RP and BSR also become nonoperational. This status results in the election of other redundant RPs and BSRs, and can disrupt IP multicast traffic flow in the network. As a best practice for multicast networks design, always configure the RP and BSR on a CLIP interface to prevent a single interface failure from causing these entities to fail.

Also, configure redundant RPs and BSRs on different switches such that these entities are on CLIP interfaces. For the successful setup of multicast streams, ensure that a unicast route exists to all CLIP interfaces from all locations in the network. A unicast route is mandatory because, for proper RP learning and stream setup on the shared RP tree, every switch in the network needs to reach the RP and BSR. You can use PIM-SM CLIP interfaces only for RP and BSR configurations, and are not intended for other purposes.

Do not configure non-SMLT IGMP leaf ports on a router to be one of the redundant RP CLIP devices. It is possible that these IGMP hosts can become isolated from the multicast data stream(s).

If you configure dual-redundant RPs (vIST peers with the same CLIP interface IP address used for the RP), the topology in the following figure does not work in link-failure scenarios. Use caution if you design a network with this topology where the vIST peers are PIM enabled, and the source and receiver edges are Layer 2.



Consider an example where one of the peers, vIST-A, is the PIM DR for the source VLAN, and the source data is hashed to vIST-A from the Layer 2 source edge. vIST-A forwards traffic to the receiver edge using the SMLT link from vIST-A to the receiver edge. If the SMLT link fails, vIST-A does not forward traffic over the vIST link to vIST-B, and the receiver edge does receive the data.

In this topology, the receiver edge sends an IGMP membership report for a group, which is recorded on both vIST peers as an IGMP LEAF on the receiver SMLT port on the receiver VLAN.

Because both of the vIST peers are the RP for the group, they do not send a (*,g) PIM JOIN message toward the other RP. The (*,g) PIM mroute does not record the vIST port as a JOIN port on either vIST device. The PIM (*,g) mroute records only a LEAF on the SMLT receiver port.

Because the source is local (Layer 2 edge), there is no PIM (s,g) JOIN message toward the source and the (s,g) PIM mroute does not record the vIST port as a JOIN port on either vIST device. The PIM (s,g) mroute records only a LEAF on the SMLT receiver port.

If the source is hashed to vIST-A, the PIM DR for the incoming VLAN, traffic is forwarded to the receiver correctly. vIST-A does not forward traffic over the vIST to vIST-B, because no JOIN exists on the vIST port. If the receiver SMLT link from the vIST-A peer is down, the traffic is not forwarded to vIST-B, and is not received by the receiver edge. Traffic resumes after the link is restored. If the source data hashes to the non-DR peer, vIST-B, no problem occurs because the non-DR always forwards traffic to the DR.

A similar situation exists in this topology when vIST-A is both the RP and the DR for the Layer 2 receiver edge. The vIST port is not in the outgoing port list because there is no JOIN message from the peer toward the source (which is not PIM enabled). Therefore, if the SMLT link from vIST-A to the receiver edge is down, the system does not forward traffic to the peer vIST-B and down to the receiver.

You can avoid the preceding problems with this topology by performing one of the following actions:

- Enable PIM on the source edge.

The vIST peers send PIM joins toward the source and the JOIN is recorded on the vIST port for the (s,g). Data is forwarded to the peer.

- Do not configure dual redundant RPs.

One vIST peer is the RP for a group.

- Do not configure one vIST peer as both the DR for the source VLAN and the RP for the receiver group.

The system forwards the traffic to the RP or to the DR, depending on which peer receives the source, and, if the SMLT link to the receiver goes down there will be no data loss.

PIM-SM and Static RP

Use static RP to provide security, interoperability, and redundancy for PIM-SM multicast networks. Consider if the administrative ease derived from using dynamic RP assignment is worth the security risks involved. For example, if an unauthorized user connects a PIM-SM router that advertises itself as a candidate RP (C-RP), it can possibly take over new multicast streams that otherwise distribute through an authorized RP. If security is important, use static RP assignment.

You can use the static RP feature in a PIM environment with devices that run legacy PIM-SMv1 and Cisco Auto-RP. For faster convergence, you can also use static RP in a PIM-SMv2 environment. If you configure static RP with PIM-SMv2, the BSR is not active.

Static RP and Auto-RP

Some legacy PIM-SMv1 networks use the auto-RP protocol. Auto-RP is a Cisco proprietary protocol that provides equivalent functionality to the legacy platform supported PIM-SM RP and BSR. You can use the static RP feature to interoperate in this environment. For example, in a mixed-vendor network, you can use auto-RP among routers that support the protocol, while other routers use static RP. In such a network, ensure that the static RP configuration mimics the information that is dynamically distributed to guarantee that multicast traffic is delivered to all parts of the network.

In a mixed auto-RP and static RP network, ensure that the legacy platform does not serve as an RP because it does not support the auto-RP protocol. In this type of network, the RP must support the auto-RP protocol.

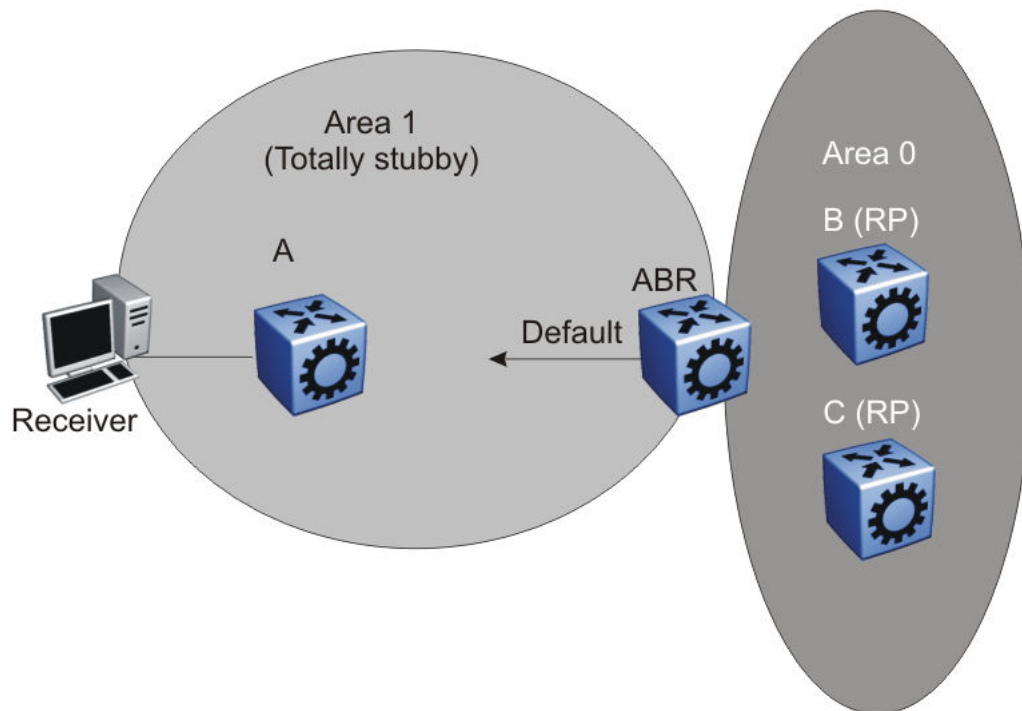
Static RP and RP Redundancy

You can provide RP redundancy through static RPs. To ensure consistency of RP selection, implement the same static RP configuration on all PIM-SM routers in the network. In a mixed vendor network, ensure that the same RP selection criteria is used among all routers. For example, to select the active RP for each group address, the switch uses a hash algorithm defined in the PIM-SMv2 standard. If a router from another vendor selects the active RP based on the lowest IP address, then the inconsistency prevents stream delivery to certain routers in the network.

If a group address-to-RP discrepancy occurs among PIM-SM routers, network outages occur. Routers that are unaware of the true RP cannot join the shared tree and cannot receive the multicast stream.

Failure detection of the active RP is determined by the unicast routing table. As long as the RP is considered reachable from a unicast routing perspective, the local router assumes that the RP is fully functional and attempts to join the shared tree of that RP.

The following figure shows a hierarchical OSPF network where a receiver is in a totally stubby area. If RP B fails, PIM-SM router A does not switch over to RP C because the injected default route in the unicast routing table indicates that RP B is still reachable.




Because failover is determined by unicast routing behavior, carefully consider the unicast routing design, as well as the IP address you select for the RP. Static RP failover performance depends on the convergence time of the unicast routing protocol. For quick convergence, ensure that you use a link state protocol, such as OSPF. For example, if you use RIP as the routing protocol, an RP failure can take minutes to detect. Depending on the application, this situation can be unacceptable.

Static RP failover time does not affect routers that have already switched over to the SPT; failover time only affects newly-joining routers.

Unsupported Static RP Configurations

If you use static RP, you disable dynamic RP learning. The following figure shows an unsupported configuration for static RP. In this example because of inter-operation between static RP and dynamic RP, no RP exists at switch 2. However, (S,G) creation and deletion occurs every 210 seconds at switch 16.



Switches 10, 15, and 16 use static RP, whereas switch 2 uses dynamic RP. The source is at switch 10, and the receivers are switches 15 and 16. The RP is at switch 15 locally. The receiver on switch 16 cannot receive packets because its SPT goes through switch 2.

Switch 2 is in a dynamic RP domain, so it cannot learn about the RP on switch 15. However, (S, G) records are created and deleted on switch 16 every 210 seconds.

Rendezvous Point Router Considerations

You can place an RP on a switch when VLANs extend over several switches. However, when you use PIM-SM, ensure that you do not span VLANs on more than two switches.

Use static group-range-to-RP mappings in an SMLT topology as opposed to RP set learning using the Bootstrap Router (BSR) mechanism. Static RP allows for faster convergence in box failure and reset, whereas there are inherent delays in the BSR mechanism as follows:


- When a router comes back up after a failover or reset, to accept and propagate (*,g) join requests from surrounding routers (either PIM join messages or local IGMP membership reports) to the RP, a PIM router must determine the address of the RP for each group for which they desire (*,g) state. The PIM router must have the unicast route to the RP address. The route to the RP address is learned by using a unicast routing protocol such as OSPF, and the RP address is either statically configured or dynamically learned using the BSR mechanism.
- When a box comes up after a reset, if the RP is not statically configured, it must wait for the BSR to select the RP from candidate RP routers, and then propagate the RP set hop-by-hop to all PIM routers. This must be done before a join message can be processed. If the PIM router receives a join message before it learns the RP set, it drops the join message, and the router waits for another join or prune message to arrive before it creates the multicast route and propagates the join message to the RP. The default Join/Prune timer is 60 seconds, and because of this and the delays inherent in BSR RP-set learning, significant multicast traffic interruptions can

occur. If the RP is statically configured, the only delay is in the unicast routing table convergence and the arrival of the Join/Prune messages from surrounding boxes.

Layer 3 Multicast Extended VLANS

Avoid using a Layer 3 multicast extended VLAN topology without SMLT.

Do not connect non-SMLT PIM routers in a linear fashion on the same VLAN. This topology is called an extended VLAN. Unlike a shared VLAN topology where all routers on the same VLAN are physically one hop away from each other, a VLAN router at one end of the extended VLAN has one or more routers in between it and the router at the far end of the extended VLAN. The following figure shows an extended VLAN.




In the preceding figure, all routers use PIM-SSM. The source connects to Router A on VLAN 10. All routers and receiver hosts connect on the same extended VLAN, VLAN 20. All routers have a receiver in VLAN 20. Router D is the PIM DR for VLAN 20 and the source host is not on VLAN 20. PIM-SSM does not require a Rendezvous Point (RP).

In this topology, each router receives an IGMP membership report from its local receiver host, and then sends a PIM SG join message towards the source on VLAN 20. VLAN flooding propagates the PIM SG join message through to Router A, the PIM DR for the source VLAN 10. Each router from Router D to Router A records a PIM join on the port on which the join message was received, and then sends out its own join message toward the source. Data then flows from the source to the receiver, as long as a join exists on those ports.

Because all routers are in the same VLAN 20, they receive joins from one another due to flooding in the VLAN. For example, Router D receives join messages from Router C on its port 'b', and Router C receives join messages from Router B on its port toward Router B. In accordance with the PIM protocol rules, suppression causes Router D to stop sending a join towards the source because it receives a join for the same group and same RP on the port (port b) of the upstream neighbor (the router towards the source). Router D does not need to send a redundant join on the same VLAN. Router D stops sending a join, and the join that is recorded on port c of Router C eventually times out and is removed from the egress list of the (s,g) multicast route entry on Router C. This removal causes Router C to stop forwarding multicast traffic to Router D, and to the receiver (R1).

The purpose of join suppression is to suppress joins on a shared VLAN, such that if all routers on the shared VLAN want to receive data from the same RP and group, then only one of them needs to send the join on the VLAN. One join is enough to pull the data from the source router to the shared VLAN for all routers to receive. The other routers can suppress sending their own joins when they see such a join on the port toward the upstream router. In this way, less protocol message congestion exists in the shared VLAN. In the following figure, Router D sends the initial join message, which is seen by Router B and Router C. Router B and Router C suppress their own join messages. Router A (the PIM DR for the source VLAN 10) sends the data to VLAN 20, which is received by Routers B, C, and D due to the shared (non-extended) VLAN topology, and traffic is forwarded to all receiver hosts.



The extended VLAN topology looks exactly like the non-extended shared VLAN topology to the router, which cannot distinguish between the two.

In the current release, you cannot disable join suppression on a router. This enhancement will be added in a future release. Until this enhancement is included, you can perform the following actions:

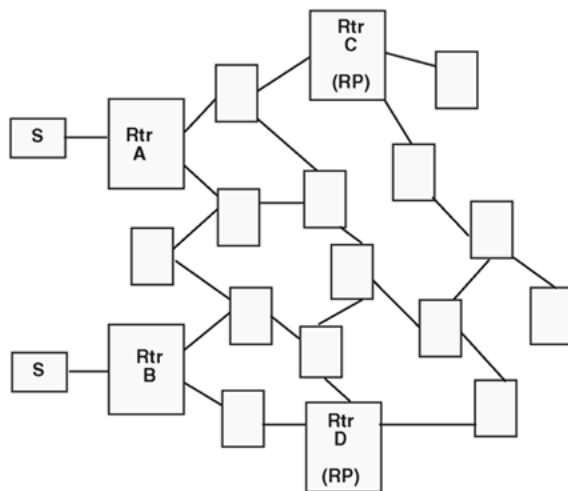
1. Avoid this type of extended VLAN topology, and instead use Layer 3 routing between the routers. Do not extend VLAN 20 throughout, but rather, create a different VLAN between each router.
2. Configure the PIM DR for VLAN 20 to be the router closer to the source (Router B) so that any join received on the VLAN 20 DR (Router B) will be recorded as an IGMP local leaf on VLAN 20 as opposed to a PIM join, which does not time out until the receiver host stops sending IGMP membership reports.

PIM-SM Design and the BSR Hash Algorithm

To optimize the flow of traffic down the shared trees in a network that uses a BSR to dynamically advertise candidate RPs, consider the hash function. The BSR uses the hash function to assign multicast group addresses to each C-RP.

The BSR distributes the hash mask used to compute the RP assignment. For example, if two RPs are candidates for the range 239.0.0.0 through 239.0.0.127, and the hash mask is 255.255.255.252, that range of addresses is divided into groups of four consecutive addresses and assigned to one or the other C-RP.

The following figure illustrates a suboptimal design where Router A sends traffic to a group address assigned to RP D. Router B sends traffic assigned to RP C. RP C and RP D serve as backups for each other for those group addresses. To distribute traffic, it is desirable that traffic from Router A use RP C and that traffic from Router B use RP D.



While still providing redundancy in the case of an RP failure, you can ensure that the optimal shared tree is used by using the following methods.

1. Use the hash algorithm to proactively plan the group-address-to-RP assignment.

Use this information to select the multicast group address for each multicast sender on the network and to ensure optimal traffic flows. This method is helpful for modeling more complex redundancy and failure scenarios, where each group address has three or more C-RPs.

2. Allow the hash algorithm to assign the blocks of addresses on the network, and then view the results using the command **show ip pim active-rp**.

Use the command output to assign multicast group addresses to senders that are located near the indicated RP. The limitation to this approach is that while you can easily determine the current RP for a group address, the backup RP is not shown. If more than one backup for a group address exists, the secondary RP is not obvious. In this case, use the hash algorithm to reveal which of the remaining C-RPs take over for a particular group address in the event of primary RP failure.

The hash algorithm works as follows:

1. For each C-RP router with matching group address ranges, a hash value is calculated according to the formula:

$$\text{Hash value [G, M, C(i)]} = \{ [1\ 103\ 515\ 245 * [(1\ 103\ 515\ 245 * (G \& M) + 12\ 345) \text{ XOR } C(i)] + 12\ 345 \} \bmod 2^{31}$$

The hash value is a function of the group address (G), the hash mask (M), and the IP address of the C-RP C(i). The expression (G&M) guarantees that blocks of group addresses hash to the same value for each C-RP, and that the size of the block is determined by the hash mask.

For example, if the hash mask is 255.255.255.248, the group addresses 239.0.0.0 through 239.0.0.7 yield the same hash value for a given C-RP. Thus, the block of eight addresses are assigned to the same RP.

2. The C-RP with the highest resulting hash value is chosen as the RP for the group. In the event of a tie, the C-RP with the highest IP address is chosen.

This algorithm runs independently on all PIM-SM routers so that every router has a consistent view of the group-to-RP mappings.

Candidate RP Considerations

The C-RP priority parameter determines an active RP for a group. The hash values for different RPs are only compared for RPs with the highest priority. Among the RPs with the highest priority value and the same hash value, the C-RP with the highest RP IP address is chosen as the active RP.

You cannot configure the C-RP priority. Each RP has a default C-RP priority value of 0, and the algorithm uses the RP if the group address maps to the grp-prefix that you configure for that RP. If a different router in the network has a C-RP priority value greater than 0, the switch uses this part of the algorithm in the RP election process.

Currently, you cannot configure the hash mask used in the hash algorithm. Unless you configure a different PIM BSR in the network with a nondefault hash mask value, the default hash mask of 255.255.255.252 is used. Static RP configurations do not use the BSR hash mask; they use the default hash mask.

For example:

RP1 = 128.10.0.54 and RP2 = 128.10.0.56. The group prefix for both RPs is 238.0.0.0/255.0.0.0. Hash mask = 255.255.255.252.


The hash function assigns the groups to RPs in the following manner:

The group range 238.1.1.40 to 238.1.1.51 (12 consecutive groups) maps to 128.10.0.56. The group range 238.1.1.52 to 238.1.1.55 (4 consecutive groups) maps to 128.10.0.54. The group range 238.1.1.56 to 238.1.1.63 (8 consecutive groups) maps to 128.10.0.56.

PIM-SM Receivers and VLANs

Some designs cause unnecessary traffic flow on links in a PIM-SM domain. In these cases, traffic is not duplicated to the receivers, but wastes bandwidth.

The following figure shows such a situation. Switch B is the DR between switches A and B. Switch C is the RP. A receiver R is on the VLAN (V1) that connects switches A and B. A source sends multicast data to the receiver.



IGMP reports that the messages that the receiver sends are forwarded to the DR, and both A and B create (*,G) records. Switch A receives duplicate data through the path

from C to A, and through the second path from C to B to A. Switch A discards the data on the second path (assuming the upstream source is A to C).


To avoid this waste of resources, do not place receivers on V1. This configuration guarantees that no traffic flows between B and A for receivers attached to A. In this case, the existence of the receivers is only learned through PIM join messages to the RP [for (*,G)] and of the source through SPT joins.

PIM Network with Non-PIM Interfaces

For proper multicast traffic flow in a PIM-SM domain, as a general rule, enable PIM-SM on all interfaces in the network (even if paths exist between all PIM interfaces). Enable PIM on all interfaces because PIM-SM relies on the unicast routing table to determine the path to the RP, BSR, and multicast sources. Ensure that all routers on these paths have PIM-SM enabled interfaces.

The following figure provides an example of this situation. If A is the RP, then initially the receiver receives data from the shared tree path (that is, through switch A).

If the shortest path from C to the source is through switch B, and the interface between C and B does not have PIM-SM enabled, then C cannot switch to the SPT. C discards data that comes through the shared path tree (that is, through A). The simple workaround is to enable PIM on VLAN1 between C and B.



Source Filtering

The system can report interest in receiving packets from only a specific source address (INCLUDE), from all but specific source addresses (EXCLUDE), or sent to specific multicast addresses. IGMPv3 interacts with PIM-SM, PIM-SSM, and snooping to provide source filtering.

Protocol Independent Multicast-Source Specific Multicast Guidelines

PIM-Source Specific Multicast (SSM) is a one-to-many model that uses a subset of the PIM-SM features. In this model, members of an SSM group can only receive multicast traffic from a specific source or sources, which is more efficient and puts less load on multicast routing devices.

IGMPv3 supports PIM-SSM by enabling a host to selectively request traffic from individual sources within a multicast group. The system can report interest in receiving packets from only specific source addresses (INCLUDE). IGMPv3 interacts with PIM-SM, PIM-SSM, and snooping to provide source filtering.

PIM-SSM Design Considerations

Use the following information when you design an SSM network:

- If you configure SSM, it affects SSM groups only. The switch handles other groups in sparse mode (SM) if a valid RP exists on the network.
- You can configure PIM-SSM only on switches at the edge of the network. Core switches use PIM-SM if they do not have receivers for SSM groups.
- For networks where group addresses are already in use, you can change the SSM range to match the groups.
- One switch has a single SSM range.
- You can have different SSM ranges on different switches.

Configure the core switches that relay multicast traffic so that they cover all of these groups in their SSM range, or use PIM-SM.

- One group in the SSM range can have multiple sources for a given SSM group.

Multicast for Multimedia

The switch provides a flexible and scalable multicast implementation for multimedia applications. Several features are dedicated to multimedia applications and in particular to television distribution.

Join and Leave Performance

For TV applications, you can attach several TVs directly to the switch, or through an IGMP-capable Ethernet switch. Base this implementation on IGMP; the set-top boxes use IGMP reports to join a TV channel and IGMP leaves to exit the channel. After a viewer changes channels, the switch issues an IGMPv2 leave for the old channel (multicast group), and sends a membership report for the new channel. If viewers change channels continuously, the number of joins and leaves can become large, particularly if many viewers attach to the switch.

The switch supports more than a thousand joins and leaves per second, which is well adapted to TV applications.



Important

For IGMPv3, ensure a join rate of 1000 per second or less. This ensures the timely processing of join requests.

If you use the IGMP proxy functionality at the receiver edge, you reduce the number of IGMP reports received by switch. This provides better overall performance and scalability.

Fast Leave

IGMP Fast Leave supports two modes of operation: single-user mode and multiple-user mode.

In single-user mode, if more than one member of a group is on the port and one of the group members leaves the group, everyone stops receiving traffic for this group.

Single-user mode does not send a group-specific query before the effective leave takes place.

Multiple-user mode allows several users on the same port or VLAN. If one user leaves the group and other receivers exist for the same stream, the stream continues. The switch tracks the number of receivers that join a given group. For multiple-user mode to operate properly, do not suppress reports. This ensures that the switch properly tracks the correct number of receivers on an interface.

The Fast Leave feature is particularly useful in IGMP-based TV distribution where only one receiver of a TV channel connects to a port. If a viewer changes channels quickly, you create considerable bandwidth savings if you use Fast Leave.

You can implement Fast Leave on a VLAN and port combination; a port that belongs to two different VLANs can have Fast Leave enabled on one VLAN (but not on the other). Thus, with the Fast Leave feature enabled, you can connect several devices on different VLANs to the same port. This strategy does not affect traffic after one device leaves a group to which another device subscribes. For example, you can use this feature when two TVs connect to a port through two set-top boxes, even if you use the single-user mode.

To use Fast Leave, you must first enable explicit host tracking. IGMP uses explicit host tracking to track all source and group members. Explicit host tracking is disabled by default. For configuration information, see [Configure Fast Leave Mode](#) on page 1881.

Last Member Query Interval Tuning

If an IGMPv2 host leaves a group, it notifies the router by using a leave message. Because of the IGMPv2 report suppression mechanism, the router cannot access information of other hosts that require the stream. Thus, the router broadcasts a group-specific query message with a maximum response time equal to the last member query interval (LMQI).

Because this timer affects the latency between the time that the last member leaves and the time the stream actually stops, you must properly tune this parameter. This timer can especially affect TV delivery or other large-scale, high-bandwidth multimedia applications. For instance, if you assign a value that is too low, this can lead to a storm of membership reports if a large number of hosts are subscribed. Similarly, assigning a value that is too high can cause unwanted high-bandwidth stream propagation across the network if users change channels rapidly. Leave latency also depends on the robustness value, so a value of 2 equates to a leave latency of twice the LMQI.

Determine the proper LMQI value for your particular network through testing. If a very large number of users connect to a port, assigning a value of 3 can lead to a storm of report messages after a group-specific query is sent. Conversely, if streams frequently start and stop in short intervals, as in a TV delivery network, assigning a value of 10 can lead to frequent congestion in the core network.

Another performance-affecting factor that you need to be aware of is the error rate of the physical medium. For links that have high packet loss, you can find it necessary to adjust the robustness variable to a higher value to compensate for the possible loss of IGMP queries and reports.

In such cases, leave latency is adversely affected as numerous group-specific queries are unanswered before the stream is pruned. The number of unanswered queries is equal to the robustness variable (default 2). The assignment of a lower LMQI can counterbalance this effect. However, if you configure the LMQI too low, it can actually exacerbate the problem by inducing storms of reports on the network. LMQI values of 3 and 10, with a robustness value of 2, translate to leave latencies of 6/10 of a second and 2 seconds, respectively.

When you choose an LMQI, consider all of these factors to determine the best configuration for the given application and network. Test that value to ensure that it provides the best performance.

**Important**

In networks that have only one user connected to each port, use the Fast Leave feature instead of LMQI, because no wait is required before the stream stops. Similarly, the robustness variable does not affect the Fast Leave feature, which is an additional benefit for links with high loss.

Layer 3 Switch Clustering and Multicast SMLT

Switch clustering is the logical aggregation of two nodes to form one logical entity known as the switch cluster. The two peer nodes in a switch cluster connect using a virtual interswitch trunk (vIST). The vIST exchanges forwarding and routing information between the two peer nodes in the cluster. This section provides guidelines for switch clusters that use multicast and Split Multilink Trunking (SMLT).

General Guidelines

The following list identifies general guidelines to follow if you use multicast and switch clustering:

- Enable Protocol Independent Multicast - Sparse Mode (PIM-SM) on the vIST VLAN for fast recovery of multicast. A unicast routing protocol is not required.
- Enable Internet Group Management Protocol (IGMP) snooping and proxy on the edge switches.

The following figure shows multicast behavior in an SMLT environment. The configuration in the following figure provides fast failover if the switch or rendezvous point (RP) fails.




Figure 133: Multicast behavior in SMLT environment

In Multicast behavior in SMLT environment the following actions occur:

1. The multicast server sends multicast data towards the source designated router (DR).
2. The source DR sends register messages with encapsulated multicast data towards the RP.
3. After the client sends IGMP membership reports towards the multicast router, the router creates a (*,G) entry.
4. The RP sends join messages towards the source DR on the reverse path.
5. After the source DR receives the join messages, it sends native multicast traffic.
6. After SW_B or SW_D receives multicast traffic from upstream, it forwards the traffic on the vIST as well as on the SMLT link. Other aggregation switches drop multicast

traffic received over the vIST at egress. This action provides fast failover for multicast traffic. Both SW_D and SW_E (Aggregation switches) have similar (S,G) records.

7. In case of SW_D or RP failure, SW_B changes only the next-hop interface towards SW_E. Because the circuitless IP (CLIP) RP address is the same, SW_B does not flush (S,G) entries and achieves fast failover.

Multicast Triangle Topology

A triangle design is an SMLT configuration that connects edge switches or SMLT clients to two aggregation switches. Connect the aggregation switches together with a vIST that carries all the SMLT trunks configured on the switches.

The switch supports the following triangle configurations:

- a configuration with Layer 3 PIM-SM routing on both the edge and aggregation switches
- a configuration with Layer 2 snooping on the client switches and Layer 3 routing with PIM-SM on the aggregation switches

To avoid using an external query device to provide correct handling and routing of multicast traffic to the rest of the network, use the triangle design with IGMP Snoop at the client switches. Use multicast routing at the aggregation switches as shown in the following figure.




Figure 134: Multicast routing using PIM-SM

Client switches run IGMP Snoop or PIM-SM, and the aggregation switches run PIM-SM. This design is simple and, for the rest of the network, PIM-SM performs IP multicast routing. The aggregation switches are the query devices for IGMP, so an external query device is not required to activate IGMP membership. These switches also act as redundant switches for IP multicast.

Multicast data flows through the vIST link when receivers are learned on the client switch and senders are located on the aggregation switches, or when sourced data

comes through the aggregation switches. This data is destined for potential receivers attached to the other side of the vIST. The data does not reach the client switches through the two aggregation switches because only the originating switch forwards the data to the client switch receivers.



Note

Always place multicast receivers and senders on the core switches on VLANs different from those that span the vIST.

The following figure shows a switch clustering configuration with a single switch cluster core and dual-connected edge devices. This topology represents different VLANs spanning from each edge device and those VLANs routed at the switch cluster core. You can configure multiple VLANs on the edge devices, 802.1Q tagged to the switch cluster core.




Figure 135: Multicast SMLT triangle

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either Virtual Router Redundancy Protocol (VRRP) BackupMaster or Routed SMLT (RSMLT) Layer 2 Edge on the switch cluster core.

Square and Full-Mesh Topology Multicast Guidelines

A square design connects a pair of aggregation switches to another pair of aggregation switches. A square design becomes a full-mesh design if the aggregation switches are connected in a full-mesh. The switch supports Layer 3 IP multicast (PIM-SM only) over a full-mesh SMLT or RSMLT configuration.

In a square design, configure all switches with PIM-SM. Place the bootstrap router (BSR) and RP in one of the four core switches; and place the RP closest to the source. If using PIM-SM over a square or full-mesh configuration, enable the **multicast smlt-square** flag.

The following three figures show switch clustering configurations with two-switch cluster cores and dual-connected edge devices.

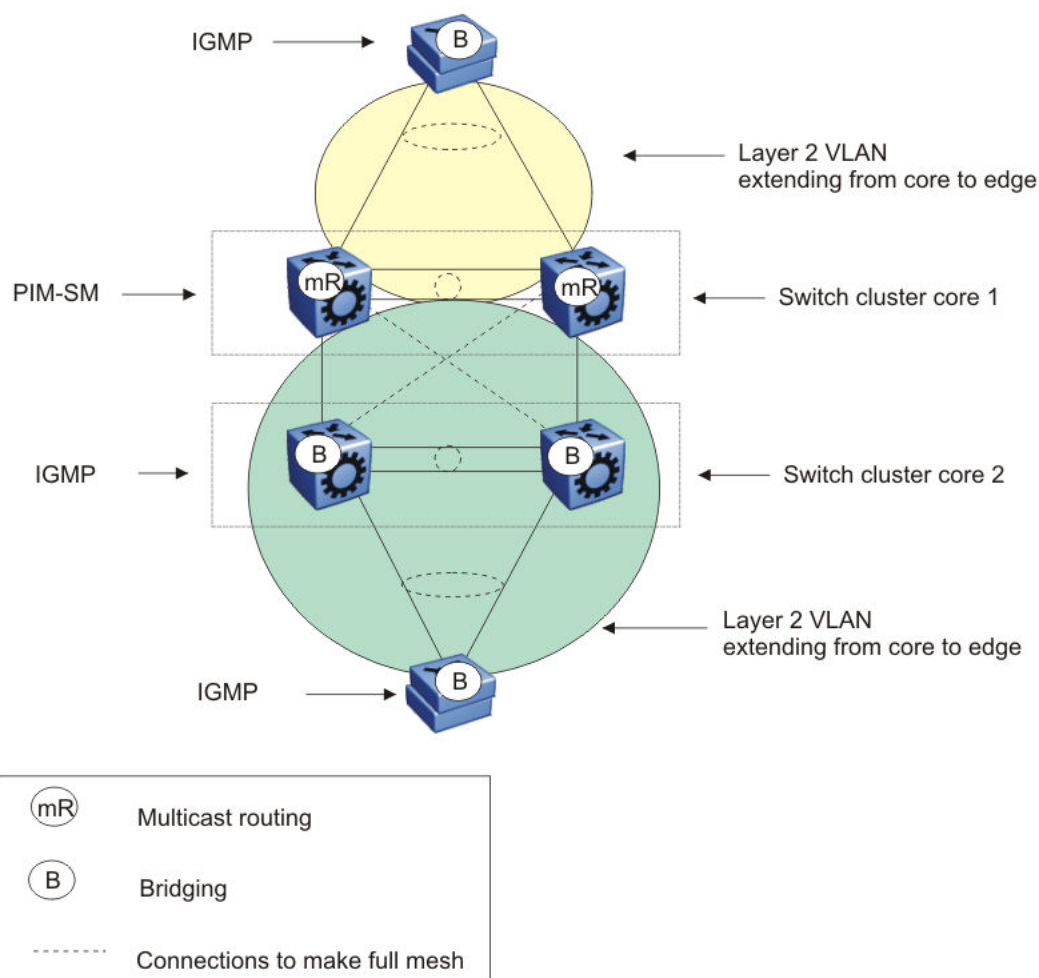


Figure 136: Multicast SMLT square 1

In the preceding figure, only one of the switch cluster cores performs Layer 3 multicast routing while the other is strictly Layer 2. Configure multiple VLANs on the edge devices, 802.1Q tagged to the switch cluster cores.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster core.




Figure 137: Multicast SMLT square 2

In the preceding figure, both of the switch cluster cores performs Layer 3 multicast routing, while the edge devices are Layer 2 IGMP.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster cores. Do not enable VRRP on the RSMLT VLAN between switch cluster cores.




Figure 138: Multicast SMLT square 3

In the preceding figure, both of the switch cluster cores and the edge devices perform Layer 3 multicast routing.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster cores. Do not enable VRRP on the RSMLT VLAN between switch cluster cores.

SMLT and Multicast Traffic Issues

If PIM-SM or other multicast protocols are used in an SMLT environment, enable the protocol on the vIST. Routing protocols in general are not run over an vIST but multicast routing protocols are an exception. When using PIM-SM and a unicast routing protocol, ensure the unicast route to the BSR and RP has PIM-SM active and enabled. If multiple OSPF paths exist and PIM-SM is not active on each pair, the BSR is learned on a path that does not have PIM-SM active. The following figure demonstrates this issue.




Figure 139: Unicast route example

The network configuration in the preceding figure is as follows:

- 5510A is on VLAN 101.
- 5510B is on VLAN 102.
- Switch B is the BSR.
- Switch A and Switch B have OSPF enabled.
- PIM is enabled and active on VLAN 101.
- PIM is either disabled or passive on VLAN 102.

In this example, the unicast route table on Switch A learns the BSR on Switch B through VLAN 102 using OSPF. The BSR is either not learned or does not provide the RP to Switch A.

Protocol Independent Multicast over IPv6

Table 152: PIM over IPv6 product support

Feature	Product	Release introduced
PIM over IPv6	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0



Note

PIM is supported in Global Routing Table (GRT) only.

Several multicast protocols are used to enable IP multicast.

Hosts use the Internet Group Management Protocol (IGMP) for IPv4 and Multicast Listener Discovery (MLD v1/v2) for IPv6 to report multicast group memberships of directly attached multicast listeners to neighboring multicast routers. MLD is the direct IPv6 replacement for the IGMP protocol used in IPv4.

Routers use Protocol Independent Multicast-Sparse Mode (PIM-SM) and PIM source Specific Mode (SSM) to exchange multicast routing information. The PIM-SM protocol is the multicast routing protocol that uses the underlying unicast routing information base to build unidirectional shared trees to group members rooted at the RP per group, and creates shortest-path trees (SPT) per source. The router forwards multicast packets along these trees. PIM-SSM does not require RP and only supports SPT.

PIM over IPv6 uses the IPv6 unicast routing table for reverse path information about source and RP.



Note

IPv4 and IPv6 multicast streams cannot interact. To configure an end-to-end PIM IPv6 network, all nodes from sender to receiver must support PIM IPv6.

PIM-SM over IPv6 Features

The following are features of PIM-SM over IPv6:

- Compliant with RFC 4601
- Multicast networks built by PIM IPv4 and PIM IPv6 do not overlap
- IPv4 receiver hosts cannot receive data from IPv6 source hosts and vice versa

- IPv4 and IPv6 multicast protocols can be enabled at the same time on the same VLAN
- PIM IPv4 and PIM IPv6 can be configured on the same VLAN
- PIM IPv4 and PIM IPv6 must be configured separately
- Supports sparse and ssm modes

Operational Note for PIM-SM over IPv6

The following are operational considerations when deploying PIM-SM over IPv6:

- You can only configure PIM-SM if you configure the `spbm_config_mode` boot flag to false.
- The following HELLO messages options are not supported:
 - GENid
 - DR priority
 - LAN-PRUNE delay
 - T-bit
- IPv6 multicast is not supported over SPBM
- IPv6 multicast routing is not virtualized, it is supported only on GRT
- IPv6 multicast configuration on SMLT VLAN is not supported. vIST peers cannot form PIM-SM over IPv6 neighbor adjacencies. Senders and receivers on the vIST peers (SMLT and non-SMLT) cannot communicate. MLT and LACP is supported.
- The switch does not support the following features:
 - Static entries
 - Bootstrap message (BSR)
 - Anycast RP
 - Virtual PIM neighbors
 - Fast join prune
 - Software forwarding
 - Passive PIM interfaces
 - IP mroute stream limit
 - Bidirectional PIM
 - Multicast Border Router (PMBR)
 - VRF support for PIM (GRT only)
 - IGMP and PIM mtrace capability

IPv6 Interface Multiple Addresses

IPv6 interfaces can have multiple addresses associated with them. A router running PIM for IPv6 has a network unique domain-wide reachable IPv6 VLAN address used for multiple hop messages. A link local address is associated with the VLAN. The link local address is a non-routable unicast IPv6 address used as source address (primary interface address) for transmitting different types of PIM messages.

IP Multicast Configuration and DvR

Configuration of IPv4 multicast is supported only on the Controller nodes of a DvR domain. You cannot configure IP multicast on the DvR Leaf nodes. The following sections detail IP multicast configuration support on DvR enabled nodes (Controllers or Leaf nodes).

For more information on DvR, see [Distributed Virtual Routing Fundamentals](#) on page 919.

Multicast Configuration Pushed from DvR Controllers to DvR Leaf Nodes

When you perform the following multicast configuration on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- IP multicast over Fabric Connect
- IGMP Layer 2 Querier parameters, such as the IGMP Layer 2 Querier version, query interval, query maximum response time, robustness value, last member query interval and compatibility mode
- Enabling and clearing of multicast route statistics

Multicast Configuration not Supported on DvR Enabled Layer 2 VSNs

- IGMP Snooping on DvR enabled Layer 2 VSNs
- SPB-PIM Gateway

For more information on SPB-PIM Gateway, see [SPB-PIM Gateway configuration](#) on page 3557.

IP Multicast Basic Configuration Using CLI

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts subscribe to multicast services using a host membership protocol. The Internet Group Management Protocol (IGMP) is an example of an IPv4 host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast–Sparse Mode (PIM–SM) and Protocol Independent Multicast–Source Specific Multicast (PIM–SSM).

Configure IP Multicast in SMLT Topologies

This procedure shows how to configure PIM and IGMP Snooping in an SMLT environment. The configuration steps show how to enable multicast, and then configure the usual PIM and IGMP Snooping related VLANs and global attributes. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

Before You Begin

SPBM must not be enabled on the vIST peers or any router that participates in the PIM network.

About This Task

The switch supports configurable VLANs in the range of 1 to 4059. VLAN 0 is invalid. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. VLAN IDs on the switch range from 2 to 4094 but, by default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the *vrf-scaling* and *spbm-config-mode* boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Disable the boot flag:

```
no boot config flags spbm-config-mode
```

The system responds with these messages:

```
Warning: Please save the configuration and reboot the switch for this  
to take effect.
```

```
Warning: Please carefully save your configuration file before  
rebooting the switch. Saving configuration file when spbm-config-  
mode is changed to disable, removes SPBM configurations from the  
configuration file.
```

3. Save the configuration and, then reboot the switch.



Important

Any change to the **spbm-config-mode** boot flag requires a reboot for the change to take effect.

4. Create the vIST VLAN:

```
vlan create <2-4059> type port-mstprstp <0-63>
```



```
interface vlan <1-4059>
```



```
ip address <A.B.C.D/X>
```
5. Configure the vIST peer address and VLAN:

```
virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
```

6. Configure the SMLT MLT:

```
mlt <1-512> enable
```

```
mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]
[slot/all][all]}
```

```
interface mlt <1-512>
```

```
smlt
```

7. Configure the vIST MLT:

```
mlt <1-512> enable
```

```
mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]
[slot/all][all]}
```

```
mlt <1-512> encapsulation dot1q
```

```
interface mlt <1-512>
```

```
virtual-ist enable
```

**Note**

The **virtual-ist enable** command enables Simplified vIST and is only available when the **spbm-config-mode** boot flag is disabled.

8. Create a customer VLAN and assign the SMLT MLT ID:

```
vlan create <2-4059>
```

```
vlan mlt <1-4059> <1-512>
```

```
interface vlan <1-4059>
```

```
ip address <A.B.C.D/X>
```

9. Configure PIM or IGMP Snooping on the SMLT VLAN:

```
interface vlan <1-4059>
```

```
ip pim enable OR ip igmp snooping
```

10. Configure PIM on the vIST VLAN:

```
interface vlan <1-4059>
```

```
ip pim enable
```

11. Enable PIM globally:

```
ip pim enable
```

**Note**

You can also configure other global PIM attributes such as **ip pim join-prune-interval**.

Example

```
enable
configure terminal
no boot config flags spbm-config-mode
```

Save the configuration and reboot the switch.

```
virtual-ist peer-ip 198.51.100.0 vlan 50

mlt 3 enable
mlt 3 member 1/35,1/36
interface mlt 3
smlt
exit
mlt 5 enable
mlt 5 member 2/1,2/2
mlt 5 encapsulation dot1q
interface mlt 5
virtual-ist enable
exit
vlan create 50 type port-mstprstp 0
interface vlan 50
ip address 198.51.100.0 255.255.255.0 1
exit
vlan create 100
vlan mlt 100 3
interface vlan 100
ip address 192.0.2.0 255.255.255.0 2
exit
interface vlan 100
ip pim enable (or ip igmp snooping)
exit
interface vlan 50
ip pim enable
exit
ip pim enable
```

Configure PIM-SM Globally

Configure PIM-SM to enable or disable PIM-SM globally on the switch and change default global parameters.

Before You Begin



Note

Before you can enable the PIM Infinite Threshold Policy feature, you must first disable the following:

- PIM-SM
- PIM-SSM
- Simplified vIST

About This Task

PIM-SM is the default mode so you do not need to configure the PIM mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable PIM-SM:

```
ip pim enable
```
3. Configure the time between bootstrap messages:

```
ip pim bootstrap-period <5-32757>
```
4. Configure the timeout to discard data:

```
ip pim disc-data-timeout <5-65535>
```
5. Enable the fast join prune interval:

```
ip pim fast-joinprune
```
6. Configure the forward cache timeout:

```
ip pim fwd-cache-timeout <10-86400>
```
7. Configure the interval for join and prune messages:

```
ip pim join-prune-interval <1-18724>
```
8. Specify how long to suppress register messages:

```
ip pim register-suppression-timeout <6-65535>
```
9. Specify how often the candidate-rendezvous point (C-RP) sends advertisements:

```
ip pim rp-c-adv-timeout <5-26214>
```
10. Configure PIM Infinite Threshold Policy:
 - a. Disable PIM-SM.

```
no pim enable
```
 - b. Enable PIM Infinite Threshold Policy:

```
ip pim spt-infinite-threshold
```
 - c. Enable PIM-SM:

```
ip pim enable
```
11. Configure the polling interval for the routing table manager (RTM):

```
ip pim unicast-route-change-timeout <2-65535>
```
12. Verify the configuration changes:

```
show ip pim
```

Example

Verify the configuration changes:

```
Switch:1>show ip pim
```

```
=====
                        Pim General Group - GlobalRouter
=====
PimStat           : disabled
Mode              : sparse
StaticRP          : disabled
FastJoinPrune     : disabled
SptInfiniteThreshold : enabled
```

```

BootstrapPeriod      : 60
CRPAdvTimeout        : 60
DiscDataTimeout      : 60
FwdCacheTimeout      : 210
RegSupprTimeout      : 60
UniRouteChangeTimeout : 5
JoinPruneInt         : 60

```

Variable Definitions

The following table describes the variables for the **ip pim** command.

Variable	Value
disc-data-timeout <5-65535>	Specifies the duration in seconds to discard data until the switch receives the join message from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The default value is 60.
bootstrap-period	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. The range is from 5–32757 and the default is 60 seconds.
enable	Enables PIM globally on the switch. The default is disabled.
fast-joinprune	Enables or disables the PIM fast join prune feature.
fwd-cache-timeout <10-86400>	Specifies the forward cache timeout value. The default value is 120.
join-prune-interval <1-18724>	Specifies the duration in seconds before the PIM router sends out the next join or prune message to its upstream neighbors. The default value is 60.
mode <sparse> <ssm>	Configures PIM mode on the switch. The default value is sparse.
register-suppression-timeout <10-65535>	Specifies the duration in seconds the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. The default value is 60.

Variable	Value
rp-c-adv-timeout	Specifies how often (in seconds) a router configured as a candidate rendezvous point router (C-RP) sends advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected bootstrap router (BSR). The range is from 5–26214 and the default is 60 seconds.
spt-infinite-threshold	Enables PIM Infinite Threshold Policy for IPv4, so that multicast traffic follows the shared tree path through a Rendezvous Point (RP) instead of automatically switching over to shortest path tree (SPT). The default is disabled.
static-rp	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism.
unicast-route-change-timeout <2-65535>	Specifies the duration in seconds the switch polls the RTM for unicast routing information updates for PIM. The default value is 5.
virtual-neighbor	Specifies to enter virtual neighbor IP to an interface globally.

Enable IPv6 PIM-SM Globally

About This Task

Use this procedure to enable IPv6 PIM-SM globally. By default, IPv6 PIM-SM is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable IPv6 PIM-SM:

```
ipv6 pim enable
```


Configure Global IPv6 PIM-SM Properties

Before You Begin



Note

Before you can enable the PIM Infinite Threshold Policy feature, you must first disable the following:

- PIM-SM
- PIM-SSM

About This Task

Use this procedure to configure the global IPv6 PIM-SM parameters on the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the timeout to discard data:

```
ipv6 pim disc-data-timeout <5-65535>
```
3. Configure the forward cache timeout:

```
ipv6 pim fwd-cache-timeout <10-86400>
```
4. Configure the interval for join and prune messages:

```
ipv6 pim join-prune-interval <1-18724>
```
5. Specify how long to suppress register messages:

```
ipv6 pim register-suppression-timeout <10-65535>
```
6. Configure PIM Infinite Threshold Policy:

```
ipv6 pim spt-infinite-threshold
```
7. Configure the polling interval for the routing table manager (RTM):

```
ipv6 pim unicast-route-change-timeout <2-65535>
```
8. Configure the PIM mode:

```
ipv6 pim mode <sparse> <:ssm>
```
9. Verify the configuration changes:

```
show ipv6 pim
```

Example

Verify the configuration changes:

```
Switch:1>show ipv6 pim
```

```
=====
                        Pim General Group - GlobalRouter
=====
PimStat           : enabled
Mode              : sparse
StaticRP          : disabled
SptInfiniteThreshold : enabled
FwdCacheTimeout   : 210
```

```

DiscDataTimeout      : 60
RegSupprTimeout      : 60
UniRouteChangeTimeout : 5
JoinPruneInt         : 60

```

Variable Definitions

The following table describes the variables for the **ipv6 pim** command.

Variable	Value
disc-data-timeout <5-65535>	Specifies the duration in seconds to discard data until the switch receives the join message from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The default value is 60.
enable	Enables PIM globally on the switch. The default is disabled.
fwd-cache-timeout <10-86400>	Specifies the forward cache timeout value. The default value is 120.
join-prune-interval <1-18724>	Specifies the duration in seconds before the PIM router sends out the next join or prune message to its upstream neighbors. The default value is 60.
mode <sparse> <ssm>	Configures PIM mode on the switch. The default value is sparse.
register-suppression-timeout <10-65535>	Specifies the duration in seconds the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. The default value is 60.
spt-infinite-threshold	Enables PIM Infinite Threshold Policy for IPv6, so that multicast traffic follows the shared tree path through a Rendezvous Point (RP) instead of automatically switching over to shortest path tree (SPT). The default is disabled.
static-rp	Add new static-rp entries and enable static-rp.
unicast-route-change-timeout <2-65535>	Specifies the duration in seconds the switch polls the RTM for unicast routing information updates for PIM. The default value is 5.

Configure PIM on a VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- You must enable PIM globally before you configure PIM on a VLAN.
- The interface uses a valid IP address.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```

2. Create a PIM interface on a VLAN:

```
ip pim enable
```

This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

```
ip pim join-prune-interval <1-18724>
```

4. Configure the time between hello messages:

```
ip pim hello-interval <0-18724>
```

5. Verify the configuration:

```
show ip pim interface vlan [<1-4059>]
```

Example

Configure the interval for join and prune messages, the time between hello messages, and then verify the configuration.

```
Switch:1(config-if)#ip pim join-prune-interval 60
Switch:1(config-if)#ip pim hello-interval 30
Switch:1>show ip pim interface vlan 10
```

```
=====
                        Vlan Ip Pim
=====
```

VLAN-ID	PIM-ENABLE	MODE	HELLOINT	JPINT	CBSRPREF	INTF TYPE
10	enable	sparse	30	60	-1 (disabled)	active

Configure PIM on a Port

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- You must enable PIM globally before you configure it on an interface.

- The interface uses a valid IP address.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create a PIM interface on a port:

```
ip pim enable
```

This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

```
ip pim join-prune-interval <1-18724>
```

4. Configure the time between hello messages:

```
ip pim hello-interval <0-18724>
```

Example

Configure the interval for join and prune messages and the time between hello messages:

```
Switch(config-if)#ip pim join-prune-interval 60  
Switch(config-if)#ip pim hello-interval 30
```

Configure IPv6 PIM on a Port or VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- Enable IPv6 interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create a PIM interface on a port or VLAN:

```
ipv6 pim enable
```

This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

```
ipv6 pim join-prune-interval <1-18724>
```

4. Configure the time between hello messages:

```
ipv6 pim hello-interval <0-18724>
```

Example

```
Switch:1(config-if)#ipv6 pim join-prune-interval 60  
Switch:1(config-if)#ipv6 pim hello-interval 30
```

Variable Definitions

The following table describes the variables for the **ipv6 pim** command.

Variable	Value
<i>hello-interval</i> <0-18724>	Specifies the duration in seconds before the PIM router sends out the next hello message to neighboring switches. The default value is 30 seconds.
<i>join-prune-interval</i> <1-18724>	Specifies the duration in seconds before the PIM router sends out the next join or prune message to its upstream neighbors. The default value is 60 seconds.

Configure SSM Globally

Configure SSM to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Before You Begin

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP, see [RIP configuration using CLI](#) on

page 3156. For more information about OSPF, see [OSPF configuration using CLI](#) on page 2783.

- Enable PIM globally.

About This Task

Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range.

For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure PIM-SSM:

```
ip pim mode ssm
```

Configure IPv6 SSM Globally

Configure IPv6 SSM to optimize IPv6 PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Before You Begin

- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM.

For more information about RIPng, see [RIPng Configuration using CLI](#) on page 3167. For more information about OSPFv3, see [OSPFv3 Configuration using CLI](#) on page 2827.

- Enable IPv6 PIM globally.

About This Task

Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model which requires only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On a SSM-enabled switch, SSM behavior is limited to the SSM group range.

For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure IPv6 PIM-SSM:

```
ipv6 pim mode ssm
```

Configure IGMP on a VLAN

Configure IGMP for each interface to change default multicasting operations.



Note

When you configure the following IGMP parameters on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- **ip igmp version**
- **ip igmp query-interval**
- **ip igmp query-max-response**
- **ip igmp robust-value**
- **ip igmp last-member-query-interval**
- **ip igmp compatibility-mode**

IGMP snooping is not supported on DvR enabled Layer 2 VSNs.

Before You Begin

- For PIM interfaces, you must enable PIM globally and on the VLAN. For snooping interfaces, do not enable PIM.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable IGMP v2-v3 compatibility mode:

```
ip igmp compatibility-mode
```

3. Configure the system to downgrade the version of IGMP:

```
ip igmp dynamic-downgrade-version
```

4. Configure message intervals and response times:

```
ip igmp last-member-query-interval <0-255> [query-interval <1-65535>]  
[query-max-response <0-255>]
```

5. Configure expected packet loss and IGMP version:

```
ip igmp robust-value <2-255> [version <1-3>]
```

6. Add multicast router ports:

```
ip igmp mrouter {slot/port[-slot/port][,...][slot/all][all]}
```

7. Enable proxy-snoop:

```
ip igmp proxy
```

8. Enable router alert:

```
ip igmp router-alert
```

9. Enable snooping:

```
ip igmp snooping
```

10. Enable SSM-snooping:

```
ip igmp ssm-snoop
```

Example

Enter VLAN Interface Configuration Mode for VLAN 10:

```
Switch:1(config)# interface vlan 10
```

Configure the last member query interval to 15 tenths of a second (equal to 1.5 seconds).

```
Switch:1(config-if)# ip igmp last-member-query-interval 15
```

Configure the query interval to 100 seconds.

```
Switch:1(config-if)# ip igmp query-interval 100
```

Configure the query maximum response time to 15 tenths of a second (equal to 1.5 seconds).

```
Switch:1(config-if)# ip igmp query-max-response 15
```

Configure the robustness value to 4 seconds.

```
Switch:1(config-if)# ip igmp robust-value 4
```

Enable proxy snoop for the VLAN.

```
Switch:1(config-if)# ip igmp proxy
```

Enable snoop for the VLAN.

```
Switch:1(config-if)# ip igmp snooping
```

Enable support for SSM on the snoop interface.

```
Switch:1(config-if)# ip igmp ssm-snoop
```

Enable IGMPv3.

```
Switch:1(config-if)# ip igmp version 3
```


Variable Definitions

Use the definitions in the following table to use the **ip igmp** command.

Variable	Value
<code>access-list WORD<1-64> {A.B.C.D/X} <eny-tx deny-rx deny-both allow-only-tx allow-only-rx allow-only-both></code>	<p>Specifies the name of the access list from 1–64 characters.</p> <p>Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.</p> <p>Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic</p>
<code>compatibility-mode</code>	<p>Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command:</p> <pre>default ip igmp compatibility-mode</pre> <p>, or use the no option to disable compatibility mode:</p> <pre>no ip igmp compatibility-mode</pre>
<code>dynamic-downgrade-version</code>	<p>Configures the version of IGMP to handle older query messages if the system downgrades. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:</p> <pre>default ip igmp dynamic-downgrade-version</pre> <p>or use the no option to disable downgrade:</p> <pre>no ip igmp dynamic-downgrade-version</pre>
<code>igmpv3-explicit-host-tracking</code>	Enables explicit host tracking on IGMPv3. The default state is disabled.
<code>immediate-leave</code>	Enables fast leave on a VLAN.

Variable	Value
<i>immediate-leave-members</i> {slot/port[/sub-port] [-slot/ port[/sub-port]] [,...] [slot/ all] [all]}	Configures IGMP fast leave members on a VLAN to specify fast-leave-capable ports. Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>last-member-query-interval</i> <0-255>	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The default is 10 tenths of a second. You should configure this value between 3-10 (equal to 0.3 – 1.0 seconds).
<i>mrdisc</i> [maxadvertinterval <2-180>] [maxinitadvertinterval <2- 180>] [maxinitadvertisements <2-15>] [minadvertinterval <3-180>] [neighdeadinterval <2-180>]	Configures the multicast router discovery options to enable the automatic discovery of multicast capable routers. The default parameter values are: <ul style="list-style-type: none"> • maxadvertinterval: 20 seconds • maxinitadvertinterval: 2 seconds • maxinitadvertisements: 3 • minadvertinterval: 15 seconds • neighdeadinterval: 60 seconds
<i>mrrouter</i> {slot/port[/sub- port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}	Adds multicast router ports. Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>proxy</i>	Activates the proxy-snoop option globally for the VLAN.

Variable	Value
<code>query-interval <1-65535></code>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
<code>query-max-response <0-255></code>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds). Important: You must configure this value lower than the query-interval.
<code>robust-value <2-255></code>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.
<code>router-alert</code>	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. Important: To maximize network performance, configure this parameter according to the version of IGMP currently in use: <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
<code>snoop-querier</code>	Enables the IGMP Layer 2 Querier feature on the VLAN. The default is disabled.
<code>snoop-querier-addr {A.B.C.D}</code>	Specifies the IGMP Layer 2 Querier source IP address.
<code>snooping</code>	Activates the snoop option for the VLAN.
<code>ssm-snoop</code>	Activates support for PIM-SSM on the snoop interface.
<code>static-group {A.B.C.D} {A.B.C.D} {slot/port[/sub-port] [-slot/ port[/sub-port]] [, ...] [slot/ all] [all]} [static blocked]</code>	Configures IGMP static members to add members to a snoop group. {A.B.C.D} {A.B.C.D} indicates the IP address range of the selected multicast group. {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]} adds ports to a static group entry. [static blocked] configures the route to static or blocked.

Variable	Value
<code>stream-limit stream-limit-max-streams <0-65535></code>	Configures multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. The default is 4.
<code>stream-limit-group {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} enable max-streams <0-65535></code>	Configures multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN. The default <code>max-streams</code> value is 4.
<code>version <1-3></code>	Configures the version of IGMP for this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configure IGMP Ports

Configure IGMP for each interface to change default multicasting operations.



Note

When you configure the following IGMP parameters on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- **`ip igmp version`**
- **`ip igmp query-interval`**
- **`ip igmp query-max-response`**
- **`ip igmp robust-value`**
- **`ip igmp last-member-query-interval`**
- **`ip igmp compatibility-mode`**

For more information on DvR, see [Distributed Virtual Routing](#) on page 915.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format `slot/port/sub-port`.

2. Enable IGMP v2-v3 compatibility mode:

```
ip igmp compatibility-mode
```

3. Configure the system to downgrade the version of IGMP:

```
ip igmp dynamic-downgrade-version
```

4. Configure message intervals and response times:

```
ip igmp last-member-query-interval <0-255> [query-interval <1-65535>]  
[query-max-response <0-255>]
```

5. Configure expected packet loss and IGMP version:

```
ip igmp robust-value <2-255> [version <1-3>]
```

6. Configure IGMP for a specific port:

```
ip igmp port {slot/port[-slot/port]}[,...][slot/all][all]}
```

7. Enable router alert:

```
ip igmp router-alert
```

Example

Configure message intervals and response times:

```
Switch(config-if)#ip igmp last-member-query-interval 30 query-interval 60 query-max-  
response 90
```

Configure expected packet loss and IGMP version:

```
Switch(config-if)#ip igmp robust-value 2 version 3
```

Configure IGMP for a specific port:

```
Switch(config-if)#ip igmp port 1/4
```

Enable router alert:

```
Switch(config-if)#ip igmp router-alert
```

Variable Definitions

Use the definitions in the following table to use the **ip igmp** command.

Variable	Value
<code>access-list WORD<1-64> {A.B.C.D/X} <eny-tx deny-rx deny-both allow-only-tx allow-only-rx allow-only-both></code>	<p>Specifies the name of the access list from 1–64 characters.</p> <p>Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.</p> <p>Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic</p>
<code>compatibility-mode</code>	<p>Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command:</p> <pre>default ip igmp compatibility-mode</pre> <p>, or use the no option to disable compatibility mode:</p> <pre>no ip igmp compatibility-mode</pre>
<code>dynamic-downgrade-version</code>	<p>Configures if the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:</p> <pre>default ip igmp dynamic-downgrade-version</pre> <p>or use the no option to disable downgrade:</p> <pre>no ip igmp dynamic-downgrade-version</pre>
<code>igmpv3-explicit-host-tracking</code>	Enables explicit host tracking on IGMPv3. The default state is disabled.
<code>immediate-leave</code>	Enables fast leave on a port.
<code>last-member-query-interval <0-255></code>	<p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decreasing the value reduces the time to detect the loss of the last member of a group. The default is 10 tenths of a second. You should configure this value between 3–10 (equal to 0.3 – 1.0 seconds).</p>

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Configures IGMP for a specific port.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>query-interval <1-65535></code>	<p>Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.</p>
<code>query-max-response <0-255></code>	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds).</p> <p>Important: You must configure this value lower than the query-interval.</p>
<code>robust-value <2-255></code>	<p>Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.</p>
<code>router-alert</code>	<p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p>Important: To maximize network performance, configure this parameter according to the version of IGMP currently in use:</p> <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable

Variable	Value
<code>stream-limit stream-limit-max-streams <0-65535></code>	Configures multicast stream limitation on a port to limit the number of concurrent multicast streams on the port. The default is 4.
<code>version <1-3></code>	Configures the version of IGMP for this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configure IGMP on a VRF

You configure IGMP on a VRF instance the same way you configure IGMP for the Global Router, except that you must use VRF Router Configuration mode.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable SSM dynamic learning:

```
ip igmp ssm dynamic-learning
```

3. Configure the range group:

```
ip igmp ssm group-range {A.B.C.D/X}
```

The system displays the following message:

```
Warning: Changing the SSM range will cause all spb-multicast and spb-
pim-gw enabled interfaces to be internally bounced. Do you wish to
continue? (y/n) ? (y/n)?
```

Enter *y* to continue.

4. Enable the SSM map table for all static entries:

```
ip igmp ssm-map all
```

5. Create a static entry for a specific group:

```
ip igmp ssm-map {A.B.C.D} {A.B.C.D} enable
```

6. Enable the generation of IGMP traps:

```
ip igmp generate-trap
```

7. Enable the generation of IGMP log messages:

```
ip igmp generate-log
```


8. Configure the fast leave mode:

```
ip igmp immediate-leave-mode {multiple-user|one-user}
```

Example

For the VRF Red context, configure a new IP multicast group address and create an SSM map table entry for the multicast group and the source at 192.32.99.151. Configure the administrative state to enable all the static SSM map table entries.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router vrf red
Switch:1(router-vrf)#ip igmp ssm group-range 232.1.1.10/32

WARNING: Changing the SSM range will cause all spb-multicast and spb-pim-gw enabled
interfaces to be internally bounced. Do you wish to continue? (y/n) ? (y/n)? y
Switch:1(router-vrf)#ip igmp ssm-map 232.1.1.10 192.32.99.151
Switch:1(router-vrf)#ip igmp ssm-map all
```

Variable Definitions

Use the definitions in the following table to use the **ip igmp** command on a VRF.

Variable	Value
<i>generate-log</i>	Enables the generation of IGMP log messages. The default is disabled.
<i>generate-trap</i>	Enables the generation of IGMP traps. The default is disabled.
<i>immediate-leave-mode</i> <i>{multiple-user one-user}</i>	<ul style="list-style-type: none"> multiple-user: Removes (from the group) the IGMP member who sent the leave message. The default is multiple-user. one-user: Removes all group members on a fast leave enabled interface port after receiving the first leave message from a member.
<i>ssm dynamic-learning</i>	Enables dynamic learning from IGMPv3 reports. The default is enabled.
<i>ssm group-range</i> <i>{A.B.C.D/X}</i>	<p>Changes the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address.</p> <p>This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.</p>
<i>ssm-map</i> <all {A.B.C.D} {A.B.C.D} enable	<p>Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that sends traffic to the group.</p> <p>Enables the administrative state for a specific entry (group). This variable does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.</p>

View IP Multicast Threshold Exceeded Statistics

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View statistics:

```
show sys stats ipmc-threshold-exceeded-cnt
```

Example

```
Switch:1>show sys stats ipmc-threshold-exceeded-cnt
SourceGroupThresholdExceeded : 7372
EgressStreamThresholdExceeded : 7331
```

IP Multicast Basic Configuration Using EDM

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts use a host membership protocol to subscribe to multicast services. The Internet Group Management Protocol (IGMP) is an example of an IPv4 host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast–Sparse Mode (PIM–SM) and Protocol Independent Multicast–Source Specific Multicast (PIM–SSM).

Configure Multicast on the Switch

Perform this procedure to configure PIM and IGMP Snooping in an SMLT environment. The configuration steps show how to enable multicast, and then configure the usual PIM and IGMP Snooping related VLANs and global attributes. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

Before You Begin

SPBM must not be enabled on the vIST peers or any router participating in the PIM network.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Boot Config** tab.

3. Clear **EnableSpbmConfigMode** to disable the boot flag.

The system responds with these messages:

```
Warning: Please save the configuration and reboot the switch for this
to take effect.
```

```
Warning: Please carefully save your configuration file before
rebooting the switch. Saving configuration file when spbm-config-
mode is changed to disable, removes SPBM configurations from the
configuration file.
```

4. Select **Apply**.
5. Save the configuration, and then reboot the switch.



Important

Any change to **EnableSpbmConfigMode** requires a reboot for the change to take effect.

6. Configure the SMLT MLT:
 - a. Expand **Configuration > VLAN > MLT/LACP**.
 - b. Select the **MultiLink/LACP Trunks** tab.
 - c. Select **Insert**.
 - d. In **Id**, type the ID number of the MLT.
 - e. In **PortMembers**, select the ellipsis (...).
 - f. In the **Port Editor: PortMembers** dialog box, select the desired ports.
 - g. Select **Ok**.
 - h. Select **Insert**.
7. Configure the vIST MLT:
 - a. Repeat steps 6a to 6g to configure the MLT.
 - b. Select **MltVistEnable** to enable Simplified vIST.



Note

MltVistEnable enables Simplified vIST and is only available when the **EnableSpbmConfigMode** boot flag is disabled.

- c. Select **Insert**.
8. Create the vIST VLAN:
 - a. Expand **Configuration > VLAN > VLANs**
 - b. In **Basic**, select **Insert**.
 - c. In **Id**, enter an unused VLAN ID, or use the ID provided.
 - d. In **MstpInstance**, select the down arrow, and then choose an MSTI instance from the list.
 - e. In **Type**, select **byPort**.
 - f. Select **OK**.
 - g. Select **Insert**.
 - h. Select the vIST VLAN from the list of VLANs, and then select **IP**.

- i. Select **Insert**.
 - j. Configure the IP address for the vIST VLAN.
9. Repeat Step 8 to create an SMLT VLAN and assign the SMLT MLT ID to it; do not use the vIST MLT ID.
10. Configure PIM or IGMP Snooping on the SMLT VLAN:
 - a. To enable PIM, select the SMLT VLAN from the list of VLANs and select **IP > PIM**. Select **Enable** and select **Apply**.
 - b. To enable IGMP Snooping, select the SMLT VLAN from the list of VLANs and select **IP > IGMP**. Select **SnoopEnable** and select **Apply**.
11. Configure PIM on the SMLT VLAN:
 - a. To enable PIM, select the SMLT VLAN from the list of VLANs and select **IP > PIM**. Select **Enable** and select **Apply**.
12. Select **IP > PIM > Globals** to enable PIM globally.
13. Select **Enable**, and then select **Apply**.

Enable IPv4 PIM-SM Globally

Enable PIM-SM to offer multicasting services. After you enable PIM-SM globally and on a particular interface, the IGMP parameters take effect.

Before You Begin



Note

Before you can enable the PIM Infinite Threshold Policy feature, you must first disable the following:

- PIM-SM
- PIM-SSM
- Simplified vIST

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Globals** tab.
4. Select **sm** (sparse mode).
5. Select the **Enable** check box.
6. Select **Apply**.
7. Configure PIM Infinite Threshold Policy:
 - a. To disable PIM, clear the **Enable** check box, and then select **Apply**.
 - b. Select **enable**, and then select **Apply**.
 - c. To enable PIM, select the **Enable** check box, and then select **Apply**.

Globals Field Descriptions

Use the descriptions in the following table to use the **Globals** tab.

Name	Description
Mode	Configures the mode on the routing switch: sm (Sparse Mode) or ssm (Source Specific Multicast).
Enable	Enables or disables PIM.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The range is from 1–18724 and the default is 60 seconds.
RegisterSuppTimer	Specifies how long (in seconds) the designated router suppresses sending registers to the rendezvous point (RP). The timer starts after the designated router receives a register-stop message from the RP. The range is from 6–65535 and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager for unicast routing information updates for PIM. The range is from 2–65535 and the default is 5 seconds. Important: If you lower this value, it increases how often the switch polls the routing table manager. This value can affect the performance of the switch, especially if a high volume of traffic flows through the switch.
DiscardDataTimeOut	Specifies how long (in seconds) to discard data until the switch receives a join message from the RP. An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The range is from 5–65535 and the default is 60 seconds.
CRPADVTimeOut	Specifies how often (in seconds) a router configured as a candidate rendezvous point router (C-RP) sends advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected bootstrap router (BSR). The range is from 5–26214 and the default is 60 seconds.
BootStrapPeriod	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. The range is from 5–32757 and the default is 60 seconds.
StaticRP	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism.

Name	Description
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value ages PIM mroutes in seconds. The range is from 10–86400 and the default value is 210. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.
FastJoinPrune	Enables or disables the PIM fast join prune feature.
SptInfiniteThreshold	Enables or disables PIM Infinite Threshold Policy, so that multicast traffic follows the shared tree path through a Rendezvous Point (RP) instead of automatically switching over to shortest path tree (SPT). The default value is disabled, which means that multicast traffic is automatically switched over to SPT.

Enable IPv6 PIM-SM Globally

Enable IPv6 PIM-SM to offer multicasting services. After you enable IPv6 PIM-SM globally and on a particular interface, the MLD parameters take effect.

Before You Begin



Note

Before you can enable the PIM Infinite Threshold Policy in **SptInfiniteThreshold**, you must first disable the following:

- PIM-SM
- PIM-SSM
- Simplified vIST

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Globals** tab.
4. Select **Enable**.
5. Select **sm** (sparse mode).
6. Select **Apply**.

Globals Field Descriptions

Use the descriptions in the following table to use the **Globals** tab.

Name	Description
Enable	Enables or disables PIM.
Mode	Configures the mode on the routing switch: sm (Sparse Mode) or ssm (Source Specific Multicast).

Name	Description
RegisterSuppTimer	Specifies how long (in seconds) the designated router suppresses sending registers to the rendezvous point (RP). The timer starts after the designated router receives a register-stop message from the RP. The range is from 10–65535 and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager for unicast routing information updates for PIM. The range is from 2–65535 and the default is 5 seconds. Important: If you lower this value, it increases how often the switch polls the routing table manager. This value can affect the performance of the switch, especially if a high volume of traffic flows through the switch.
DiscardDataTimeOut	Specifies how long (in seconds) to discard data until the switch receives a join message from the RP. An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The range is from 5–65535 and the default is 60 seconds.
StaticRP	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism.
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value ages PIM mroutes in seconds. The range is from 10–86400 and the default value is 210. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The range is from 1–18724 and the default is 60 seconds.
SptInfiniteThreshold	Enables or disables PIM Infinite Threshold Policy, so that multicast traffic follows the shared tree path through a Rendezvous Point (RP) instead of automatically switching over to shortest path tree (SPT). To change this value, you must first disable: <ul style="list-style-type: none"> • PIM-SM • PIM-SSM • Simplified vIST The default value is disabled, which means that multicast traffic is automatically switched over to SPT.

Enable PIM on a Port

Enable PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- You must enable PIM globally before you enable it on an interface.
- The interface uses a valid IP address.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **PIM** tab.
5. Select **Enable**.
6. Select **Apply**.

PIM Field Descriptions

Use the data in the following table to use the **PIM** tab.

Name	Description
Enable	Enables (true) or disables (false) PIM for the specified port.
Mode	Displays the mode currently running on the routing switch.
IntfType	Indicates the interface type as active or passive.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724 seconds.
CBSRPreference	Configures the preference for this local interface to become a candidate BSR (C-BSR). The C-BSR with the highest BSR priority and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR. The range is -1-255.

Enable IPv6 PIM on a Port

Enable IPv6 PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- You must enable IPv6 interface before you enable PIM on a port.

About This Task

You can also right-click the port and use the **Edit IPv6** shortcut menu to reach this same tab.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select the **PIM** tab.
5. Select **Enable**.
6. Select **Apply**.

PIM Field Descriptions

Use the data in the following table to use the **PIM** tab.

Name	Description
Address	Specifies the IPv6 address of the PIM interface.
NetMask	Specifies the network mask for the IPv6 address of the PIM interface.
Enable	Enables (true) or disables (false) PIM for the specified port.
Mode	Displays the mode currently running on the routing switch.
DR	Specifies the designated router on this PIM interface.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724 seconds.
OperState	Specifies the current operational state of this PIM interface.
Type	Specifies the type of interface.

Enable SSM Globally

Enable Source Specific Multicast (SSM) to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

Before You Begin

- Configure a unicast protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP, see [RIP configuration using EDM](#) on page 3173. For more information about OSPF, see [OSPF Configuration Using EDM](#) on page 2857.
- Enable PIM globally.



Important

After you enable PIM in SSM mode, the IGMP parameters take effect. To take full advantage of SSM, enable IGMPv3 if hosts that attach to the switch run IGMPv3 or configure the SSM table.

About This Task

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Globals** tab.
4. Select **ssm** (source specific multicast).
5. Select the **Enable** check box.
6. Select **Apply**.

The system displays the following message:

```
Are you sure you want to change the PIM mode? The traffic will not be
stopped immediately. All Static Source Group entries in the SSM range
will be deleted. Do you wish to continue?
```

7. Select **Yes**.

Enable IPv6 SSM Globally

Enable Source Specific Multicast (SSM) to optimize IPv6 PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the IPv6 PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

Before You Begin

- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM.

For more information about RIPng, see [RIPng Configuration using EDM](#) on page 3181. For more information about OSPFv3, [OSPFv3 Configuration using EDM](#) on page 2896.

- Enable PIM globally.



Important

After you enable IPv6 PIM in SSM mode, the MLD parameters take effect. To take full advantage of SSM, enable MLDv2 if hosts that attach to the switch run MLDv2.

About This Task

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On a SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Globals** tab.
4. Select the **Enable** check box.
5. Select **ssm** (source specific multicast).
6. Select **Apply**.

The system displays the following message:

```
Warning: RP entries in the SSM range will be deleted
```

```
Do you wish to continue? (y/n)?
```

7. Select **Yes**.

Enable PIM on a VLAN Interface

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- You must enable PIM globally before you enable it on an interface.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Select **IP**.
6. Select the **PIM** tab.
7. Select **Enable**.

8. Select **Apply**.

PIM Field Descriptions

Use the data in the following table to use the **PIM** tab.

Name	Description
Enable	Enables (true) or disables (false) PIM.
Mode	Displays the mode that currently runs on the switch. The valid modes are SSM and Sparse. This variable is a read-only field.
IntfType	Specifies the type of interface: active or passive.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724.
CBSRPreference	Configures the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR. The range is -1-255.

Enable PIM on a CLIP Interface

Enable Protocol Independent Multicasting (PIM) on a CLIP interface so that it can participate in PIM routing.

Before You Begin

- You must globally enable PIM.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select the required CLIP interface.
5. Select **PIM**.
6. Select **Enable**.

You must enable PIM on the CLIP interface for PIM to function. The mode is indicated on this tab.

7. Select **Apply**.

Circuitless PIM Field Descriptions

Use the data in the following table to use the **Circuitless PIM** tab.

Name	Description
Enable	Enables PIM on the CLIP interface.
Mode	Specifies the PIM mode.

Enable IPv6 PIM on a VLAN Interface

Configure IPv6 PIM for each interface to enable the interface to perform multicasting operations.

Before You Begin

- You must enable IPv6 PIM globally before you enable it on an interface.

Procedure

- In the navigation pane, expand **Configuration > VLAN**.
- Select **VLANS**.
- Select the **Basic** tab.
- Select the VLAN ID that you want to configure with PIM.
- Select **IPv6**.
- Select the **PIM** tab.
- Select **Enable**.
- Select **Apply**.

PIM Field Descriptions

Use the descriptions in the following table to use the **PIM** tab.

Name	Description
IfIndex	Specifies the interface index for PIM.
Address	Specifies the IPv6 address of the PIM interface.
Netmask	Specifies the network mask for the IPv6 address of the PIM interface.
Enable	Enables (true) or disables (false) PIM.
Mode	Displays the mode that currently runs on the switch. The valid modes are SSM and Sparse. This variable is a read-only field.
DR	Specifies the designated router on this PIM interface.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724.

Name	Description
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724.
OperState	Specifies the current operational state of this PIM interface.
Type	Specifies the type of interface.

Configure IGMP Parameters on a Port

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **IGMP** tab.
5. Edit the appropriate values.



Note

When you configure the following IGMP parameters on the DvR Controllers in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

- Version
- QueryInterval
- QueryMaxResponseTime
- Robustness
- LastMembQueryIntvl
- CompatibilityModeEnable



Note

To use the fast leave feature on IGMP, enable explicit-host-tracking.

6. Select **Apply**.

IGMP Field Descriptions

Use the data in the following table to use the **IGMP** tab.

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds).</p> <p>Important: You must configure this value lower than the QueryInterval.</p>
Robustness	<p>Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value.</p> <p>The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.</p>
LastMembQueryIntvl	<p>Configures the maximum response time (in 1/10 seconds) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Configure this parameter to values greater than 3. If you do not require a fast leave process, Use values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)</p>
SnoopEnable	Enables snoop on the interface. The default is disabled.
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
ProxySnoopEnable	Enables proxy snoop on the interface. The default is disabled.
Version	Configures the version of IGMP (1, 2 or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this port.

Name	Description
Maximum Number Of Stream	Configures the maximum number of streams this port permits. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This variable is a read-only value.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router. Important: Configure this variable only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.
RouterAlertEnable	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. To maximize network performance, configure this parameter according to the version of IGMP currently in use: <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
DynamicDowngradeEnable	Configures if the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.

Configure IGMP Parameters on a VLAN

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select **IGMP**.
7. Configure the relevant variables.



Note

When you configure the following IGMP parameters on the DvR Controllers in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

- Version
- QueryInterval
- QueryMaxResponseTime
- Robustness
- LastMembQueryIntvl
- CompatibilityModeEnable

Configuration of IGMP snooping is not supported on DvR enabled Layer 2 VSNs.

8. Select **Apply**.

IGMP Field Descriptions

Use the data in the following table to use the **IGMP** tab.

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the IGMP host query packets transmit on the interface. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds.)</p> <p>Important: You must configure this value lower than the QueryInterval.</p>

Name	Description
Robustness	<p>Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value.</p> <p>The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.</p>
LastMembQueryIntvl	<p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Configure this parameter to values greater than 3. If you do not require a fast leave process, use values greater than 10. (The value 3 is equal to 0.3 seconds, and 10 is equal to 1 second.)</p>
SnoopEnable	Enables snoop on the interface. The default is disabled.
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
ProxySnoopEnable	Enables proxy snoop on the interface. The default is disabled.
Version	Configures the version of IGMP (1, 2, or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables or disables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this VLAN.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this VLAN. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	<p>Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router.</p> <p>Important: Configure this field only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.</p>

Name	Description
RouterAlertEnable	<p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p>To maximize network performance, configure this parameter according to the version of IGMP currently in use:</p> <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
DynamicDowngradeEnable	<p>Configures if the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.</p>
CompatibilityModeEnable	<p>Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.</p>
ExplicitHostTrackingEnable	<p>Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.</p>
SnoopQuerierEnable	<p>Enables snoop querier. The default is disabled.</p> <p>When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups.</p> <p>Enable Layer 2 Querier on only one node in the VLAN.</p>
SnoopQuerierAddr	<p>Specifies the pseudo IP address of the IGMP snoop querier. The default IP address is 0.0.0.0.</p>

Multicast Listener Discovery

Table 153: Multicast Listener Discovery product support

Feature	Product	Release introduced
Multicast Listener Discovery (MLD)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

MLD Fundamentals

MLD is an asymmetric protocol that specifies separate behaviors for multicast address listeners (hosts or routers that listen to multicast packets) and multicast routers. Each multicast router learns, for each directly attached link, which multicast addresses and which sources have listeners on that link. The information that MLD gathers is provided to the multicast routing protocols that the router uses. This information ensures that multicast packets arrive at all links where listeners require such packets.

A multicast router can itself be a listener of one or more multicast addresses; that is, the router performs both the multicast router role and the multicast address listener part of the protocol. The router collects the multicast listener information needed by the multicast routing protocol and informs itself and other neighboring multicast routers of the listening state.

IPv6 routers use MLD to discover:

- the presence of multicast listeners on directly attached links
- multicast addresses required by neighboring nodes

MLD Versions

The purpose of the MLD protocol in the IPv6 multicast architecture is to allow an IPv6 router to discover the presence of multicast listeners on directly-attached links and to discover which multicast addresses are of interest to neighboring nodes. MLD is the direct IPv6 replacement for the IGMP protocol used in IPv4. The MLD implementation described in this document is based on the MLDv2 standard, which is a backward-compatible update to the MLDv1 standard.

There are three versions of IGMP, and two versions of MLD. IGMPv2 is equivalent in function to MLDv1 and IGMPv3 is equivalent to MLDv2.

MLD Querier

MLD Querier is similar to IGMP Querier. A multicast query router communicates with hosts on a local network by sending MLD queries. This router periodically sends a general query message to each local network of the router. This is standard multicast behavior.

**Note**

Queries are sent only if PIM is enabled globally and on the interface. PIM and snooping cannot be enabled at the same time.

Each VLAN using MLD multicast must have a router performing multicast queries. Networks with no stand-alone devices currently have no capability for implementing the pruning of multicast traffic. A dedicated querier must be available on the network.

There are several behavioral differences between a traditional query router and a switch or stack using the MLD Querier functionality. The following are the differences:

- There is no election process. When a switch or stack restarts, queries are sent as part of MLD startup. This process stops other devices from sending queries while they detect the new device starting up. The last active device sending queries on the network is the active one. This is not the case with Layer 3 MLD behavior.
- If the current active device stops sending queries, a timeout period must elapse before another device takes over. This can result in an ageout of groups, and subsequent flooding, before a new query is sent and the pruning process restarts. This occurs only during the transition between active query devices. Once the new device is established, queries are sent as configured in the Query Interval and Robust Values fields.
- Multiple active query devices are not supported. Enabling multiple devices establishes one active device and other devices listening to take over should the active device fail.

The querier version is determined by the received query version and establishes the interface operational version. By default, the interface operational version is MLDv1. If the interface operational version is downgraded from MLDv2 to MLDv1 (when operational version is MLDv2 and a MLDv1 query is received), then all MLDv2 listeners (registered by MLDv2 reports) are removed and all incoming MLDv2 reports are dropped.

MLD Snooping

MLD snooping is an IPv6 multicast constraining mechanism running on Layer 2 devices. When MLD snooping is enabled on a VLAN, the switch examines the MLD messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. Based on the learning, the switch forwards multicast traffic only to those interfaces in the VLAN that are connected to the interested receivers instead of flooding traffic to all the interfaces.

The following figure shows an example of this scenario. On the left side of the figure, IPv6 multicast packets are transmitted when MLD snooping is not enabled. All the hosts that are interested and not interested receive the IP Multicast traffic consuming

bandwidth. Whereas, on the right side of the figure, when MLD snooping is enabled and IPv6 multicast packets are transmitted, only the interested hosts receive the IP multicast packets.




Figure 140: IPv6 multicast packet transmission when MLD snooping is enabled and not enabled

The following figure shows IPv6 multicast packets transmitted when MLD v2 snooping is enabled and not enabled.




Figure 141: IPv6 multicast packet transmission when MLD v2 snooping is enabled and not enabled

MLD Snooping Configuration Guidelines and Restrictions

You can perform the following configurations to manage and control IPv6 multicast groups using the MLD snooping feature:

- Enable or disable MLD snooping on each VLAN. MLD snooping can be enabled on a maximum of 512 VLANs.
- Enable IGMP snooping and MLD snooping on the same VLAN.

Limitations

Following are the limitations for MLD snooping configuration:

- The maximum (S,G,V) entries supported in the IPv6 multicast routing table (L3_ENTRY_IPV6_MULTICAST) is 512.

MLD snooping shares the (S,G,V) entries with IGMP snooping, where the (S,G,V) entries number = (G,V) MLD_V1 type entries number + (S,G,V) MLD_V2 type entries number + (*,G,V) MLD_V2 type entries number + number of groups without (*,G,V) registered listeners.

- IPv6 MLD proxy functionality is not supported.

- Multicast Flood Control (MFC) is not supported.
- Static mrouter ports cannot be configured.
- IPv6 MLD send query functionality is not supported.
- Configure static router ports is not supported.
- Multicast Listener Discovery (MLDv1/v2) is not supported with vIST/SMLT.
- Multicast Listener Discovery (MLDv1/v2) is not supported with Fabric Connect.
- Protocol Independent Multicast (PIM) with IPv6 is not supported vIST/SMLT.
- Protocol Independent Multicast (PIM) with IPv6 is not supported with Fabric Connect (SPBM).

MLD Configuration Using the CLI

This section provides procedures to configure Multicast Listener Discovery (MLD) using the command line interface (CLI).

Enable MLD Trap Generation

About This Task

Use this procedure to enable MLD traps.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable MLD trap generation:
`ipv6 mld generate-trap`

Enable MLD Log Status

About This Task

Use this procedure to enable MLD traps.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable MLD log status:
`ipv6 mld generate-log`

Configure MLD Version

About This Task

Use this procedure to configure MLD version.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure MLD version:

```
ipv6 mld version <1-2>
```



Note

For MLD to function correctly, the MLD version must be the same on all routers in the network.

Variable Definitions

The following table describes the variables for the **ipv6 mld version** command.

Variable	Value
<1-2>	Indicates the version of MLD that runs on this interface.

Configure the MLD Last Listener Query Interval

About This Task

Use this procedure to configure the last listener query interval in seconds for the MLD interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the last listener query interval:

```
ipv6 mld last-listener-query-interval <0-60>
```

Variable Definitions

The following table describes the variables for the **ipv6 mld last-listener-query-interval** command.

Variable	Value
<0-60>	Indicates the last listener query interval in seconds.

Configure the MLD Query Interval

About This Task

Use this procedure to configure the query interval for the MLD interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} OR interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the query interval for the MLD interface:

```
ipv6 mld query-interval <1-65535>
```

Variable Definitions

The following table describes the variables for the **ipv6 mld query-interval** command.

Variable	Value
<1-65535>	Indicates the frequency at which MLD host query packets transmit on this interface.

Configure the MLD Query Maximum Response Time

About This Task

Use this procedure to configure the query maximum response time for mld interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the query maximum response time for mld interface:

```
ipv6 mld query-max-response-time <0-60>
```

Variable Definitions

The following table describes the variables for the **ipv6 mld query-max-response-time** command.

Variable	Value
<0-60>	Indicates the query maximum response interval time in seconds.

Configure the MLD Robustness

About This Task

The robustness value allows the tuning for the expected packet loss on a link. If a link expects packet loss, increase the robustness variable value.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the MLD robustness:

```
ipv6 mld robust-value <2-255>
```

Variable Definitions

The following table describes the variables for the **ipv6 mld robust-value** command.

Variable	Value
<2-255>	Specifies a numerical value for MLD snooping robustness.

Enable MLD Snooping on a VLAN

About This Task

Use this procedure to enable MLD snooping on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable MLD snooping:

```
ipv6 mld snooping
```

Enable MLD SSM-Snooping on a VLAN

About This Task

Use this procedure to enable IPv6 MLD ssm-snooping on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable MLD snooping:

```
ipv6 mld ssm-snoop
```

Display MLD Snooping Configuration Status

About This Task

Use this procedure to display information about the MLD snooping configuration for the switch.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the switch MLD snooping configuration status:

```
show ipv6 mld snooping
```

Example

```
Switch:1>show ipv6 mld snooping
=====
Mld Snooping - GlobalRouter
=====
IFINDEX  SNOOP   SSM    ACTIVE      MROUTER
          ENABLE SNOOP   MROUTER     EXPIRATION
          ENABLE PORTS      TIME
-----
V666     False  False  NONE        0
1 out of 1 entries displayed
```

Display MLD Snooping Tracing Information

About This Task

Use this procedure to display MLD snooping tracing information.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the MLD snooping tracing information:
show ipv6 mld snoop-trace

Example

```
Switch:1>show ipv6 mld snoop-trace
=====
Mld Snoop Trace - GlobalRouter
=====
GROUP/
SOURCE          IN    IN    OUT    OUT    TYPE
ADDRESS         VLAN  PORT VLAN  PORT
-----
ff10:0:0:0:0:0:1/ 10    1/15 10    1/16   ACCESS
5051:0:0:0:0:1:84:51
```

Display MLD Interface Information

About This Task

Use this procedure to display MLD snooping interface parameters.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display MLD interface information:
show ipv6 mld interface [gigabitethernet {slot/port[/sub-port]}] [vlan <1-4059>]

Examples

```
Switch:1>show ipv6 mld interface
=====
Mld Interface - GlobalRouter
=====
IF      STATUS  VERS  OPER  VERS  QUERIER                                Wrong Query JOINS  MODE
-----
Pl/3    inact   2     2     2001:0db8:3c4d:0015:0000:0000:1a2f:1aaa  0              0    pim
V666    inact   2     2     2001:0db8:3c4d:0015:0000:0000:1a2f:1bbb  0              0    pim

Switch:1>show ipv6 mld interface vlan 10
=====
Vlan IPv6 Mld
=====
VLAN  QUERY  QUERY  ROBUST  VERSION  LAST  SNOOP  SSM      DYNAMIC
ID    INTVL  MAX                    LIST  ENABLE  SNOOP  DOWNGRADE
      RESP                    QUERY    ENABLE
-----
10    125    10     2        1        1     false  false    enabled

Switch:1>show ipv6 mld interface gigabitethernet 1/11
=====
Port IPv6 MLD
=====
PORT  QUERY  QUERY  ROBUST  VERSION  LAST  DYNAMIC
NUM   INTVL  MAX                    LIST  DOWNGRADE
      RESP                    QUERY
-----
1/11  125    10     2        1        1     enabled

1 out of 1 entries displayed
```

Variable Definitions

The following table describes the variables for the **show ipv6 mld interface** command.

Variable	Value
vlan <1-4059>	Displays MLD snooping information for the configured VLANs.
gigabitEthernet {slot/port[/sub-port]}	Displays MLD snooping information on a specific interface.

Display MLD System Parameters

About This Task

Use this procedure to display information about the MLD traps and logs.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the system parameters:
show ipv6 mld sys

Example

```
Switch:1#show ipv6 mld sys
=====
Mld System Parameters - GlobalRouter
=====
generate-trap : disable
generate-log  : disable
```

*Display MLD Cache Information***About This Task**

Use this procedure to display the learned multicast groups in the cache.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the learned multicast groups in the cache:
show ipv6 mld cache

Example

```
Switch:1>show ipv6 mld cache
=====
Mld Cache Information
=====
GRPADDRESS/LASTREPORTER      INTERFACE EXPIRATION
-----
ff03:0:0:0:0:0:0:0/          Vlan10    0 day(s), 00h:04m:12s
fe80:0:0:0:200:9aff:fe68:3dd5
1 out of 1 entries displayed
```

*Display the MLD Group Information***About This Task**

Use this procedure to display the MLD group information to show the learned multicast groups and the attached ports.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the MLD group information:
show ipv6 mld group [count] [group] [member-subnet]

Examples

```
Switch:1>show ipv6 mld group
=====
Mld Group - GlobalRouter
=====
GRPADDR/MEMBER              INPORT      EXPIRATION
-----
ff1e:0000:0000:0000:0000:0002:4444/ V666-1/4    0
```

```

2001:0db8:3c4d:0015:0000:0000:1a2f:1a2c

1 out of 1 group Receivers displayed

Total number of unique groups 1 text
Switch:1>show ipv6 mld group group ff1e:0000:0000:0000:0000:0002:4444 detail

=====
                        Mld Group Detail - GlobalRouter
=====

Interface:                Vlan666-1/4
MLDv2 Group:              ff1e:0000:0000:0000:0000:0002:4444
Interface Group Mode:     EXCLUDE
Interface Compatibility Mode: MLD_V2
Interface Group Timer:    258
Vl Host Timer:            Not Running
Interface Group Include Source List:
    Source Address                Expires
    2001:0db8:3c4d:0015:0000:0000:1a2f:1aaa        258
Interface Group Exclude Source List :
    Source Address                Expires
    2001:0db8:3c4d:0015:0000:0000:1a2f:1bbb        N/A

```

View IPv6 MLD Host Cache

View the learned multicast group addresses in the host cache.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View IPv6 MLD host cache:

```
show ipv6 mld-host-cache
```

MLD Configuration Using EDM

This section provides procedures to configure Multicast Listener Discovery (MLD) using Enterprise Device Manager (EDM).

Configure MLD Globally

About This Task

Use the following procedure to configure MLD parameters for the switch.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Globals** tab.
4. Configure the MLD global parameters as required.
5. Select **Apply** to save the changes.
6. Select **Refresh** to update the changes.

Globals Field Description

Use the data in the following table to use the **Globals** tab.

Field	Description
GenerateTrap	Enables MLD to generate traps.
GenerateLog	Enables MLD to generate logs.

View the MLD SSM Global Information

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Ssm Globals** tab.

Ssm Globals Field Description

Use the data in the following table to use the **Ssm Globals** tab.

Field	Description
RangeGroup	Specifies the ssm range.
RangeMask	Specifies the ssm range mask.

MLD Interface Configuration

Configure the interfaces so that the switch forwards multicast traffic only to those interfaces in the VLAN that are connected to the interested receivers instead of flooding traffic to all the interfaces.

Configure an MLD Interface

Perform this procedure to change the configuration of existing MLD interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Interfaces** tab.
4. Select **Insert**.
5. Configure the MLD interface parameters.
6. Select **Insert**.
7. Select **Apply** to save the changes.
8. Select **Refresh** to update the changes.

Interfaces Field Description

Use the data in the following table to use the **Interfaces** tab.

Field	Description
IfIndex	Specifies the internetwork layer interface value of the interface for which MLD is enabled.
QueryInterval	Specifies the frequency at which MLD host-query packets are transmitted on this interface. Values range from 1 to 65535.
Version	Indicates the MLD version.
Querier	Specifies the address of the MLD Querier on the IPv6 subnet to which this interface is attached.
QueryMaxResponseDelay	Specifies the maximum query response time advertised in MLD queries on this interface. Values range from 0 to 60.
Joins	Specifies the number of times a group membership has been added on this interface.
Groups	Specifies the current number of entries for this interface in the cache table.
Robustness	Specifies the robustness variable tuning for the expected packet loss on a subnet. If a subnet is expected to experience loss, the robustness variable can be increased. Values range from 2 to 255.
LastListenQueryIntvl	Specifies the maximum response delay inserted into the group-specific queries sent in response to the leave group messages. It also indicates the amount of time between group-specific query messages. Values range from 0 to 60. This value can be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
SnoopEnable	Indicates if snooping is enabled.
FlushAction	Specifies the MLD flush action as one of the following: <ul style="list-style-type: none"> flushGrpMember flushMrouter flushSender
SsmEnable	Indicates if ssm is enabled.
NewQuerier	Specifies the IPv6 address of the new MLD querier.

Field	Description
DynamicDowngradeEnable	Enables dynamic downgrade of the MLD version when older version query message is received.
OperVersion	Specifies the operational version of the MLD running on this interface.
McastMode	Specifies the MLD interface mode as one of the following: <ul style="list-style-type: none"> • snoop • pim • snoopSpb • routerSpb • dvmrp • none

Configure MLD on a Port

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select the **MLD** tab.
5. Configure the MLD interface parameters.
6. Select **Apply** to save the changes.
7. Select **Refresh** to update the changes.

MLD Field Description

Use the data in the following table to use the **MLD** tab.

Field	Description
QueryInterval	Specifies the frequency at which MLD host-query packets are transmitted on this interface. The default is 125.
Version	Indicates the MLD version. The default is 1.
Querier	Specifies the address of the MLD Querier on the IPv6 subnet to which this interface is attached.
QueryMaxResponseDelay	Specifies the maximum query response time advertised in MLD queries on this interface. The default is 10.
Joins	Specifies the number of times a group membership has been added on this interface.
Groups	Specifies the current number of entries for this interface in the cache table.

Field	Description
Robustness	Specifies the robustness variable tuning for the expected packet loss on a subnet. If a subnet is expected to experience loss, the robustness variable can be increased. The default is 2.
LastListenQueryIntvl	Specifies the maximum response delay inserted into the group-specific queries sent in response to the leave group messages. It also indicates the amount of time between group-specific query messages. The default is 1. This value can be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
SnoopEnable	Indicates if snooping is enabled.
FlushAction	Specifies the MLD flush action as one of the following: <ul style="list-style-type: none"> flushGrpMember flushMrouter flushSender
SsmEnable	Indicates if ssm is enabled.
NewQuerier	Specifies the IPv6 address of the new MLD querier.
DynamicDowngradeEnable	Enables dynamic downgrade of the MLD version when older version query message is received.
OperVersion	Specifies the operational version of the MLD running on this interface.
McastMode	Specifies the MLD interface mode as one of the following: <ul style="list-style-type: none"> snoop pim snoopSpb routerSpb dvmrp none

Configure MLD on a VLAN

Procedure

1. In the navigation pane, expand **Configuration > VLAN**
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN from the list.
5. Select the **IPv6** tab.

6. Select the **MLD** tab.
7. Configure the MLD interface parameters.
8. Select **Apply** to save the changes.
9. Select **Refresh** to update the changes.

MLD Field Description

Use the data in the following table to use the **MLD** tab.

Field	Description
QueryInterval	Specifies the frequency at which MLD host-query packets are transmitted on this interface. The default is 125.
Version	Indicates the MLD version. The default is 1.
Querier	Specifies the address of the MLD Querier on the IPv6 subnet to which this interface is attached.
QueryMaxResponseDelay	Specifies the maximum query response time advertised in MLD queries on this interface. The default is 10.
Joins	Specifies the number of times a group membership has been added on this interface.
Groups	Specifies the current number of entries for this interface in the cache table.
Robustness	Specifies the robustness variable tuning for the expected packet loss on a subnet. If a subnet is expected to experience loss, the robustness variable can be increased. The default is 2.
LastListenQueryIntvl	Specifies the maximum response delay inserted into the group-specific queries sent in response to the leave group messages. It also indicates the amount of time between group-specific query messages. The default is 1. This value can be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
SnoopEnable	Indicates if snooping is enabled.
FlushAction	Specifies the MLD flush action as one of the following: <ul style="list-style-type: none"> • flushGrpMember • flushMrouter • flushSender
SsmEnable	Indicates if ssm is enabled.

Field	Description
NewQuerier	Specifies the IPv6 address of the new MLD querier.
DynamicDowngradeEnable	Enables dynamic downgrade of the MLD version when older version query message is received.
OperVersion	Specifies the operational version of the MLD running on this interface.
McastMode	Specifies the MLD interface mode as one of the following: <ul style="list-style-type: none"> • snoop • pim • snoopSpb • routerSpb • dvmrp • none

Enable MLD Snooping

About This Task

Use the following procedure to enable MLD snooping on the switch.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select **Snooping** tab.
4. Double-click **SnoopEnable** for the desired port.
5. Select **true** to enable MLD snooping.
6. Double-click **SsmEnable**.
7. Select **true** to enable SSM.
8. Select **Apply**.

Snooping Field Description

Use the data in the following table to use the **Snooping** tab.

Field	Description
IfIndex	Specifies the interface on which you enabled MLD snooping. It specifies the port number if the interface is a brouter port or the VLAN number if the interface is a VLAN.
SnoopEnable	Indicates the status of MLD snooping on the specified interface: <ul style="list-style-type: none">• True – MLD snooping is enabled• False – MLD snooping is disabled
SsmEnable	Indicates the status of SSM on the specified interface: <ul style="list-style-type: none">• True – SSM is enabled• False – SSM is disabled

View MLD Snoop Trace Information

About This Task

Use this procedure to display information about the multicast groups traversing the snoop enabled router.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Snoop Trace** tab.

Snoop Trace Field Description

Use the data in the following table to use the **Snoop Trace** tab.

Field	Description
GrpAddr	Specifies the IP multicast address of the group traversing the router.
SrcAddr	Specifies the IP source address of the multicast group address.
OutVlan	Specifies the egress VLAN ID for the multicast group.
OutPort	Specifies the egress port of the multicast group.
InVlan	Specifies the ingress VLAN ID for the multicast source.

Field	Description
InPort	Specifies the ingress port for the multicast group.
Type	Specifies the port type on which the snoop entry is learnt.

View MLD Cache Information

About This Task

Use this procedure to display information about the learned multicast groups in the cache.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Cache** tab.

Cache Field Description

Use the data in the following table to use the **Cache** tab.

Field	Description
Address	The IPv6 multicast group address for which this entry contains information.
IfIndex	Indicates the internetwork-layer interface for which this entry contains information for an IPv6 multicast group address.
LastReporter	Indicates the source IPv6 address of the last membership report received for this IPv6 Multicast group address on this interface. If membership report is not received, the value is 0::0
ExpiryTime	Indicates the minimum amount of time remaining before the entry ages out.

View MLD V2 Cache Information

About This Task

Use this procedure to display information about the MLDv2 corresponding to each interface, port, and multicast group paired on a router.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **V2 Cache** tab.

V2 Cache Field Description

Use the data in the following table to use the **V2 Cache** tab.

Field	Description
GroupAddress	Specifies the multicast group address that others want to join. A group address can be the same for many incoming ports.
Ifindex	Identifies a physical interface or a logical interface (VLAN), which has received group reports from various sources.
InPort	Identifies a physical interface or a logical interface (VLAN), which has received group reports from various sources.
Version1HostTimer	Specifies the time remaining until the local router assumes that there are no more MLDv1 members on the IP subnet attached to the interface. This is applicable only for MLDv1 hosts. Upon receiving an MLDv1 report, this value is reset to the group membership timer.
SourceFilterMode	Specifies the current group state applicable on MLDv2 compatible nodes.

View IPv6 MLD Host Cache

View the learned multicast group addresses in the host cache.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Host Cache** tab.

Host Cache Field descriptions

Use the data in the following table to use the **Host Cache** tab.

Name	Description
Ifindex	Shows the index value that uniquely identifies the interface to which this entry applies.
GrpAddress	Shows the IP address for the multicast group.
GrpLocallyRegistered	Shows the Group Locally Registered for an IPv6 MLD host-cache entry.
GrpLastReporter	Shows the Group Last Reporter address for an IPv6 MLD host-cache entry.
GrpUpTime	Shows the Group Uptime for an IPv6 MLD host-cache entry.

Name	Description
GrpExpiryTime	Shows the Group Expiry Time for an IPv6 MLD host-cache entry.
GrpFilterMode	Shows the Group Filter Mode for an IPv6 MLD host-cache entry.

View MLD Source Information

About This Task

Use this procedure to display information about the MLD source.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Source** tab.

Source Field Description

Use the data in the following table to use the **Source** tab.

Field	Description
GroupAddress	Specifies the IPv6 multicast group address for which this entry contains information.
Ifindex	Specifies the interface for which this entry contains information for an IP multicast group address.
InPort	Identifies a physical interface or logical interface (VLAN), which has received group reports for this source.
HostAddress	Specifies the host address to which this entry corresponds.
MemberAddress	Specifies the IPv6 address of a member that has sent source specific report wishing to join this source.
Expire	Specifies the state of this entry.
Mode	Specifies the current member state. This is applicable to MLDv2 compatible nodes.
MemberExpire	Specifies the time until the member for this source expires.

View MLD Sender Information

About This Task

Use this procedure to display information about the multicast senders.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Sender** tab.

Sender Field Description

Use the data in the following table to use the **Sender** tab.

Field	Description
GrpAddr	Specifies the IPv6 multicast group address.
Ifindex	Specifies the interface index of the sender.
MemberAddr	Specifies the IPv6 host address.
Action	Specifies the MLD action as one of the following: <ul style="list-style-type: none">• none• flushEntry• flushGrp
Port	Specifies the MLD sender port.

View MLD Group Information

About This Task

Use this procedure to display information about the groups configured in this device.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 MLD**.
3. Select the **Group** tab.

Group Field Description

Use the data in the following table to use the **Group** tab.

Field	Description
IPv6Address	Specifies the multicast group address that others want to join to. A group address can be the same for many incoming ports.
Members	Specifies the IP address of a source that has sent group report wishing to join this group.
InPort	Identifies a physical interface or a logical interface which has received group reports from various sources.

Field	Description
Expiration	Specifies the time left before group report expires on this port. This is updated upon receiving a group report.
IfIndex	Identifies a physical interface or a logical interface which has received group reports from various sources.

PIM Configuration Using the CLI

The switch supports two modes of Protocol Independent Multicast (PIM): Sparse Mode (SM) and Source Specific Multicast (SSM).

- PIM-SM supports multicast groups spread out across large areas of a company or the Internet.
- PIM-SSM optimizes PIM-SM by simplifying the many-to-many model (servers-to-receivers).



Important

The **spbm-config-mode** boot flag must be disabled before you can configure PIM or IGMP. To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

Before You Begin

For an IPv4 PIM configuration using the CLI:

- Configure an IPv4 interface.

For more information, see [IP routing configuration using the CLI](#) on page 2115.

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM.

For more information about RIP, see [RIP configuration using CLI](#) on page 3156. For more information about OSPF, see [OSPF configuration using CLI](#) on page 2783.

- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- You must first configure and enable PIM on an IP interface, which can be circuitless, before you can utilize that interface as a candidate rendezvous point (RP). To configure PIM-SM RP for an IP interface, see [Configure a Candidate Rendezvous Point](#) on page 1842.
- Configure one or more bootstrap routers (BSR) to propagate RP information to all switches in the network.

For an IPv6 PIM configuration using the CLI:

- Configure an IPv6 interface.

For more information, see [Configure an IPv6 Interface](#) on page 2194.

- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM.

For more information about RIPng, see [RIPng Configuration using CLI](#) on page 3167. For more information about OSPFv3, see [OSPFv3 Configuration using CLI](#) on page 2827.

- Enable IPv6 PIM-SM globally
- Enable IPv6 PIM-SM on individual interfaces.

Change the Interface Status to Passive

Change the PIM interface status to passive to deny PIM control traffic on the interface.

Before You Begin

- The PIM interface is disabled.

About This Task

The command you use depends on the required administrative state of the interface (enable or disable).

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create a passive interface and enable it simultaneously:

```
ip pim passive
```

3. Create a passive interface in the disabled state:

```
ip pim interface-type passive
```

You must manually enable the interface.

4. Enable a disabled interface:

```
ip pim enable
```

Variable Definitions

Use the data in the following table to use the **ip pim** command.

Variable	Value
<i>active</i>	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control traffic transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.
<i>passive</i>	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control traffic transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.

Change the Interface Status to Active

Change the PIM interface status to active to allow PIM control traffic on the interface.

Before You Begin

- The PIM interface is disabled.

About This Task

The command you use depends on the required administrative state of the interface (enable or disable).

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create an active interface in the disabled state:

```
ip pim interface-type active
```

You must manually enable the interface.

3. Create an active interface and enable it simultaneously using either of the following commands:

- `ip pim active`
- `ip pim enable`



Note

The second command enables an active interface only if this is the first PIM interface you create on the port or VLAN, or you created an active interface in the disabled state. If you already created a passive interface in the disabled state, the second command enables that passive interface.

Variable Definitions

Use the data in the following table to use the `ip pim` command.

Variable	Value
<i>active</i>	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.
<i>passive</i>	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active.

Configure the PIM Virtual Neighbor

Configure a PIM virtual neighbor if the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the PIM virtual neighbor:
`ip pim virtual-neighbor <A.B.C.D> <A.B.C.D>`

Example

Configure the PIM virtual neighbor:

```
Switch:1(config)#ip pim virtual-neighbor 192.0.2.0 198.51.100.0
```

Variable Definitions

Use the definitions in the following table to use the **ip pim virtual-neighbor** command.

Variable	Value
<code>{A.B.C.D} {A.B.C.D}</code>	The first IP address indicates the IP address of the selected interface. The second IP address indicates the IP address of the neighbor.

Configure a Candidate Rendezvous Point

Configure a candidate rendezvous point (C-RP) to serve as backup to the RP router.

About This Task

You can configure only one interface on the switch for multiple groups. You cannot configure multiple interfaces for multiple groups.

With the mask value, you can configure a C-RP router for several groups in one configuration.

For example, if you use a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Add a candidate rendezvous point:
`ip pim rp-candidate group <A.B.C.D> <A.B.C.D> rp <A.B.C.D>`
3. Remove a candidate rendezvous point:
`no ip pim rp-candidate group <A.B.C.D> <A.B.C.D>`

4. Display information about the candidate rendezvous points for the PIM-SM domain:

```
show ip pim rp-candidate
```

Example

Add a candidate rendezvous point:

```
Switch:1(config)#ip pim rp-candidate group 224.1.1.0 255.255.255.0 rp 198.51.100.0
```

Variable Definitions

Use the definitions in the following table to use the **ip pim rp-candidate** command.

Variable	Value
<i>group</i> {A.B.C.D} {A.B.C.D}	Specifies the IP address and the address mask of the multicast group. After the IP address and group mask are combined, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
<i>rp</i> {A.B.C.D}	Specifies the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Configure Static RP

Configure a static RP to ignore the bootstrap router (BSR) mechanism and use the statically configured RPs.

Before You Begin

- Enable PIM-SM globally.

About This Task

Static RP-enabled switches use this feature to communicate with switches from other vendors that do not use the BSR.



Important

You cannot configure a static RP-enabled switch as a BSR or as a C-RP router. All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable static RP:

```
ip pim static-rp
```

The system displays the following message:

```
WARNING: RP information learnt dynamically through BSR functionality will be lost.
Do you wish to enable Static RP? (y/n) ?
```

3. Enter y.

4. Configure a static RP entry:

```
ip pim static-rp {A.B.C.D/X} {A.B.C.D}
```

5. Configure all the switches in the network (including switches from other vendors) to map to the same RP.

6. Display information about the candidate rendezvous points for the PIM-SM domain:

```
show ip pim static-rp
```

Example

Configure a static RP:

```
Switch:1(config)# ip pim static-rp 239.255.0.0/255.255.0.0 198.51.100.0
```

Variable Definitions

Use the definitions in the following table to use the **ip pim static-rp** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and address mask of the multicast group. When combined, the IP address and address mask identify the range of the multicast addresses that the RP handles.
{A.B.C.D}	Specifies the IP address of the static RP.

Configure IPv6 PIM Static RP

On IPv6 PIM BSR mechanism is not supported so static RP must be configured.

Before You Begin

Enable IPv6 PIM-SM globally.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable static RP:

```
ipv6 pim static-rp
```

3. Configure an IPv6 static RP entry:

```
ipv6 pim static-rp WORD<0-255> WORD<0-255>
```

4. Configure all the switches in the network (including switches from other vendors) to map to the same RP.
5. Display information about the candidate rendezvous points for the PIM-SM domain:

```
show ipv6 pim static-rp
```

Variable Definitions

The following table describes the variables for the **ipv6 pim static-rp** command.

Variable	Value
WORD<0-255>	Specifies the IPv6 address and address mask of the multicast group. When combined, the IPv6 address and address mask identify the range of the multicast addresses that the RP handles.
WORD<0-255>	Specifies the IPv6 address of the static RP.

Configure a Candidate BSR on a Port

Configure additional routers as candidate BSRs (C-BSR) to provide backup protection in the event that the primary BSR fails. PIM-SM cannot run without a BSR.

Before You Begin

- Static RP is disabled.

About This Task

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure a candidate BSR:

```
ip pim bsr-candidate preference <0-255>
```

Example

Configure a candidate BSR:

```
Switch:1(config-if)#ip pim bsr-candidate preference 2
```

Variable Definitions

Use the definitions in the following table to use the **ip pim bsr-candidate** command.

Variable	Value
<i>preference</i> <0-255>	Activates the C-BSR on this interface and configures its preference value, from 0-255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR.

Configure a Candidate BSR on a VLAN

Configure additional routers as candidate BSRs (C-BSR) to provide backup protection in the event that the primary BSR fails. PIM-SM cannot run without a BSR.

Before You Begin

- Static RP is disabled.

About This Task

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```



```
configure terminal
```



```
interface vlan <1-4059>
```
2. Configure a candidate BSR on a VLAN:

```
ip pim bsr-candidate preference <0-255>
```

Example

Configure a candidate BSR on a VLAN:

```
Switch:1(config-if)#ip pim bsr-candidate preference 5
```

Variable Definitions

Use the definitions in the following table to use the **ip pim bsr-candidate** command.

Variable	Value
<i>preference</i> <0-255>	Activates the C-BSR on this interface and configures its preference value, from 0-255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR.

Enable Square-SMLT Globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

About This Task



Important

The following command also activates full-mesh configurations.



Note

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable square-SMLT:

```
multicast smlt-square
```

Use PIM Debugging Commands

Use Protocol Independent Multicast (PIM) traces to aid in PIM troubleshooting.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Start debug trace message output:

```
debug ip pim pimdbgtrace
```

3. Stop debug trace message output:
`no debug ip pim pimdbgtrace`

`default debug ip pim pimdbgtrace`
4. Configure the system to display trace messages forwarded by the device:
`debug ip pim send-dbg-trace`
5. Stop the system from displaying trace messages forwarded by the device:
`no debug ip pim send-dbg-trace`

`default debug ip pim send-dbg-trace`
6. Configure the system to display trace messages received by the device:
`debug ip pim rcv-dbg-trace`
7. Stop the system from displaying trace messages received by the device:
`no debug ip pim rcv-dbg-trace`

`default debug ip pim rcv-dbg-trace`
8. Configure the system to display hello messages forwarded or received by the device:
`debug ip pim hello`
9. Stop the system from displaying hello messages forwarded or received by the device:
`no debug ip pim hello`

`default debug ip pim hello`
10. Configure the system to display and log debug trace messages:
`debug ip pim pimdbglog`
11. Stop the system from displaying and logging debug trace messages:
`no debug ip pim pimdbglog`

`default debug ip pim pimdbglog`
12. Configure the system to display register messages forwarded or received by the device:
`debug ip pim register`
13. Stop the system from displaying register messages forwarded or received by the device:
`no debug ip pim register`

`default debug ip pim register`
14. Configure the system to display debug trace messages after an enabled message type, for example, hello or register, is received from a specific sender IP address:
`debug ip pim source {A.B.C.D}`

Variable Definitions

The following table defines parameters for the **debug ip pim** command.

Variable	Value
<i>assert</i>	Displays the assert debug traces. The default is false (disabled).
<i>bstrap</i>	Displays bootstrap debug traces. The default is false (disabled).
<i>group {A.B.C.D}</i>	Displays debug traces from a specific group IP address. The default is 0.0.0.0 (disabled).
<i>hello</i>	Displays hello debug traces. The default is false (disabled).
<i>joinprune</i>	Displays join and prune debug traces. The default is false (disabled).
<i>pimdbglog</i>	Logs debug traces. The default is false (disabled).
<i>pimdbgtrace</i>	Displays PIM debug traces. The default is false (disabled).
<i>rcv-dbg-trace</i>	Displays trace messages received by the switch. The default is false (disabled).
<i>register</i>	If enabled, the system displays register debug traces. The default is false (disabled).
<i>regstop</i>	Displays register stop debug traces. The default is false (disabled).
<i>rp-adv</i>	Displays RP advertisement debug traces. The default is false (disabled).
<i>send-dbg-trace</i>	Displays trace messages forwarded by the switch. The default is false (disabled).
<i>source {A.B.C.D}</i>	Displays debug traces from a specific source IP address. The default is 0.0.0.0 (disabled).

PIM Configuration Using EDM

The switch supports two modes of Protocol Independent Multicast (PIM): Sparse Mode (SM) and Source Specific Multicast (SSM).

- PIM-SM supports multicast groups spread out across large areas of a company or the Internet.
- PIM-SSM optimizes PIM-SM by simplifying the many-to-many model (servers-to-receivers).



Important

The **EnableSpbmConfigMode** boot flag must be disabled before you can configure PIM or IGMP. To verify the configuration, navigate to **Configuration > Edit > Chassis** and select on the **Boot Config** tab.

Before You Begin

For an IPv4 PIM configuration using EDM:

- Configure an IP interface.

For more information, see [IP routing configuration using Enterprise Device Manager](#) on page 2142.

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM.

For more information about RIP, see [RIP configuration using EDM](#) on page 3173. For more information about OSPF, see [OSPF Configuration Using EDM](#) on page 2857.

- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- Configure one or more rendezvous points (RP) for the groups that multicast applications use in the network.



Important

If you configure the rendezvous point (RP) to be the address of a circuitless IP (CLIP) interface, then you must first configure and enable PIM on the CLIP interface before you can utilize that interface as a candidate RP. To configure a PIM-SM RP for a circuitless IP interface, see [Configure a Candidate RP](#) on page 1861.

- Configure one or more bootstrap routers (BSR) to propagate RP information to all switches in the network.

For an IPv6 PIM configuration using EDM:

- Configure an IPv6 interface. For more information, see [Configure an IPv6 Interface](#) on page 2229.
- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM. For more information about RIPng, see [RIPng Configuration using EDM](#) on page 3181. For more information about OSPFv3, see [OSPFv3 Configuration using EDM](#) on page 2896.
- Enable IPv6 PIM-SM globally.
- Enable IPv6 PIM-SM on individual interfaces.

Enable Static RP

Enable static RP to avoid the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.

3. Select the **Globals** tab.
4. Select **sm** (sparse mode).
5. Select **Enable**.
6. Select **Static RP**.
7. Select **Apply**.

The system displays the following message:

```
RP information learnt dynamically through BSR functionality will be
lost. Do you wish to enable Static RP?
```

8. Select **Yes**.

Enable IPv6 Static RP

Use this procedure to enable IPv6 static RP.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Globals** tab.
4. Select **Enable**.
5. Select **sm** (sparse mode).
6. Select **StaticRP**.
7. Select **Apply**.
8. Select **Yes**.

Configure a Static RP

Configure a static RP to ignore the BSR mechanism and use the statically configured RPs only. A static RP-enabled switch uses this feature to communicate with switches from other vendors that do not use the BSR mechanism.

Before You Begin

- Before you can configure a static RP, you must enable the following:
 - PIM-SM
 - static RP

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Static RP** tab.
4. Select **Insert**.
5. Type the required information in each box.
6. Select **Insert**.

Static RP Field Descriptions

Use the descriptions in the following table to use the **Static RP** tab.

Name	Description
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, this value identifies the range of the multicast addresses that the RP handles.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the range of the multicast addresses that the RP handles.
Address	Configures the IP address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid if the switch uses a unicast route to the network for the static RP and is invalid otherwise.

Job Aid

Keep in mind the following configuration considerations:

- Static RPs do not age; they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, you must configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP for certain group range.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of vendor switches across the network, ensure that all switches or routers use the same active RP because vendors use different algorithms to elect the active RP. This switch uses the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.
- Static RP on the switch is active as long as the switch uses a unicast route to the network for the static RP. If the switch loses this route, the static RP is invalidated, and the hash algorithm is invoked to remap all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm is invoked to remap the affected groups.

Configure an IPv6 Static RP Entry

Configure an IPv6 static RP to use the statically configured RPs. A static RP-enabled switch uses this feature to elect the active RP only from the statistically configured switches, without any relation to the RP information of other switches.

Before You Begin

- Before you can configure a static RP, you must enable the following:
 - IPv6 PIM-SM
 - IPv6 static RP

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Static RP** tab.
4. Select **Insert**.
5. Type the required information in each box.
6. Select **Insert**.

Static RP Field Descriptions

Use the descriptions in the following table to use the **Static RP** tab.

Name	Description
GroupAddress	Configures the IPv6 address of the multicast group. When combined with the group mask, this value identifies the range of the multicast addresses that the RP handles.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the range of the multicast addresses that the RP handles.
Address	Configures the global IPv6 address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid if the switch uses a unicast route to the network for the static RP.

Job Aid

Keep in mind the following configuration considerations:

- Static RPs do not age; they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, you must configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP for certain group range.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of vendor switches across the network, ensure that all switches or routers use the same active RP because vendors use different algorithms to elect the active RP. This switch uses the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.
- Static RP on the switch is active as long as the switch uses a unicast route to the network for the static RP. If the switch loses this route, the static RP is invalidated, and the hash algorithm is invoked to remap all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm is invoked to remap the affected groups.

View the Active RP

Perform this procedure to view information about the active RP for all the running multicast groups on the switch.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Active RP** tab.

Active RP Field Descriptions

Use the data in the following table to use the **Active RP** tab.

Name	Description
GroupAddress	Shows the IP address of the multicast group.
Address	Shows the IP address of the RP router. This address must be one of the local PIM-SM enabled interfaces.
Priority	Shows the priority of the RP.

View the IPv6 Active RP

Perform this procedure to view information about the IPv6 active RP for all the running multicast groups on the switch.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Active RP** tab.

Active RP Field Descriptions

Use the data in the following table to use the **Active RP** tab.

Name	Description
GroupAddress	Shows the IPv6 address of the multicast group.
Address	Shows the IPv6 address of the RP router. This address can be one of the local PIM-SM enabled interfaces or any reachable global IPv6 address configured using the static-rp CLI command. Note: IPv6 link local address is always used as the PIM interface address.
Priority	Shows the priority of the RP.

Configure a Candidate Bootstrap Router

Configure routers as candidate bootstrap routers (C-BSR) to provide backup protection in case the primary BSR fails. PIM-SM cannot operate without a BSR. A PIM-SM domain can use only one active BSR.

About This Task

The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **PIM** tab.
5. Select **Enable**.
6. Type the preference in **CBSRPreference**.
The C-BSR with the highest BSR-preference and address becomes the active BSR. The default is -1, which indicates that the current interface is not a C-BSR.
7. Select **Apply**.

View Current BSR Information

View the current BSR information to review the configuration.

Before You Begin

- You must disable static RP.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select **Current BSR** tab.

Current BSR Field Descriptions

Use the descriptions in the following table to use the **Current BSR** tab.

Name	Description
Address	Shows the IP address of the current BSR for the local PIM domain.
FragmentTag	Shows a randomly generated number that distinguishes fragments that belong to different bootstrap messages. Fragments that belong to the same bootstrap message carry the same fragment tag.
HashMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. The hashmask allows a small number of consecutive groups to always hash to the same RP.
Priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
BootStrapTimer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

Change the PIM State on a VLAN Interface

Change the state (active or passive) of PIM on a VLAN interface.

Before You Begin

- Before you change the state of PIM on a VLAN interface, you must first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when the interface receives streams.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Select **IP**.
6. Select the **PIM** tab.
7. Clear the **Enable** check box.
8. Select **Apply**.
9. Select **active** or **passive**.
10. Reenable PIM on the VLAN interface.
11. Select **Apply**.

Edit PIM Interface Parameters

Edit PIM parameters for an interface to customize the PIM configuration.

Before You Begin

- Before you change the state (active or passive) of a PIM interface, first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when the interface receives streams.

Procedure

- In the navigation pane, expand **Configuration > IP**.
- Select **PIM**.
- Select the **Interfaces** tab.
- Configure the required parameters.
- Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	Shows the interface Index. This variable is a read-only field.
Address	Shows the IP address of the PIM interface. This variable is a read-only field.
NetMask	Shows the network mask for the IP address of the PIM interface. This variable is a read-only field.
Mode	Shows the configured mode of this interface. The valid modes are SSM and sparse. This variable is a read-only field.
InterfaceType	Specifies if the interface is active or passive.
DR	Shows the router with the highest IP address on a LAN designated to perform these tasks.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds.
CBSRPreference	Configures the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is -1, which indicates that the current interface is not a C-BSR.
OperState	Indicates the status of PIM on this interface: Up or Down.

Edit IPv6 PIM Interface Parameters

Edit the IPv6 PIM parameters for an interface to customize the IPv6 PIM configuration.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Interfaces** tab.
4. Configure the required parameters
5. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	Shows the interface Index. This variable is a read-only field.
Address	Shows the IPv6 address of the PIM interface. This variable is a read-only field.
NetMask	Shows the network mask for the IPv6 address of the PIM interface. This variable is a read-only field.
Enable	Shows the configured mode of this PIM interface. sparseDense mode is valid only for PIMv1.
Mode	Shows the configured mode of this interface. The valid modes are SSM and sparse. This variable is a read-only field.
DR	Shows the router with the highest IPv6 address on a LAN designated to perform these tasks.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds.
OperState	Indicates the status of PIM on this interface: Up or Down.
Type	Specifies the interface type.

Configure the PIM Virtual Neighbor

Configure a PIM virtual neighbor if the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Virtual Neighbors** tab.
4. Select **Insert**.
5. Specify the IP address of the virtual neighbor.
6. Specify the interface index for the PIM interface.
7. Select **Insert**.

Virtual Neighbors Field Descriptions

Use the descriptions in the following table to use the **Virtual Neighbors** tab.

Name	Description
Address	Specifies the IP address of the neighbor.
IfIndex	Specifies the IP address of the PIM interface.

View PIM-SM Neighbor Parameters

View PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Neighbors** tab.

Neighbors Field Descriptions

Use the descriptions in the following table to use the **Neighbors** tab.

Name	Description
Address	Shows the IP address of the PIM neighbor.
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Shows the time since this neighbor became a neighbor of the local router.
ExpiryTime	Shows the time remaining before the neighbor expires.

View IPv6 PIM-SM Neighbor Parameters

View IPv6 PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Neighbors** tab.

Neighbors Field Descriptions

Use the descriptions in the following table to use the **Neighbors** tab.

Name	Description
Address	Shows the IPv6 address of the PIM neighbor.
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Shows the time since this neighbor became a neighbor of the local router.
ExpiryTime	Shows the time remaining before the neighbor expires.

View IPv6 Neighbor Secondary Address

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **Neighbor Secondary Address** tab.

Neighbor Secondary Address Field Descriptions

Use the descriptions in the following table to use the **Neighbor Secondary Address** tab.

Name	Description
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
Type	Shows the address type of this PIM neighbor.
Primary	The primary IPv6 address of this PIM neighbor.
SecAddress	The secondary IPv6 address of this PIM neighbor.

Viewing RP Set Parameters

View the RP set to see a list of rendezvous point addresses. The BSR constructs this list from C-RP advertisements, and then distributes it to all PIM routers in the PIM domain for the BSR. View the parameters for troubleshooting purposes.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **RP Set** tab.

RP Set Field Descriptions

Use the descriptions in the following table to use the **RP Set** tab.

Name	Description
GroupAddress	Shows the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GroupMask	Shows the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
Address	Shows the IP address of the C-RP router.
HoldTime	Shows the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
ExpiryTime	Shows the time remaining before this C-RP router times out.

Configure a Candidate RP

Configure a C-RP router to add it to the RP Set.

About This Task

You can configure only one interface on a switch for multiple groups.

Use the GroupMask value to configure a candidate RP for several groups in one configuration. For example, if you use a C-RP configuration with a GroupAddress value of 224.0.0.0 and a GroupMask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **PIM**.
3. Select the **Candidate RP** tab.
4. Select **Insert**.
5. Configure the required parameters.
6. Select **Insert**.

Candidate RP Field Descriptions

Use the descriptions in the following table to use the **Candidate RP** tab.

Name	Description
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
InterfaceAddress	Configures the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Enable Square-SMLT Globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

About This Task

**Important**

The following configuration also activates full-mesh configurations.

**Note**

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > IP**
2. Select **Multicast**.
3. Select the **Globals** tab.
4. Select **MulticastSquareSmltEnable**.
Clear this check box if you want to disable square-SMLT globally.
5. Select **Apply**.

View IPv6 RP Set Parameters

View the IPv6 RP set to see a list of rendezvous point addresses and to view the parameters for troubleshooting purposes.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 PIM**.
3. Select the **RP Set** tab.

RP Set *Field Descriptions*

Use the descriptions in the following table to use the **RP Set** tab.

Name	Description
GroupAddress	Specifies the IPv6 address of the multicast group. When combined with the group mask, this value identifies a group prefix for which the address is a static RP.
GroupMask	Specifies the address mask of the multicast group. When combined with the group address, this value identifies a group prefix for which the address is a static RP.
Address	Specifies the IPv6 address of the static RP.
HoldTime	Specifies the hold time of the static RP. The value is 0.
ExpiryTime	Specifies the minimum time remaining before the static RP is down. The value is 0.

View IPv6 Mroute Interface Information

Use the following procedure to view IPv6 Mroute information for an interface.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 Mroute**.
3. Select the **Interfaces** tab.

Interfaces *Field Descriptions*

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	Displays the slot and port number or VLAN ID for this entry.
Ttl	Displays the datagram time-to-live (TTL) threshold for the interface. IPv6 multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means that all multicast packets are forwarded out of the interface.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none"> • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

View IPv6 Mroute Next Hop Information

Use the following procedure to view IPv6 Mroute next hop information.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 Mroute**.
3. Select the **Next Hop** tab.

Next Hop Field Descriptions

Use the data in the following table to use the **Next Hop** tab.

Name	Description
Group	Displays the IPv6 multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
IfIndex	Displays the slot and port number or VLAN ID for this entry.
Address	Displays the address of the next hop specific to this entry. For most interfaces, it is identical to the next-hop group. Non Broadcast Multiple Access (NBMA) interfaces, however, can use multiple next hop addresses out of a single outgoing interface.
State	Displays whether the outgoing interface and next hop represented by this entry currently forward IPv6 datagrams. A value of forwarding indicates the information is currently used; pruned indicates it is not used.
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.

Name	Description
ClosestMemberHops	Displays the minimum number of hops between this router and members of the IPv6 multicast group reached through the next hop on this outgoing interface. IPv6 multicast datagrams for the group that use a time-to-live less than this number of hops are not forwarded to the next hop.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none"> • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

Configure Resource Usage Counter for IPv6 Mroute

Configure the resource usage counters to query the number of ingress and egress IPv6 multicast streams traversing the switch. After you configure the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive notification through a trap on the console, a logged message, or both.



Important

If you do not configure the thresholds, EDM displays only the ingress and egress records that are currently in use.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 Mroute**.
3. Select the **Resource Usage** tab.
4. Configure the ingress and egress thresholds.
5. Configure the notification methods.
6. Select **Apply**.

Resource Usage Field Descriptions

Use the data in the following table to use the **Resource Usage** tab.

Name	Description
Ingress Records In-Use	Displays the number of ingress records (source or group) traversing the switch.
Egress Records In-Use	Displays the number of egress records traversing the switch.
Ingress Threshold	Configures the ingress threshold level (0–32767).
Egress Threshold	Configures the egress threshold level (0–32767).

Name	Description
SendTrapAndLog	Sends both trap and log notification messages after the number of streams exceeds a threshold level.
SendTrapOnly	Sends only trap notification messages after the number of streams exceeds a threshold level. You can configure only one notification type.
LogMsgOnly	Sends only log notification messages after the number of streams exceeds a threshold level.

View IPv6 Multicast Route Information

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 Mroute**.
3. Select the **Route** tab.

Route Field Descriptions

Use the data in the following table to use the **Route** tab.

Name	Description
Group	Displays the IPv6 multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
UpStreamNeighbor	Shows the address of the upstream neighbor from which the IPv6 datagrams from these sources are received.
IfIndex	Displays the slot and port number or VLAN ID for this entry.
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none"> • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

IGMP Configuration Using the CLI

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.



Important

The **spbm-config-mode** boot flag must be disabled before you can configure PIM or IGMP. To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

Before You Begin

- Complete one of the following tasks:
 - Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
 - Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Protocol Independent Multicast-Sparse Mode (PIM-SM) or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).



Important

To configure and use IGMP on a VRF instance you must first select and launch the VRF context.

To select and launch the VRF context, see [Configure IGMP on a VRF](#) on page 1794.

Configure Multicast Stream Limitation on an Ethernet Port

Configure multicast stream limitation on an Ethernet port to limit the number of concurrent multicast streams on the port. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

About This Task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the port drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.



Note

Configuration of multicast stream limitation is not supported on a node configured as the DvR Leaf within a DvR domain.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable multicast stream limitation and configure the maximum number of allowed streams:

```
ip igmp stream-limit stream-limit-max-streams <0-65535>
```

3. If stream-limit is already enabled on the interface, change the maximum number of allowed streams:

```
ip igmp stream-limit stream-limit-max-streams <0-65535>
```

4. Display multicast stream limitation information for the ports on a specific interface:

```
show ip igmp stream-limit interface
```

Example

Enable multicast stream limitation on the Ethernet port and configure the maximum number of allowed streams to 8.

```
Switch:1(config-if)# ip igmp stream-limit  
Switch:1(config-if)# ip igmp stream-limit stream-limit-max-streams 8
```

Variable Definitions

Use the data in the following table to use the **ip igmp stream-limit-max-streams** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this port. The range is from 0–65535 and the default is 4.

Configure Multicast Stream Limitation on a VLAN

Configure multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

About This Task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the VLAN drops joins to new streams. A service provider

uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

**Note**

Configuration of multicast stream limitation is not supported on a node configured as the DvR Leaf within a DvR domain.

Procedure

1. Enter VLAN Interface Configuration mode:
`enable`

`configure terminal`

`interface vlan <1-4059>`
2. Enable multicast stream limitation and configure the maximum number of allowed streams:
`ip igmp stream-limit stream-limit-max-streams <0-65535>`
3. If stream-limit is already enabled on the VLAN, change the maximum number of allowed streams:
`ip igmp stream-limit stream-limit-max-streams <0-65535>`
4. Display multicast stream limitation information for the ports on a specific interface:
`show ip igmp stream-limit port`

Example

Enable multicast stream limitation and configure the maximum number of allowed streams to 8.

```
Switch:1(config-if)# ip igmp stream-limit
Switch:1(config-if)# ip igmp stream-limit stream-limit-max-streams 8
```

Variable Definitions

Use the data in the following table to use the **ip igmp stream-limit** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this VLAN. The range is from 0–65535 and the default is 4.

Configure VLAN Multicast Stream Limitation Members

Configure multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```

2. Configure multicast stream limitation members on a VLAN:

```
ip igmp stream-limit-group {slot/port[/sub-port] [-slot/port[/sub-
port]][,...]} enable max-streams <0-65535>
```

Example

Enable multicast stream limitation on ports 2/1 to 2/3, and configure the maximum allowed number of streams to 6 for this interface.

```
Switch:1(config-if)#ip igmp stream-limit-group 2/1-2/3 max-streams 6
```

Variable Definitions

Use the data in the following table to use the **ip igmp stream-limit-group** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams for the specified ports on this VLAN. The range is from 0-65535 and the default is 4.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure Multicast Router Discovery Options

Configure the multicast router discovery options to enable the automatic discovery of multicast-capable routers.

About This Task



Important

The switch does not support the Multicast Router Discovery (MRDISC) protocol on brouter ports.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable multicast router discovery:

```
ip igmp mrdisc
```

3. Configure the maximum advertisement intervals between successive advertisements:

```
ip igmp mrdisc maxadvertinterval <2-180> maxinitadvertinterval <2-180>
```

4. Configure the maximum advertisements after initialization:

```
ip igmp mrdisc maxinitadvertisements <2-15>
```

5. Configure the minimum advertisement interval between successive advertisements:

```
ip igmp mrdisc minadvertinterval <3-180>
```

6. Configure the time allowed before a neighbor is declared dead:

```
ip igmp mrdisc neighdeadinterval <2-180>
```

Example

Configure the maximum advertisement intervals between successive advertisements:

```
Switch:1(config-if)#ip igmp mrdisc maxadvertinterval 30 maxinitadvertinterval 5
```

Configure the maximum advertisements after initialization:

```
Switch:1(config-if)#ip igmp mrdisc maxinitadvertisements 8
```

Configure the minimum advertisement interval between successive advertisements:

```
Switch:1(config-if)#ip igmp mrdisc minadvertinterval 30
```

Configure the time allowed before a neighbor is declared dead:

```
Switch:1(config-if)#ip igmp mrdisc neighdeadinterval 60
```

Variable Definitions

Use the data in the following table to use the **ip igmp mrdisc** command.

Variable	Value
<code>maxadvertinterval <2-180></code>	Configures the maximum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration, and then reset the switch. The default is 20.
<code>maxinitadvertinterval <2-180></code>	Configures the maximum number (in seconds) between successive initial advertisements. For this change to take effect, you must save the configuration, and then reset the switch. The default is 2.
<code>maxinitadvertisements <2-15></code>	Configures the maximum number of initial multicast advertisements after initialization. For this change to take effect, you must save the configuration, and then reset the switch. The default is 3.
<code>minadvertinterval <3-180></code>	Configures the minimum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration, and then reset the switch. The default is 15.
<code>neighdeadinterval <2-180></code>	Configures the multicast router discovery dead interval—the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down. The default is 60.

Configure Explicit Host Tracking

Configure explicit host tracking to track all the source and group members.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure explicit host tracking:

```
ip igmp igmpv3-explicit-host-tracking
```

3. Display all the tracked members for a specific group:

```
show ip igmp group group <A.B.C.D> tracked-members [member-subnet  
<A.B.C.D/X>] [source-subnet <A.B.C.D/X>] [port {{slot/port[/sub-port]  
[-slot/port[/sub-port]][,...][slot/all][all]]} [vlan <1-4059>]
```

4. Display the IGMPv3 specific data:

```
show ip igmp group group <A.B.C.D> detail port {{slot/port[/sub-port]  
[-slot/port[/sub-port]][,...][slot/all][all]]} vlan <1-4059>
```

Examples

Configure explicit host tracking:

```
Switch:1(config-if)#ip igmp igmpv3-explicit-host-tracking
```

Display all the tracked members:

```
Switch:1#show ip igmp group
```

```
=====
                        Igmp Group - GlobalRouter
=====
```

GRPADDR	INPORT	MEMBER	EXPIRATION	TYPE	L2ISID
224.5.2.1	V701-1/4	62.0.1.1	214	Dynamic	40400
224.5.2.2	V702-1/4	62.0.2.1	221	Dynamic	40400
224.5.2.3	V703-1/4	62.0.3.1	217	Dynamic	40400
224.5.2.4	V704-1/4	62.0.4.1	223	Dynamic	40400

```
4 out of 4 group Receivers displayed
```

```
Total number of unique groups 2
```

Display all the tracked members for a specific group:

```
Switch:1(config-if)#show ip igmp group group 232.1.1.1 tracked-members
```

```
=====
                        Members of Channels/Groups - GlobalRouter
=====
```

INTERFACE	CHANNEL/GROUP	MEMBER	MEMBER MODE	EXP
Vlan333-1/30	*/232.1.1.1	133.133.133.200	IS_EXCLUDE	205

Note:

The "*" attached to the interface (if any) indicates that the interface has explicit host tracking disabled.

Display IGMPv3 specific data:

```
Switch:1(config-if)#show ip igmp group group 232.32.32.10 detail
```

```
=====
                        Igmp Group Detail - GlobalRouter
=====
```

Interface:	Vlan222-1/10
IGMPv3 Group:	232.32.32.10
Interface Group Mode:	INCLUDE
Interface Compatibility Mode:	IGMP_V3
V2 Host Timer:	Not Running
V1 Host Timer:	Not Running

```
Interface Group Include Source List:
Source Address Expires
133.133.133.200 114
```

Variable Definitions

Use the data in the following table to use the **ip igmp igmpv3-explicit-host-tracking** command.

Variable	Value
<i>explicit-host-tracking</i>	Enables explicit host tracking on IGMPv3. The default state is disable.
<i><A.B.C.D></i>	Specifies the IP address of the group of the tracked member.

Configure IGMP Static Members

Configure IGMP static members to add members to a snoop group. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams are always forwarded to the multicast router within the VLAN, in addition to the ports in this static entry.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Configure interface static members:

```
ip igmp static-group {A.B.C.D} {A.B.C.D} {port {slot/port[-slot/port]
[,...][slot/all][all]}} [static|blocked]
```

Example

Configure interface static members:

```
Switch:1(config-if)#ip igmp static-group 239.1.1.1 239.1.2.1 port 2/1 static
```

Variable Definitions

Use the data in the following table to use the **ip igmp static-group** command.

Variable	Value
<i>{A.B.C.D} {A.B.C.D}</i>	Indicates the IP address range of the selected multicast group.
<i>port</i>	Adds ports to a static group entry

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Creates a static group entry. Specifies the port or list of ports that is a member of the VLAN interface being configured to which you want to redirect the multicast stream for this multicast group.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code><static blocked></code>	Configures the route to static or blocked.

Configure SSM Dynamic Learning and Range Group

Configure SSM dynamic learning and a range group to enable the IGMPv3 dynamic learning feature and to extend the default SSM range of 232/8 to include an IP multicast address. As new SSM channels are learned, the system displays them in the SSM channel table.

Before You Begin

- To define the range group, you must first disable PIM.

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable SSM dynamic learning:

```
ip igmp ssm dynamic-learning
```

3. Configure the range group:

```
ip igmp ssm group-range <A.B.C.D/X>
```

The system displays the following message:

```
Warning: Changing the SSM range will cause all spb-multicast and spb-
pim-gw enabled interfaces to be internally bounced. Do you wish to
continue? (y/n) ? (y/n)?
```

Enter **y** to continue.

Example

Define the SSM range group address (234.0.0.0) and mask (255.0.0.0). Enable dynamic learning from IGMPv3 reports.

```
Switch:1(config)#ip igmp ssm group-range 234.0.0.0/255.0.0.0
```

```
WARNING: Changing the SSM range will cause all spb-multicast and spb-pim-gw enabled
interfaces to be internally bounced. Do you wish to continue? (y/n) ? (y/n)? y
```

```
Switch:1(config)#ip igmp ssm dynamic-learning
```

Variable Definitions

Use the data in the following table to use the **ip igmp ssm** command.

Variable	Value
{A.B.C.D/X}	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Change the SSM Range Group

Change the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address.

Before You Begin

Before you disable or delete an ssm-map, always send IGMPv1 or IGMPv2 leave messages from hosts that operate in IGMPv1 or IGMPv2. If you do not perform this action, receiving and processing reports in SSM range on an IGMP interface enabled with IGMPv1 or IGMPv2 can lead to unexpected behavior.

About This Task



Important

This procedure reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), it also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable PIM:

```
no ip pim enable
```

If you forget to disable PIM, the system displays the following error message:

```
Error: PIM is enabled in SSM mode, disable PIM
```

3. Delete each entry in the SSM channel table:

```
no ip igmp ssm-map [all] [{A.B.C.D} enable]
```

If you forget to delete the SSM channels, the system displays the following error message:

```
Error: SSM source group table not empty
```

4. Configure the new IP multicast group address:

```
ip igmp ssm group-range {A.B.C.D/X}
```

The system displays the following message:

```
Warning: Changing the SSM range will cause all spb-multicast and spb-  
pim-gw enabled interfaces to be internally bounced. Do you wish to  
continue? (y/n) ? (y/n)?
```

Enter **y** to continue.

5. Enable PIM:

```
ip pim enable
```

Example

Configure the new IP multicast group address:

```
Switch:1(config)#ip igmp ssm group-range 232.0.0.0/16
```

```
WARNING: Changing the SSM range will cause all spb-multicast and spb-pim-gw enabled  
interfaces to be internally bounced. Do you wish to continue? (y/n) ? (y/n)? y
```

Variable Definitions

Use the data in the following table to use the **ip igmp ssm group-range** and **ip igmp ssm** commands.

Variable	Value
<code>{A.B.C.D/X}</code>	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Configure the SSM Map Table

Configure the SSM map table to map groups to their sending source. SSM maps cannot conflict with static source groups. After you configure an SSM map or a static source group, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) to different sources or multiple sources to the same group for both static source group and an SSM map.

About This Task

The consistency check applies to all SSM map entries, even if they are disabled. If you disable an entry, it becomes inactive. If you do not delete the entry, you can reenabling it later.

After you disable an SSM map, the switch stops multicast traffic from the specified source to the specified group. You can use this static configuration as a security feature to block traffic from a certain source to a specific group.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable the SSM map table for all static entries:

```
ip igmp ssm-map all
```
3. Create a static entry for a specific group:

```
ip igmp ssm-map {A.B.C.D} {A.B.C.D} enable
```

Example

Create an SSM map table entry for the multicast group 234.0.1.0 and the source at 192.32.99.151. Configure the administrative state to enable all the static SSM map table entries.

```
Switch:1(config)#ip igmp ssm-map 234.0.1.0 192.32.99.151
Switch:1(config)#ip igmp ssm-map all
```

Variable Definitions

Use the data in the following table to use the **ip igmp ssm-map** command.

Variable	Value
<code>{A.B.C.D} {A.B.C.D}</code>	Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that sends traffic to the group.
<code>{A.B.C.D} enable</code>	Enables the administrative state for a specific entry (group). This variable does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.

Configure Multicast Access Control for an IGMP Ethernet Port

Configure multicast access control for an IGMP Ethernet port to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure multicast access control:

```
ip igmp access-list WORD<1-64> {A.B.C.D/X} <deny-tx|deny-rx|deny-both|
allow-only-tx|allow-only-rx|allow-only-both>
```

3. Change an existing access list:

```
ip igmp access-list WORD<1-64>> {A.B.C.D/X} mode <deny-tx|deny-rx|
deny-both|allow-only-tx|allow-only-rx|allow-only-both>
```

Variable Definitions

Use the data in the following table to use the **ip igmp access-list** command

Variable	Value
<i>{A.B.C.D/X}</i>	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
<i>deny-tx deny-rx deny-both allow-only-tx allow-only-rx allow-only-both</i>	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
<i>mode</i>	Changes the access control group configuration.
<i>WORD<1-64></i>	Specifies the name of the access list from 1-64 characters.

Configure Multicast Access Control for a VLAN

Configure multicast access control for an IGMP VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Configure multicast access control:

```
ip igmp access-list WORD<1-64> [A.B.C.D/X] <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>
```

3. Change an existing access list:

```
ip igmp access-list WORD<1-64> [A.B.C.D/X] mode <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>
```

Variable Definitions

Use the data in the following table to use the **ip igmp access-list** command.

Variable	Value
<i>{A.B.C.D/X}</i>	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
<i>deny-tx deny-rx deny-both allow-only-tx allow-only-rx allow-only-both</i>	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
<i>mode</i>	Changes the access control group configuration.
<i>WORD<1-64></i>	Specifies the name of the access list from 1-64 characters.

Configure Fast Leave Mode

Configure fast (immediate) leave mode to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces. Normal IGMP behavior is skipped. Fast leave mode provides one command that controls all IGMP fast leave enabled interfaces.

Before You Begin

- You must enable explicit-host-tracking before configuring fast-leave mode for IGMPv3. For more information on enabling explicit-host-tracking, see [Configure Explicit Host Tracking](#) on page 1872.

About This Task

If a single user connects to an interface, you do not need to track if other users exist on the interface to perform the fast leave. In cases like this, you must change the mode to one-user.



Important

Fast leave mode applies only to fast leave enabled IGMP interfaces.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. View the current fast leave mode:

```
show ip igmp sys
```



Note

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

3. Configure fast leave mode:

```
ip igmp immediate-leave-mode <multiple-user|one-user>
```

Example

Change the mode to one-user.

```
Switch:1(config)#ip igmp immediate-leave-mode one-user
```

Variable Definitions

Use the data in the following table to use the **ip igmp immediate-leave-mode** command.

Variable	Value
<i>multiple-user one-user</i>	<p><i>multiple-user</i> removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This configuration is the default.</p> <p><i>one-user</i> removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.</p>

Enable Fast Leave Mode on a Port

Enable fast (immediate) leave mode to specify if a port receives a leave message from a member of a group. If you enable fast leave mode on a port, it uses the global fast leave mode configuration.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable fast leave:

```
ip igmp immediate-leave
```

Configure IGMP Fast Leave Members on a VLAN

Configure IGMP fast leave members on a VLAN to specify fast leave capable ports.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable fast leave on the VLAN:

```
ip igmp immediate-leave
```

3. Configure fast leave members on a VLAN:

```
ip igmp immediate-leave-members {slot/port[/sub-port] [-slot/port[/sub-  
port]][,...]}
```

Variable Definitions

Use the data in the following table to use the **ip igmp immediate-leave-members** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Enable IGMP Layer 2 Querier

When no multicast router exists in your network, you can use IGMP Layer 2 Querier to allow the Layer 2 switch to act as a multicast router so that the system can participate in multicast environments where multicast routing is not required.

Before You Begin

- You must enable IGMP snooping.

About This Task

When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups.

By default, IGMP Layer 2 Querier is disabled.

Enable Layer 2 Querier on only one node in the VLAN.

On Shortest Path Bridging (SPB) Customer VLANs (CVLAN), IGMP Querier is enabled automatically when you enable snooping on the VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable IGMP Layer 2 Querier:

```
ip igmp snoop-querier
```

What to Do Next

You must enable the IGMP Layer 2 Querier address. See [EnablingIGMPLayer2QuerierAddress](#)

Enable IGMP Layer 2 Querier Address

To use the IGMP Layer 2 Querier feature you must designate the IGMP Layer 2 Querier source IP address, the address the system uses in the query message.

Before You Begin

- Enable IGMP Layer 2 Querier.

About This Task

You must configure the IGMP Layer 2 Querier address to an IP address in the IP subnet that IGMP hosts, and to which IGMP snoopers in the VLAN belong.

The default IP address is 0.0.0.0 when the IGMP Layer 2 Querier is disabled.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```

2. Enable the IGMP Layer 2 Querier address:

```
ip igmp snoop-querier-addr {A.B.C.D}
```

3. Verify the configuration:

```
show ip igmp snooping [vrf WORD<1-16>] [vrfsids WORD<0-512>]
```

Example

Enable the IGMP Layer 2 Querier feature for VLAN 4, and configure the querier address. Verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 4
Switch:1(config-if)#ip igmp snoop-querier
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
Switch:1#show ip igmp snooping

=====
                        Igmp Snooping - GlobalRouter
=====
IFINDEX  SNOOP   PROXY  SSM    STATIC      ACTIVE      MROUTER
          ENABLE SNOOP  SNOOP  MROUTER      MROUTER      EXPIRATION
```

	ENABLE	ENABLE	PORTS	PORTS	TIME
V2	false	false	false		0
V3	false	false	false		0
V4	true	false	false		0
V200	false	false	false		0
IFINDEX	SNOOP	CONFIGURED	DYNAMIC	COMPATIBILITY	
	QUERIER	QUERIER	DOWNGRADE	MODE	
	ENABLE	ADDRESS	VERSION		
V2	false	0.0.0.0	enable	disable	
V3	false	0.0.0.0	enable	disable	
V4	true	192.0.2.1	enable	disable	
V200	false	0.0.0.0	enable	disable	
4 out of 4 entries displayed					

Querier Not Elected

If a Querier is not elected, use the following procedure to troubleshoot the issue.

Procedure

1. Enter Privileged EXEC mode:
enable
2. As the IGMP Layer 2 Querier is based on IGMP snoop, check whether IGMP snoop is enabled on the VLAN:
show ip igmp interface vlan

If IGMP snoop is disabled, the Layer 2 Querier cannot work until IGMP snoop and IGMP Layer 2 Querier are reenabled.

Example

Check whether IGMP snoop is enabled on the VLAN:

```
Switch:1>enable
Switch:1#show ip igmp interface vlan
```

Vlan Ip Igmp												
VLAN ID	QUERY INTVL	QUERY MAX	ROBUST LEAVE	VERSION	LAST MEMB	PROXY SNOOP	SNOOP ENABLE	SNOOP ORIGIN	SSM SNOOP	UPnP FILTER	FAST LEAVE	FAST
		RESP			QUERY	ENABLE			ENABLE	Enable	ENABLE	
PORTS												
1	125	100	2	2	10	false	false	RADIUS	false	false	false	false
2	125	100	2	2	10	false	false	RADIUS	false	false	false	false
3	125	100	2	2	10	false	false	RADIUS	false	false	false	false
4	125	100	2	2	10	false	false	RADIUS	false	false	false	false
5	125	100	2	2	10	false	false	RADIUS	false	false	false	false
10	125	100	2	2	10	false	false	RADIUS	false	false	false	false
100	125	100	2	2	10	false	false	RADIUS	false	false	false	false
200	125	100	2	2	10	false	false	RADIUS	false	false	false	false
300	125	100	2	2	10	false	false	RADIUS	false	false	false	false
444	125	100	2	2	10	false	false	RADIUS	false	false	false	false

All 10 out of 10 Total Num of Igmp entries displayed

VLAN ID	SNOOP QUERIER ENABLE	SNOOP QUERIER ADDRESS	DYNAMIC DOWNGRADE VERSION	COMPATIBILITY MODE	EXPLICIT HOST TRACKING	UPnP FILTER ADDRESS
1	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
2	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
3	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
4	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
5	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
10	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
100	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
200	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
300	false	0.0.0.0	enable	disable	disable	239.255.255.250/32
444	false	0.0.0.0	enable	disable	disable	239.255.255.250/32

All 10 out of 10 Total Num of Igmp entries displayed

Enable Trace Messages for IGMP Layer 2 Querier Troubleshooting

If the preceding information does not address your issue, you can also use the following trace command to view additional information related to Layer 2 querier.



Caution

Using the trace tool inappropriately can cause CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation. If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Use the following trace command to begin the trace operation for additional information related to Layer 2 querier:
trace level 23 <1-4>
3. Stop tracing:
trace shutdown
4. View the trace results:
trace screen enable
5. View trace saved to a file:
show trace file [tail]

Variable Definitions

Use the data in the following table to use the **trace** command.

Variable	Value
<code>level</code> <code>[<Module_ID>]</code> <code>[<1-4>]</code>	Starts the trace by specifying the module ID and level. Module ID 23 represents the IGMP module <Module_ID> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see <i>Fabric Engine and VOSS Feature Matrix</i> . To see the list of module IDs, use the show trace modid-list command or CLI command completion Help. <0-4> specifies the trace level: <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
<code>shutdown</code>	Stops the trace operation.
<code>screen</code> <code>{disable enable}</code>	Enables or disables the display of trace output to the screen. Important: Avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

Use the data in the following table to use the **show trace** command.

Variable	Value
<code>file [tail]</code>	Displays the trace results saved to a file.
<code>level</code>	Displays the current trace level for all modules.
<code>modid-list</code>	Specifies the module ID list.

Troubleshoot IGMPv3 Backwards Compatibility

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv2 and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message. If it is a v2 message, IGMP snoop processes handle the message.

To troubleshoot issues with the IGMPv3 backwards compatibility feature, perform the following procedure.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Verify that the SSM static channel is configured for the v1/v2 joins received. Display the configured SSM static channels:
`show ip igmp ssm-map`

3. Verify that the SSM group range is configured for the v1/v2 joins received. Display the configured SSM group range:

```
show ip igmp ssm
```

Example

Display the configured SSM static channels and display the configured SSM group range:

```
Switch:1#show ip igmp ssm-map
```

Igmp Ssm Channel - GlobalRouter				
GROUP	SOURCE	MODE	ACTIVE	STATUS
233.252.0.1	192.0.2.200	dynamic	false	enabled
233.252.0.2	192.0.2.200	dynamic	false	enabled
233.252.0.3	192.0.2.200	dynamic	false	enabled
233.252.0.4	192.0.2.200	dynamic	false	enabled
233.252.0.5	192.0.2.200	dynamic	false	enabled
233.252.0.6	192.0.2.200	dynamic	false	enabled
233.252.0.7	192.0.2.200	dynamic	false	enabled
233.252.0.8	192.0.2.200	dynamic	false	enabled
233.252.0.9	192.0.2.200	dynamic	false	enabled
233.252.0.10	192.0.2.200	dynamic	false	enabled

10 out of 10 entries displayed

```
Switch:1#show ip igmp ssm
```

Igmp Ssm Global - GlobalRouter	
DYNAMIC LEARNING	SSM GROUP RANGE
enable	232.0.0.0/255.0.0.0

View IGMP Interface Information

Perform this procedure to view the IGMP interface table.

About This Task

If an interface does not use an IP address, the system does not display it in the IGMP table. One exception is an IGMP snooping interface, which does not require an interface IP address.

If an interface uses an IP address, but neither IGMP snoop or PIM is enabled, the system displays the interface as inactive in the Status field.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View IGMP interfaces:

```
show ip igmp interface [gigabitethernet {slot/port[-slot/port]}
[,...][slot/all][all]][,...]|vlan <1-4059>] [vrf WORD<1-16>][vrfids
WORD<0-512>]
```

Example

View IGMP interfaces:

```
Switch:1#show ip igmp interface vrf l3vsn-100
```

```
=====
                        Igmp Interface - VRF l3vsn-100
=====
```

IF	QUERY INTVL	STATUS	VERS.	OPER VERS	QUERIER	QUERY MAXRSPT	WRONG QUERY	JOINS	ROBUST	LASTMEM QUERY	MODE	L2ISID	SCOPE ISID
V10	125	inact	2	2	0.0.0.0	100	0	0	2	10	routed-spb	0	100

Variable Definitions

The following table defines parameters for the **show ip igmp interface** command.

Variable	Value
<i>gigabitethernet</i> <i>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <p>If you do not specify a slot and port, the command output includes all IGMP interfaces.</p>
<i>vlan</i> <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <p>If you do not specify a VLAN ID, the command output includes all IGMP interfaces.</p>
<i>vrf</i> <i>WORD</i> <1-16>	Optionally, identifies the VRF name. If you do not specify a VRF name, the results display information for the Global Router. If you specify a VRF name, the results display information only for the VRF you specify.
<i>vrfids</i> <i>WORD</i> <0-512>	Optionally, identifies the VRF ID. If you do not specify a range of VRF IDs, the results display information for the Global Router. If you specify a VRF ID or range of VRF IDs, the results display information only for the VRF you specify.

View Multicast Group Trace Information for IGMP Snoop

About This Task

Multicast group trace tracks the data flow path of the multicast streams.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the multicast group trace for an IGMP snoop-enabled interface:
show ip igmp snoop-trace [source {A.B.C.D}] [group {A.B.C.D}]

Example

Display the multicast group trace for an IGMP snoop-enabled interface:

```
Switch:1# show ip igmp snoop-trace
```

```
=====
                        Snoop Trace - GlobalRouter
=====
```

GROUP ADDRESS	SOURCE ADDRESS	IN VLAN	IN PORT	OUT VLAN	OUT PORT	TYPE
233.252.0.1	192.0.2.6	500	1/1	500	1/5	ACCESS
233.252.0.10	192.0.2.7	500	1/1	500	1/10	ACCESS

Variable Definitions

The following table defines parameters for the **show ip igmp snoop-trace** command.

Variable	Value
<i>group {A.B.C.D}</i>	Specifies the group IP address in the format a.b.c.d.
<i>source {A.B.C.D}</i>	Specifies the source IP address in the format a.b.c.d.

View IGMP Group Information

View information about IGMP groups to see the current group operation on the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View IGMP group information:
show ip igmp group group <A.B.C.D> detail [port {slot/port[/sub-port]}/-slot/port[/sub-port]][,...]] [vrf WORD <1-16>] [vrffids WORD <0-512>]

show ip igmp group group <A.B.C.D> tracked-members [member-subnet <A.B.C.D./X>] [port {slot/port[/sub-port]}/-slot/port[/sub-port]] [,...]] [source-subnet <A.B.C.D/X>] [vrf WORD <1-16>] [vrffids WORD <0-512>]

Example

View IGMP group information:

```
Switch:1>enable
Switch:1#show ip igmp group group 232.0.0.0
```

```

=====
                        Igmp Group - GlobalRouter
=====
GRPADDR      INPORT      MEMBER      EXPIRATION TYPE
-----
232.0.0.0    V1015-1/2    200.0.15.53  258        Dynamic

1 out of 271 group Receivers displayed

Total number of unique groups 271

```

Variable Definitions

The following table defines parameters for the **show ip igmp group** command.

Variable	Value
<i>count</i>	Displays the number of entries in the IGMP group.
<i>group</i> <A.B.C.D>	Specifies the address of the IGMP group.
<i>member-subnet</i> {default <A.B.C.D>}]	Specifies the IP address and mask of the IGMP member.
<i>vrf</i> WORD<1-16>	Specifies the VRF name.
<i>vrfids</i> WORD<0-512>	Specifies the VRF ID.

The following table defines parameters for the **show ip igmp group group** command.

Variable	Value
<code>detail</code> <code>[port {slot/port[/sub-port]</code> <code>[-slot/port[/sub-port]][,...]</code> <code>[slot/all][all]} vlan <1-4059> </code> <code>vrfWORD <1-16> vrfidsWORD <0-</code> <code>255>]</code>	<p>Use the <i>detail</i> parameter to show IGMPv3-specific data.</p> <p>For data related to a specific interface use the following:</p> <ul style="list-style-type: none"> <code>port{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> — Specifies the port list. <code>vlan <1-4059></code> — Specifies the VLAN. <p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <ul style="list-style-type: none"> <code>vrf WORD<1-16></code> <p>— Specifies the VRF name.</p> <ul style="list-style-type: none"> <code>vrfids WORD<0-255></code> — Specifies the VRF ID.
<i>tracked-members</i>	Use the <i>tracked-members</i> parameter to view all the tracked members for a specific group.
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

The following table defines parameters for the **show ip igmp group group <A.B.C.D> tracked-members** command.

Variable	Value
<code>member-subnet {default <A.B.C.D>}]</code>	Specifies the IP address and mask of the IGMP member.
<code>port {slot/port[/sub-port]</code> <code>[-slot/port[/sub-port]][,...]</code> <code>[slot/all][all]}</code>	Specifies the port list.
<code>source-subnet <A.B.C.D/X></code>	Specifies the source IP address and the subnet mask.

Variable	Value
<code>vlan <1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

Determine the Protocol Configured on the Added VLAN

Use this procedure to determine the protocol configured on the added VLAN.

The protocol configured on the added VLAN can be one of the following values:

- snoop
- snoop-spb
- route-spb
- pim

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Determine the protocol configured on the added VLAN:

```
show ip igmp interface [gigabitethernet {slot/port[-slot/port]}[,...]  
[slot/all][all]][vlan <1-4059>][vrf WORD<1-16>][vrfids WORD<0-512>]
```

The protocol displays under the Mode column of the command output.

Example

Determine the protocol configured on the added VLAN:

```
Switch:1#show ip igmp interface vrf l3vsn-100
```

```
=====
                        Igmp Interface - VRF l3vsn-100
=====
```

IF	QUERY INTVL	STATUS	VERS.	OPER VERS	QUERIER	QUERY MAXRSPT	WRONG QUERY	JOINS	ROBUST	LASTMEM QUERY	MODE	L2ISID	SCOPE ISID
V10	125	inact	2	2	0.0.0.0	100	0	0	2	10	routed-spb	0	100

Variable Definitions

The following table defines parameters for the **show ip igmp interface** command.

Variable	Value
<code>gigabitethernet{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>vlan<1-4059></code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p>
<code>vrfWORD<1-16></code>	Specifies the VRF instance by the VRF name.
<code>vrfidsWORD<0-512></code>	Specifies the VRF ID for which to display statistics.

IGMP Configuration Using EDM

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.



Important

The **EnableSpbmConfigMode** boot flag must be disabled before you can configure PIM or IGMP. To verify the configuration, navigate to **Configuration > Edit > Chassis** and select on the **Boot Config** tab.

Before You Begin

- Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
- Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Protocol Independent Multicast-Sparse Mode (PIM-SM), or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).



Important

To configure and use IGMP on a VRF instance you must first select and launch the VRF context.

To select and launch the VRF context, see [Select and Launch a VRF Context View](#) on page 4100.

Enable IGMP snoop on a VLAN

Enable IGMP snooping on a VLAN to optimize the multicast data flow for a group within a VLAN to only those that are members of the group that uses IGMP snoop.

About This Task

The switch listens to group reports from each port and builds a database of multicast group members for each port. The switch suppresses the reports heard by not forwarding them to other hosts, forcing the members to continuously send their own reports.

The switch relays group membership from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN. The switch multicasts data only to the participating group members and to the multicast routers within the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **IGMP** tab.
7. Select the **SnoopEnable** check box.
8. Select the **ProxySnoopEnable** check box.
9. In **StreamLimitEnable**, select **enable**.
10. Select **Apply**.

Configure IGMP Interface Static Members

Configure IGMP interface static members to add members to a snoop group.

About This Task

You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams always forward to the multicast router within the VLAN, in addition to the ports in this static entry.



Important

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Static** tab.
4. Select **Insert**.
5. Configure the required parameters.
6. Select **Insert**.

Static Field Descriptions

Use the data in the following table to use the **Static** tab.

Name	Description
IfIndex	Shows the interface where the IGMP entry is enabled.
GrpAddr	Indicates the start of the IP multicast address range of the multicast stream. Within the indicated valid range (224.0.0.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.
ToGrpAddr	Indicates the end of the IP multicast address range of the multicast stream. If an address is not entered, the IP address in the GrpAddr field is the single address.
MemberPorts	Specifies the ports to which you want to redirect the multicast stream for this multicast group. The ports must be member ports of the VLAN.
NotAllowedToJoin	Specifies the ports that do not receive the multicast stream for this multicast group.

Configure the SSM Map Table

Configure the SSM map table to map groups to their sending source. SSM maps cannot conflict with static source groups. After you configure an SSM map or a static source group, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) or multiple groups to different sources for both static source group and an SSM channel.

Before You Begin

Before you disable or delete an ssm-map, always send IGMPv1 or IGMPv2 leave messages from hosts that operate in IGMPv1 or IGMPv2. If you do not perform this action, receiving and processing reports in SSM range on an IGMP interface enabled with IGMPv1 or IGMPv2 can lead to unexpected behavior.

About This Task

The consistency check applies to all SSM channel entries, even if they are disabled. If you disable an entry, it becomes inactive. If you do not delete the entry, you can reenable it later.

After you disable an SSM map, the switch stops multicast traffic from the specified source to the specified group. You can use this static configuration as a security feature to block traffic from a certain source to a specific group.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Ssm Map** tab.
4. Select **Insert**.
5. Type the IP address for the multicast group and source.
6. Select **Insert**.

You can change the default status of an SSM map from enable to disable in **AdminState**.

Ssm Map Field Descriptions

Use the data in the following table to use the **Ssm Map** tab.

Name	Description
IpMulticastGrp	Specifies an IP multicast address that is within the SSM range.
IpSource	Specifies the IP address of the source that sends traffic to the group.
LearningMode	Displays whether the entry is statically configured (Static) or dynamically-learned from IGMPv3 (Dynamic). This variable a read-only field.
Activity	Displays the current activity of the selected (S,G) entry. True indicates that traffic is flowing to the switch, otherwise, the system displays it false. This variable a read-only field.
AdminState	Configures the administrative state for the selected static entry. This state determines whether the switch uses the static entries. Configure this field to enable (default) to use the entry or disable to save for future use.

Configure SSM Range and Global Parameters

Configure the SSM range parameter to extend the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without changing their group configurations.

Before You Begin

- To change the RangeGroup configuration, you must first disable PIM.
- To change the RangeGroup configuration, you must delete all entries in the SSM channel table before you configure the new IP multicast group address.

About This Task

The other global parameters enable the IGMPv3 dynamic learning feature and configure the administrative state for all the entries in the SSM channel table.



Important

If you change the RangeGroup configuration, the switch reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), this procedure also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Ssm Global** tab.
4. Configure the appropriate fields.
5. Select **Apply**.

Ssm Global Field Descriptions

Use the data in the following table to use the **Ssm Global** tab.

Name	Description
DynamicLearning	Activates the dynamic learning of SSM channel (S,G) pairs from IGMPv3 reports. As new SSM channels are learned, the system displays them in the SSM channel table.
RangeGroup	Configures the IP multicast group address. The lowest group address is 224.0.0.0 and the highest is 239.255.255.255. The default is 232.0.0.0.

Name	Description
RangeMask	Configures the address mask of the multicast group. The default is 255.0.0.0.
SsmMapAdminAction	Configures the administrative state, which determines whether the switch uses the table entries: <ul style="list-style-type: none"> • enableAll—Globally activates all the static entries in the SSM channel table. This value does not affect the dynamically learned entries. • disableAll—Globally inactivates all the static entries in the SSM channel table. This value does not affect the dynamically learned entries.

Configure Multicast Stream Limitation on an Interface

Configure multicast stream limitation to limit the number of concurrent multicast streams on the interface. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific interface and control access to multicast streams.

About This Task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the interface drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a specific limit of multicast streams on an interface.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **StreamLimit** tab.
4. To change the status of an interface, double-click **StreamLimitEnable** for the interface, and then select **enable** or **disable**.
5. (Optional) If the interface is enabled, you can edit **Maximum Number of Stream**.
6. Select **Apply**.

StreamLimit Field Descriptions

Use the data in the following tab to use the **StreamLimit** tab.

Name	Description
Interface	Displays the slot and port number or VLAN ID for this interface.
StreamLimitEnable	Enables or disables stream limitation on this interface.

Name	Description
Maximum Number Of Stream	Configures the maximum number of streams allowed on this interface. The range is from 0–65535, and the default is 4.
Current Number Of Stream	Displays the current number of streams received on this interface. This value is a read-only value.

Configure Multicast Stream Limitation on a VLAN

Configure multicast stream limitation to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific VLAN and control access to multicast streams.

About This Task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the VLAN drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a specific limit of multicast streams on an interface.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **IGMP** tab.
7. In **StreamLimitEnable**, select **enable**.
8. Configure the maximum number of streams.
9. Select **Apply**.

Configure Multicast Stream Limitation on a Port

Configure multicast stream limitation to limit the number of concurrent multicast streams on the port. Limit the number of streams to protect the bandwidth on a specific port and control access to multicast streams.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **IGMP** tab.
5. In **StreamLimitEnable**, select **enable**.
6. Configure the maximum number of streams.
7. Select **Apply**.

Configure Multicast Stream Limitation Members

Configure multicast stream limitation members on ports of the specified interface to configure the maximum number of streams on the interface.

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **StreamLimit Members** tab.
4. Select **Insert**.
5. Type the number of the VLAN to which you want to add a member or select **Vlan** to select an ID from the list.
6. Type the number of the slot and port that you want to add as a member or select **Port**, and then select a port from the graphic display. If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.



Important

You must select one of the ports in the VLAN that you selected in step 4.

7. Type a maximum number of streams or accept the default of 4.
8. Select **Insert**.

StreamLimit Members Field Descriptions

Use the data in the following table to use the **StreamLimit Members** tab.

Name	Description
IfIndex	Displays the ID of the VLAN.
Port	<p>Lists each slot and port number for this interface with stream limitation enabled.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Name	Description
MaxStreams	Configures the maximum number of allowed streams for this specific port. The number of allowed streams cannot exceed the maximum number for the interface. The range is from 0–65535 and the default is 4.
NumStreams	Displays the current number of streams received on this interface. This value is a read-only value.

Delete Multicast Stream Limitation Member

Delete a multicast stream limitation member from an interface to remove it from the configuration.

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **StreamLimit Members** tab.
4. Select the row that lists the member you want to delete.
5. Select **Delete**.

Configure the IGMP Interface

Configure the IGMP interface to change global IGMP values for the interface. Use the Interface tab to view or edit the IGMP interface table.

About This Task

If an interface does not use an IP address, the system does not display it in the IGMP table. If an interface uses an IP address, but neither IGMP snoop or PIM is enabled, the system displays the interface as inactive.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Interface** tab.
4. Edit the appropriate information.
5. Select **Apply**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The default is 125.
Status	Shows the IGMP row status. If an interface uses an IP address and PIM-SM is enabled, the status is active. Otherwise, it is notInService.
Version	Configures the version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	Shows the version of IGMP that currently runs on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface attaches.
QueryMaxResponseTime	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds.)</p> <p>Important: You must configure this value lower than the QueryInterval.</p>
WrongVersionQueries	Shows the number of queries received with an IGMP version that does not match the interface. You must configure all routers on a LAN to run the same version of IGMP. If the interface receives queries with the wrong version, this value indicates a version mismatch.
Joins	Shows the number of times this interface added a group membership, which is the same as the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	<p>Tunes for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, increase the robustness value.</p> <p>The default value of 2 means that the switch drops one query for each query interval without the querier aging out.</p>

Name	Description
LastMembQueryIntvl	<p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of second. As a best practice, configure this parameter to values greater than 3. If you do not need a fast leave process, you can configure values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)</p>
OtherQuerierPresent Timeout	Shows the length of time that must pass before a multicast router determines that no other querier exists. If the local router is the querier, the value is 0.
FlushAction	<p>Configures the flush action to one of the following:</p> <ul style="list-style-type: none"> • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	<p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p>Important: To maximize network performance, configure this parameter according to the version of IGMP currently in use.</p> <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop.
SnoopQuerierEnable	Enables IGMP Layer 2 Querier.
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts for each channel or group. The default is disabled. You must select this parameter to use fast leave for IGMPv3.

Name	Description
McastMode	Specifies the protocol configured on the VLAN. <ul style="list-style-type: none"> snoop — Indicates IGMP snooping is enabled on a VLAN. snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over Fabric Connect for a Layer 2 VSN). pim — Indicates PIM is enabled. routed-spb — Indicates IP multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts.
ExtnUpnpFilterEnable	Enables Universal Plug and Play (uPnP) Filtering to filter multicast packets destined for a specific range. The default is disabled.
ExtnUpnpFilterAddress	Specifies the multicast destination IP address to filter on an IGMP-enabled interface. The default is 239.255.255.250/32.
ExtnUpnpFilterAddressMask	Specifies the IGMP uPnP Filtering IP subnet to which this interface is attached.
ExtnSnoopOrigin	Specifies the origin of IGMP Snooping configuration on the port. The supported values are: <ul style="list-style-type: none"> config - Configured by the user. radius - Configured by the Remote Authentication Dial-In User Service (RADIUS) attribute.
ExtnRoutedSpbQuerierAddr	Specifies the IP address of the SPB querier.

Configure IGMP Sender Entries

Configure IGMP sender entries to identify a source that sends multicast data to a multicast group.

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Sender** tab.
4. Change the appropriate options.
5. Select **Apply**.

Sender Field Descriptions

Use the data in the following table to use the **Sender** tab.

Name	Description
IfIndex	Specifies the interface where you enabled the IGMP entry.
GrpAddr	Specifies the multicast group address of the multicast stream. Within the indicated valid range (224.0.0.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an error message.
MemberAddr	Specifies the IP address of a host.
Action	Flushes an entry or a group.
TPort	Identifies the T port.
State	Indicates whether a sender exists because of an IGMP access filter. The options are filtered and not filtered.
L2Isid	Specifies the Layer 2 I-SID of the C-VLAN.
Scopelsid	Indicates the scope I-SID of SPB routed multicast sender. If the sender does not support an SPB scope I-SID, the field displays an invalid value of 0xFFFFFFFF.
Datalsid	Indicates the data I-SID of SPB routed multicast sender. If the sender does not support an SPB data I-SID, the field displays an invalid value of 0xFFFFFFFF.

Configure Fast Leave Mode

Configure fast leave mode to control all IGMP fast leave enabled interfaces.

Before You Begin

- You must enable explicit-host-tracking before configuring fast-leave mode. To enable explicit-host-tracking, see [Configure IGMP Parameters on a Port](#) on page 1808 and [Configure IGMP Parameters on a VLAN](#) on page 1810.

About This Task

Fast leave relies on an alternative leave process where the switch stops sending traffic for the group immediately after it receives a leave message, without issuing a query to check if other group members exist on the network. Use this global parameter to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces.

**Important**

Fast leave mode applies only to fast leave enabled IGMP interfaces.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Global** tab.
4. Select the mode.
5. Select **Apply**.

Global Field Descriptions

Use the data in the following table to use the **Global** tab.

Name	Description
FastLeaveMode	Configures the mode to one of the following values: <ul style="list-style-type: none">• multipleUser: Removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This value is the default.• oneUser: Removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.
GenerateTrap	Generates a trap. The default is disable.
GenerateLog	Generates a log message. The default is disable.

Configure Multicast Access Control for an Interface

Configure multicast access control for a selected IGMP interface or VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Access Control** tab.
4. Select **Insert**.
5. Type the number of the slot and port or the VLAN ID that you want to add as a member or select **Port** or **Vlan** and then select a port or VLAN from the graphic display.
6. Select ... near **PrefixListId**.
7. Select a prefix list ID.
8. Select **OK**.
9. Type the host address and host mask.

10. Select the action mode that you want for the specified host.
11. Select **Insert**.

Access Control Field Descriptions

Use the data in the following table to use the **Access Control** tab.

Name	Description
IfIndex	Specifies the interface where the IGMP entry is enabled.
PrefixListId	Specifies a numeric string that identifies the prefix list.
HostAddr	Specifies the IP address of the host.
HostMask	Specifies the subnet mask that determines the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
PrefixListName	Specifies the name of the prefix list.
ActionMode	Specifies the action for the host identified by HostAddr. The options include the following: <ul style="list-style-type: none">• denied IP multicast transmitted traffic (deny-tx).• denied IP multicast received traffic (deny-rx).• denied both IP multicast transmitted and received traffic (deny-both).• allowed IP multicast transmitted traffic (allow-only-tx).• allowed IP multicast received traffic (allow-only-rx).• allowed both IP multicast transmitted and received traffic (allow-only-both).

View IGMP Cache Information

View IGMP cache information to view the group for which members exist on a specific interface.

About This Task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Cache** tab.

Cache Field Descriptions

Use the data in the following table to use the **Cache** tab.

Name	Description
Address	Shows the IP multicast group address for this entry that contains this information.
IfIndex	Shows the interface from which the corresponding multicast group address is heard.
LastReporter	Shows the IP address of the source of the last membership report received for this IP multicast group address on this interface. If no membership report is received, the object uses the value 0.0.0.0.
ExpiryTime	Shows the amount of time (in seconds) that remain before this entry ages out.
Version1HostTimer	Shows the time that remains until the local router assumes that no IGMPv1 members exist on the IP subnet attached to the interface. Upon hearing IGMPv1 membership report, this value resets to the group membership timer. When the time that remains is nonzero, the local router ignores IGMPv2 leave messages for this group that it receives on this interface.
Type	Shows the type of IGMP entry.
StaticPorts	Shows the static ports associated with the entry.

View IGMPv3 Cache

View the IGMPv3 specific data corresponding to each interface, port, and multicast group pair on a router.

About This Task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **IGMPv3 Cache** tab to view the IGMPv3 cache information.

IGMPv3 Cache Field Descriptions

Use the data in the following table to use the **IGMPv3 Cache** tab.

Name	Description
GroupAddress	Specifies the Multicast group Address (Class D) that others want to join. A group address can be the same for many incoming ports.
IfIndex	An unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports from various sources.
InPort	An unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports from various sources.
ModeExpiryTimer	Represents the time remaining before the interface EXCLUDE state expires and the interface state transitions to INCLUDE mode. This value is applicable only to IGMPv3-compatible nodes.
Version1HostTimer	Specifies the time remaining until the local router assumes that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. This entry only applies to IGMPv1 hosts. Upon hearing any IGMPv1 report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.
Version2HostTimer	Specifies the time remaining until the local router assumes that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. Assuming no IGMPv1 hosts have been detected, the local router does not ignore any IGMPv2 Leave messages for this group that it receives on this interface.
SourceFilterMode	Specifies the current group state, applicable to IGMPv3-compatible nodes. The value indicates whether the state is INCLUDE or EXCLUDE.

View and Edit Multicast Router Discovery Information

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Multicast Router Discovery** tab.

- To edit the current configuration, double-click the value, make the change, and then select **Apply**.

Multicast Router Discovery *Field Descriptions*

Use the data in the following table to use the **Multicast Router Discovery** tab.

Name	Description
Interface	Shows the interface where IGMP is enabled.
MrdiscEnable	Enables (true) or disables (false) the router interface to listen for multicast router discovery messages to determine where to send multicast source data and IGMPv2 reports. If you enable snoop, you automatically enable multicast router discovery.
DiscoveredRouterPorts	Lists ports that the Multicast Router Discovery (MRDISC) protocol discovers. Important: The switch does not support the MRDISC protocol on brouter ports.
MaxAdvertiseInterval	Shows the maximum time allowed between sending router advertisements from the interface, in seconds. The range is from 2–180 seconds. The default is 20 seconds.
MinAdvertiseInterval	Shows the minimum time allowed between sending unsolicited router advertisements from the interface, in seconds. This value must be more than 3 seconds but no greater than the value assigned to the MaxAdvertiseInterval value.
MaxInitialAdvertiseInterval	Configures the maximum number (in seconds) of multicast advertisement intervals that you can configure on the switch.
MaxInitialAdvertisements	Configures the maximum number of initial multicast advertisements that you can configure on the switch.
NeighborDeadInterval	Shows the time interval (in seconds) before the router interface drops traffic after a user leaves the multicast group.

View the IGMP Router Source List

View the source list entries corresponding to each interface and multicast group pair on a router.

About This Task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Igmp Router Source List** tab to view the IGMPv3 cache information.

Igmp Router Source List Field Descriptions

Use the data in the following table to use the **Igmp Router Source List** tab.

Name	Description
GroupAddress	Specifies the IP multicast group address for which this entry contains information.
IfIndex	Specifies the interface for which this entry contains information for an IP multicast group address.
InPort	Specifies a unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports for this source.
HostAddress	Specifies the host address to which this entry corresponds.
MemberAddress	Specifies the IP Address of a member that has sent source specific report wishing to join this source.
Expire	This value indicates the relevance of the source list entry, where a non-zero value indicates this is an INCLUDE state value, and a zero value indicates this to be an EXCLUDE state value.
Mode	Specifies the current member state, applicable to IGMPv3-compatible nodes. The value indicates whether the state is INCLUDE or EXCLUDE.
MemberExpire	This value indicates the time until the member for this source expires.

View IGMP Snoop Information

View information about IGMP snoop to see the current configuration.

About This Task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Snoop** tab.

Snoop Field Descriptions

Use the data in the following table to use the **Snoop** tab.

Name	Description
Interface	Shows the VLAN ID for the VLAN.
SnoopEnable	Shows the status of IGMP snoop. IGMP snoop works only if a multicast router exists in the VLAN.
SsmSnoopEnable	Shows the status of SSM snoop.
ProxySnoopEnable	Indicates whether the IGMP report proxy feature is enabled. If you enable this feature, the switch forwards reports from hosts to the multicast router once for each group for each query interval, or after new group information is available. If you disable this feature, the switch forwards all reports from different hosts to multicast routers, and can forward more than one group report for the same multicast group for each query interval. The default is enabled.
FastLeaveEnable	Shows the status of fast leave for this port.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router. Important: Configure this variable only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.
SnoopActiveMRouterPorts	Shows the active multicast router ports. Active multicast router ports are ports that directly attach to a multicast router. These ports include the querier port and all ports in the forwarding state that you configure as well as those that were dynamically learned through receiving queries.
SnoopMRouterExpiration	Indicates the time that remains before the multicast router ages out. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The query maximum response interval (obtained from the queries received) is used as the timer resolution.

View IGMP Snoop Trace Information

View the multicast group trace to track the data flow path of multicast streams.

Procedure

1. In the navigation pane, expand **Configuration > IP**.

2. Select **IGMP**.
3. Select the **Snoop Trace** tab.

Snoop Trace Field Descriptions

Use the data in the following table to use the **Snoop Trace** tab.

Name	Description
GrpAddr	Displays the IP multicast address of the group traversing the router.
SrcAddr	Displays the IP source address of the multicast group.
OutVlan	Displays the egress VLAN ID for the multicast group.
InPort	Displays the ingress port for the multicast group.
InVlan	Displays the ingress VLAN ID for the multicast group.
OutPort	Displays the egress port of the multicast group.
Type	Displays the port type on which the snoop entry is learned.

View IGMP Group Information

View information about IGMP groups to see the current group operation on the switch.

About This Task



Note

The following procedure displays the dynamically learned IGMP groups. **IP > IGMP > Static** displays statically configured IGMP groups. This is in contrast to the CLI command **show ip igmp group**, which displays both dynamically learned and statically configured IGMP groups, and the CLI command **show ip igmp static**, which displays only the statically configured groups.

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Groups** tab.

Groups Field Descriptions

Use the data in the following table to use the **Groups** tab.

Name	Description
IpAddress	Shows the multicast group address (Class D). A group address can be the same for many incoming ports.
Members	Shows the IP address of the host that issues the membership report to this group.

Name	Description
InPort	Shows the port that receives the group membership report.
IfIndex	Shows a unique value that identifies a physical interface or a logical interface (VLAN) that receives the membership report.
Expiration	Shows the time left before the group report expires on this port. This variable is updated after the port receives a group report.

Route Management Using the CLI

With multicast route commands, you can configure and view IP multicast routing parameters on the switch.

Configure Multicast Stream Limits

Limit the number of multicast streams to protect the CPU from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

About This Task

You can enable or disable the mroute stream limit for the entire device or for individual ports when the switch is operating. If you enable the mroute stream limit for the device and for an individual port, only the periodic check is performed for that port.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable stream limitation globally:

```
ip mroute stream-limit
```
3. Enter GigabitEthernet Interface Configuration mode.

```
interface gigabitethernet {slot/port[-slot/port]}[,...][slot/all][all]}
```
4. Enable stream limits:

```
ip mroute stream-limit
```
5. For Gigabit Ethernet interfaces, configure the maximum number of streams and the interval at which to sample:

```
ip mroute max-allowed-streams <1-32768> max-allowed-streams-timer-check <1-3600>
```

6. Show the mroute stream limit configuration:

```
show ip mroute interface gigabitethernet {slot/port[-slot/port]}[,...]  
[slot/all][all]}
```

Example

```
Switch:1>enable  
Switch:1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#ip mroute stream-limit  
Switch:1(config)#interface gigabitethernet 1/6  
Switch:1(config-if)#ip mroute stream-limit  
Switch:1(config-if)#ip mroute max-allowed streams 1000 max-allowed-streams-timer-check 20
```

Variable Definitions

Use the data in the following table to use the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Use the data in the following table to use the **ip mroute** command.

Variable	Value
max-allowed-streams <1-32768>	Configures the maximum number of streams on the specified port. The port is shut down if the number of streams exceeds this limit. The default value is 1984 streams.
max-allowed-streams-timer-check <1-3600>	Configures the sampling interval, which checks if the number of ingress multicast streams to the CPU is under a configured limit or if the port needs to shut down. The default value is 10 seconds.

Configure Multicast Static Source Groups

Configure static source group entries in the Protocol Independent Multicast (PIM) multicast routing table. The PIM cannot prune these entries from the distribution tree.

Before You Begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
 - PIM-Sparse Mode (SM)
 - PIM-Source Specific Multicast (SSM)

About This Task

Even if no receivers exist in the group, the multicast stream for a static source group entry remains active.

The maximum number of static source groups must not exceed 1024.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Configure a static source group entry:


```
ip mroute static-source-group <A.B.C.D> <A.B.C.D/X>
```

Example

Create a static source group for two multicast groups: 224.32.2.1 and 226.50.2.2. The static source group for group 224.32.2.1 is for a source subnet 10.10.10.0/24. The static source group for group 226.50.2.2 is for the host 20.20.20.100/32.

```
Switch:1(config)# ip mroute static-source-group 224.32.2.1 10.10.10.0/24
Switch:1(config)# ip mroute static-source-group 226.50.2.2 20.20.20.100/32
```

Variable Definitions

Use the definitions in the following table to use the **ip mroute static-source-group** command.

Variable	Value
<i>A.B.C.D</i>	Specifies the IP address of the multicast group.
<i>A.B.C.D/X</i>	Specifies the multicast source IP address and subnet mask for the static source group entry. You cannot create duplicate groups. How you configure the source address depends on the protocol and mode you use.

Configure IP Multicast Software Forwarding

When you use the IP multicast software forwarding feature you can avoid initial data loss experienced by multicast applications; this is suitable for low bandwidth conditions.

When you configure the IP multicast software forwarding feature the system forwards the initial packets of an IP multicast data stream it receives and creates a corresponding hardware record for subsequent packets.

By default, multicast software forwarding is disabled.

About This Task



Note

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

IP multicast software forwarding is a global system configuration feature that is only applicable to traditional PIM protocol and IGMP Snooping protocols, not SPB-PIM Gateway or Layer 3 VSN SPB Multicast. If you enable IP multicast software forwarding, the hardware continues to forward IP multicast traffic. The software only forwards initial data traffic.

After a new data stream arrives, the first data packet is sent to the CPU, which programs the multicast route in hardware, and all packets that arrive subsequent to this programming are forwarded by hardware only.

If you enable software forwarding, all initial packets received before hardware programming is complete are sent to the CPU for forwarding and packet suppression by the hardware is disabled.

If you do not enable software forwarding, only the first data packet is sent to the CPU and subsequent packets are suppressed by the hardware so that the CPU is not overwhelmed with traffic. During this time, packets suppressed by the hardware are dropped.



Important

To avoid overloading the CPU, ensure that you do not use the IP multicast software forwarding feature for video multicast applications.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable software forwarding:
`multicast software-forwarding`
3. Show the software forwarding configuration:
`show multicast software-forwarding`

Example

```
Switch:1#show multicast software-forwarding

=====
                        Mcast Software Forwarding - GlobalRouter
=====
McastSoftwareForwarding      :enabled
```

Configure the Resource Usage Counter for Multicast Streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing the switch.

About This Task**Note**

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

After you configure the counter thresholds for ingress and egress records, if the record usage exceeds the threshold, you receive notification by a trap on the console, a logged message, or both.

If you do not configure the thresholds, the system displays only the ingress and egress records currently in use.

You can configure the resource usage counter on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the thresholds:

```
ip mroute resource-usage egress-threshold <0-32767> ingress-threshold <0-32767>
```

3. Configure one of the following notification methods:

- Configure a log-only notification method:

```
ip mroute resource-usage log-msg
```

- Configure a trap-only notification method:

```
ip mroute resource-usage trap-msg
```

- Configure both notification methods:

```
ip mroute resource-usage log-msg trap-msg
```

Example

Configure the egress threshold to 200.

```
Switch:1(config)# ip mroute resource-usage egress-threshold 200
```

Configure the ingress threshold to 100.

```
Switch:1(config)# ip mroute resource-usage ingress-threshold 100
```

Enable the log message notification method.

```
Switch:1(config)# ip mroute resource-usage log-msg
```

Variable Definitions

Use the data in the following table to use the **ip mroute resource-usage** command.

Variable	Value
<i>egress-threshold</i> <0-32767>	Configures the egress record threshold (S,G). The system sends a notification message after the number of streams exceeds a threshold level. The default is 0.
<i>ingress-threshold</i> <0-32767>	Configures the ingress record threshold. The system sends a notification message after the number of streams exceeds a threshold level. The default is 0.

Configure Prefix Lists

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

About This Task



Important

When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a prefix list:

```
ip prefix-list WORD<1-64> {A.B.C.D/X} [ge <0-32>] [le <0-32>]
```

3. (Optional) Rename an existing prefix list:

```
ip prefix-list WORD<1-64> name WORD<1-64>
```

4. Display the prefix list:

```
show ip prefix-list [prefix {A.B.C.D}] [vrf WORD<1-16>] [vrfids WORD<0-512>] [WORD <1-64>]
```

Example

Configure a prefix-list. Display the prefix list.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list LIST1 47.17.121.50/255.255.255.0
Switch:1(config)#show ip prefix-list LIST1
=====
                        Prefix List - GlobalRouter
=====

```

PREFIX	MASKLEN	FROM	TO

List 1	LIST1:		
	47.17.121.50	24	24 24
1 Total Prefix List entries configured			

```
Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
```

Variable Definitions

The following table defines parameters for the **ip prefix-list** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats: <ul style="list-style-type: none"> a.b.c.d/x a.b.c.d/x.x.x.x default
ge <0-32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
le <0-32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
name WORD<1-64>	Renames the specified prefix list. The name length is 1-64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

The following table defines parameters for the **show ip prefix-list** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the prefix to include in the command output.
<i>vrf WORD<1-16></i>	Specifies the name of the VRF.
<i>vrfids WORD<0-512></i>	Specifies the ID of the VRF and is an integer in the range of 0–512.
<i>WORD<1-64></i>	Specifies a prefix list, by name, to use for the command output.

The following table defines parameters for the **show ip prefix-list** command output.

Variable	Value
PREFIX	Indicates the member of a specific prefix list.
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
TO	Indicates the prefix mask endpoint in bits.

View the Hardware Resource Usage

About This Task

The switch can query the number of ingress and egress IP multicast streams traversing the switch. After you configure the thresholds for ingress and egress records, if the record-usage goes beyond the threshold, the device notifies you by way of a trap on the console, logged message, or both.

If you do not configure the thresholds, the switch displays only the ingress and egress records currently in use.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the hardware resource usage:

```
show ip mroute hw-resource-usage
```

Example

Show the hardware resource usage:

```
Switch:1>show ip mroute hw-resource-usage
=====
Multicast Hardware Resource Usage
=====
EGRESS      INGRESS      EGRESS      INGRESS      LOG MSG      SEND TRAP      SEND TRAP
REC IN-USE  REC IN-USE  THRESHOLD   THRESHOLD   ONLY         ONLY          AND LOG
```

0	0	0	0	false	false	false		

Route Management Using EDM

View or edit interface configuration information for Layer 3 IP multicast protocols on the switch.

View Multicast Route Information

View multicast route information for troubleshooting purposes.

This tab shows multicast routing information for IP datagrams from a particular source and addressed to a particular IP multicast group address.

About This Task



Note This procedure is supported on a DvR Controller; it is not supported on a DvR Leaf node.

You can view the multicast routes for a Layer 3 Virtual Services Network (VSN) the same way you view the Global Router except that you must first launch the appropriate VRF context.

Use the following procedure to determine the data stream learned when IP multicast is configured on a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > IP > Multicast**.
2. Select the **Routes** tab.

Routes Field Descriptions

Use the data in the following table to use the **Routes** tab.

Name	Description
Group	Displays the IP multicast group address for this entry that contains multicast routing information.
Source	Displays the network address that, when combined with the corresponding route SourceMask value, identifies the source that contains multicast routing information.
SourceMask	Displays the network mask that, when combined with the corresponding route Source value, identifies the multicast source.
UpstreamNeighbor	Shows the address of the upstream neighbor from which the IP datagrams from these sources are received. The address is 0.0.0.0 if the network is local.

Name	Description
Interface	Displays the interface, slot and portnumber, or VLAN ID where IP datagrams sent by these multicast sources to this multicast address are received.
ExpiryTime	Displays the amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none"> • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb (12) • spbpimgw(13)

View Multicast Next-Hop Information

View all multicast next-hop information.

This tab shows information about the next hops used by outgoing interfaces to route IP multicast datagrams. Each entry is one in a list of next hops on outgoing interfaces for particular sources that send to a particular multicast group address.

About This Task

You can view the multicast routes for a Layer 3 Virtual Services Network (VSN) the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Next Hops** tab.

Next Hops Field Descriptions

Use the data in the following table to use the **Next Hops** tab.

Name	Description
Group	Displays the IP multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.

Name	Description
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
ReceiverPort	Displays the receiver port for this next hop.
OutInterface	Displays the interface slot and portnumber or VLAN ID for the outgoing interface for this next hop.
Address	Displays the address of the next hop specific to this entry. For most interfaces, it is identical to the next-hop group. Non Broadcast Multiple Access (NBMA) interfaces, however, can use multiple next hop addresses out of a single outgoing interface.
State	Displays whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. A value of forwarding indicates the information is currently used; pruned indicates it is not used.
UpTime	Displays the up time for this entry.
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	Displays the minimum number of hops between this router and members of the IP multicast group reached through the next hop on this outgoing interface. IP multicast datagrams for the group that use a time-to-live less than this number of hops are not forwarded to the next hop.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none"> • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb
Pkts	Displays the number of next hop packets.

View Multicast Interface Information

View multicast interface information to verify the multicast configuration.

This tab shows multicast routing information specific to interfaces.

About This Task

You can view multicast interface information for a Layer 3 VSN the same way you view the Global Router except that you must first launch the appropriate VRF context.



Note

This procedure is supported on a DvR Controller; it is not supported on a DvR Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Interfaces** tab.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Interface	Displays the slot and port number or VLAN ID for this entry.
Ttl	Displays the datagram time-to-live (TTL) threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means that all multicast packets are forwarded out of the interface.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none">• other(1): none of the following• local(2): manually configured• netmgmt(3): configured by a network management protocol• pimSparseMode(8): PIM-SMv2• igmpOnly(10)• pimSsmMode(11)• spb

Add New Static Source Groups

Add a new static source group to create an entry that the switch cannot prune from the distribution tree.

Before You Begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
 - PIM-SM
 - PIM-SSM

About This Task

An attempt to add a duplicate of an existing source-group entry results in an error message.



Note

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

The switch supports PIM only in the Global Router. You cannot configure static source groups for specific VRF contexts.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Static Source Group** tab.
4. Select **Insert**.
5. Complete the information in the dialog box.
6. Select **Insert**.

Edit Static Source Groups

Configure static source-group entries in the PIM multicast routing table. PIM cannot prune these entries from the distribution tree. In other words, even if no receivers exist in the group, the multicast stream for a static source-group entry stays active.

Before You Begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
 - PIM-Sparse Mode (SM)
 - PIM-Source Specific Multicast (SSM)

About This Task



Note

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

The maximum number of static source groups must not exceed 1024.

The switch supports PIM only in the Global Router. You cannot configure static source groups for specific VRF contexts.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Static Source Group** tab.

4. Edit the required information.
5. Select **Apply**.

Static Source Group Field Descriptions

Use the data in the following table to use the **Static Source Group** tab.

Name	Description
GroupAddress	Configures the multicast group IP address for this static source-group entry.
SourceSubnet	Configures the multicast source address for this static source-group entry. How you configure the source address depends on the protocol and mode you use.
SrcSubnetMask	Configures the subnet mask of the source for this static source-group entry.

Configure IP Multicast Software Forwarding

Configure IP multicast software forwarding to enable the system to initially forward IP multicast data until a hardware record is created. The system forwards the initial packets of a stream it receives and creates a corresponding hardware record for subsequent packets. The advantage of this feature is that it avoids initial data loss experienced by multicast applications and is most suited for low bandwidth.

About This Task



Note

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

IP multicast software forwarding is a global system configuration feature that is only applicable to traditional PIM protocol and IGMP Snooping protocols, not SPB-PIM Gateway or Layer 3 VSN SPB Multicast. If you enable IP multicast software forwarding, the hardware still forwards IP multicast traffic. The software forwards only initial data traffic.

After a new data stream arrives, the first data packet is sent to the CPU, which programs the multicast route in hardware, and all packets that arrive subsequent to this programming are forwarded by hardware only. If you enable software forwarding, all initial packets received before hardware programming is complete are sent to the CPU for forwarding. If you enable software forwarding, packet suppression by the hardware is disabled. If you do not enable software forwarding, only the first data packet is sent to the CPU and subsequent packets are suppressed by the hardware so that the CPU is not overwhelmed with traffic. During this time, packets suppressed by the hardware are dropped.

By default, the feature is disabled.

**Important**

To avoid overloading the CPU, do not use the IP multicast software forwarding feature for video multicast applications.

If you configure multicast software forwarding from within a VRF context, the configuration applies to the Global Router and all VRF contexts. You cannot change the multicast software forwarding configuration for individual VRF contexts.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Globals** tab.
4. Select the **SWForwardingEnable** check box.
5. Select **Apply**.

Globals *Field Descriptions*

Use the data in the following table to use the **Globals** tab.

Name	Description
SWForwardingEnable	Enables the system to initially forward IP multicast data until a hardware record is created. The default is disabled.
StatsEnabled	Enables or disables multicast route statistics. The default is disabled.
StatsClear	Clears multicast route statistics.

Configure Mroute Stream Limit

Limit the number of multicast streams to protect a CPU from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Mroute Stream Limit** tab.
5. Select the **StreamLimitEnable** box.
6. Edit other fields as required.
7. Select **Apply**.

Mroute Stream Limit Field Descriptions

Use the data in the following table to use the **Mroute Stream Limit** tab.

Name	Description
StreamLimitEnable	Enables or disables mroute stream limit on the port.
StreamLimit	Specifies the maximum number of multicast streams allowed to enter the CPU through this port.
StreamTimerCheck	Specifies the sampling period, in seconds, to check the number of multicast streams that enter the CPU through this port.

Configure Mroute Stream Limit on an Extreme Integrated Application Hosting Port

**Note**

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to limit the number of multicast streams to protect a Central Processing Unit (CPU) from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization, or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the Extreme Integrated Application Hosting (IAH) port you want to configure.
3. Select the **Mroute Stream Limit** tab.
4. Select **StreamLimitEnable**.
5. Configure other fields as required.
6. Select **Apply**.

Mroute Stream Limit Field Descriptions

Use data in the following table to configure the **Mroute Stream Limit** tab.

Name	Description
StreamLimitEnable	Enables or disables mroute stream limit on the Extreme Integrated Application Hosting (IAH) port. The default is disabled.
StreamLimit	Specifies the maximum number of multicast streams allowed to enter the CPU through the IAH port. The default value is 1984.
StreamTimerCheck	Specifies the sampling period, in seconds, to check the number of multicast streams that enter the CPU through the IAH port. The default is 10 seconds.

Configure Resource Usage Counter for Multicast Streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing the switch. After you configure the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive notification through a trap on the console, a logged message, or both.

About This Task**Note**

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

**Important**

If you do not configure the thresholds, EDM displays only the ingress and egress records that are currently in use.

You can configure the resource usage counter on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Resource Usage** tab.
4. Configure the ingress and egress thresholds.
5. Configure the notification methods.
6. Select **Apply**.

Resource Usage Field Descriptions

Use the data in the following table to use the **Resource Usage** tab.

Name	Description
Egress Records In-Use	Displays the number of egress records traversing the switch.
Ingress Records In-Use	Displays the number of ingress records (source or group) traversing the switch.
Egress Threshold	Configures the egress threshold level (0–32767).
Ingress Threshold	Configures the ingress threshold level (0–32767).
SendTrapOnly	Sends only trap notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type. You can configure only one notification type.
SendTrapAndLog	Sends both trap and log notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type.
LogMsgOnly	Sends only log notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type.

Multicast route statistics configuration using the CLI

The following sections provide procedural information you can use to configure multicast route statistics using the Command Line Interface (CLI).

Enable IP Multicast Route Statistics

Enable the collection and display of IP multicast route statistics.

These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on multicast group classification. By default, collection of multicast route statistics is disabled.



Note

When you enable IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Enable the collection of IP multicast route statistics.

```
ip mroute stats enable
```

3. (Optional) Set the IP multicast route statistics to default.

```
default ip mroute stats enable
```

4. (Optional) Disable the collection of IP multicast route statistics.

```
no ip mroute stats enable
```

5. View the IP multicast route statistics.

```
show ip mroute stats [WORD<3-160> {A.B.C.D[,E.F.G.H][,...]}]
```



Note

The maximum number of multicast group IP addresses is 10.

Example

Enable the collection of IP multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip mroute stats enable
```

View the IP multicast route statistics:

```
Switch:1#show ip mroute stats

=====
                        Multicast Stats
=====
-----
                        Statistics : Enabled
```

View the statistics for the multicast group IP address 225.0.0.1:

```
Switch:1#show ip mroute stats 225.0.0.1

=====
                        Multicast Stats - GlobalRouter
=====
GroupAddress      SourceCounter      IngressPackets      IngressBytes      AverageSize
-----
225.0.0.1          1                   30452198             3897881344         128
```

View the statistics for multiple (up to a maximum of 10) group IP addresses.

```
Switch:1#show ip mroute stats
225.0.0.1,225.0.0.2,225.0.0.3,225.0.0.4,225.0.0.5,225.0.0.6,225.0.0.7,225.0.0.8,225.0.0.9,
225.0.0.10

=====
                        Multicast Stats - GlobalRouter
=====
GroupAddress      SourceCounter      IngressPackets      IngressBytes      AverageSize
-----
225.0.0.1          1                   32446194             4153112832         128
225.0.0.2          1                   32446196             4153112960         127
225.0.0.3          1                   32446197             4153113088         127
225.0.0.4          1                   32446198             4153113216         127
225.0.0.5          1                   32446199             4153113472         128
```

225.0.0.6	1	32446200	4153113600	128
225.0.0.7	1	32446201	4153113728	128
225.0.0.8	1	32446203	4153113856	127
225.0.0.9	1	32446203	4153113856	127
225.0.0.10	1	32446203	4153113984	128

Variable definitions

Use the data in the following table to use the **show ip mroute stats** command.

Variable	Definition
WORD<3-160> {A.B.C.D[,E.F.G.H][,...]}	Specifies the multicast group IP address for which to display statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses. You can specify a maximum of 10 groups.

Clear IP Multicast Route Statistics

Use this procedure to clear the IP multicast route statistics. This resets the IP multicast statistics counters.



Note

When you clear IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Clear the IP multicast route statistics:
`clear ip mroute stats`

Example:

Clear the IP multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clear ip mroute stats
```

Monitor IP Multicast Route Statistics

Use this procedure to monitor the IP multicast route statistics at regular intervals.

Procedure

1. Enter Privileged EXEC mode:
`enable`

2. Monitor the IP multicast route statistics:

```
monitor ip mroute stats WORD<7-160> {A.B.C.D[,E.F.G.H][,...]} [vrf
WORD<1-32>] [vrfid <0-255>]
```



Note

You can monitor a maximum of 10 group IP addresses.

Examples

Monitor the IP multicast route statistics for the group IP address 225.0.0.1. In this example, the statistics are monitored at intervals of 5 seconds for a duration of 300 seconds. The following example displays monitoring output from three consecutive intervals.

```
Switch:1>en
Switch:1#monitor ip mroute stats 225.0.0.1
```

```

                                MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Mon Dec 21 16:12:07 2015

=====
                        Multicast Stats - GlobalRouter
=====
GroupAddress      SourceCounter      IngressPackets      IngressBytes      AverageSize
-----
225.0.0.1         1                   4716624              603727872         128
                                MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Mon Dec 21 16:12:13 2015

=====
                        Multicast Stats - GlobalRouter
=====
GroupAddress      SourceCounter      IngressPackets      IngressBytes      AverageSize
-----
225.0.0.1         1                   4767325              610217600         128
                                MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Mon Dec 21 16:12:19 2015

...
...

```

Monitor the IP multicast route statistics for a maximum of 10 group IP addresses. The statistics are monitored at intervals of 5 seconds for a duration of 300 seconds. The following example displays monitoring output from three consecutive intervals.

```
Switch:1#monitor ip mroute stats
225.0.0.1,225.0.0.2,225.0.0.3,225.0.0.4,225.0.0.5,225.0.0.6,225.0.0.7,225.0.0.8,225.0.0.9,225.0.0.10
                                MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Mon Dec 21 16:22:07 2015

=====
                        Multicast Stats - GlobalRouter
=====

```

GroupAddress	SourceCounter	IngressPackets	IngressBytes	AverageSize
225.0.0.1	1	9532039	1220100992	128
225.0.0.2	1	9532041	1220101120	127
225.0.0.3	1	9532042	1220101248	127
225.0.0.4	1	9532043	1220101376	127
225.0.0.5	1	9532044	1220101632	128
225.0.0.6	1	9532045	1220101760	128
225.0.0.7	1	9532046	1220101888	128
225.0.0.8	1	9532047	1220101888	127
225.0.0.9	1	9532048	1220102016	127
225.0.0.10	1	9532048	1220102144	128

MULTICAST STATISTIC

Monitor Interval: 5sec | Monitor Duration: 300sec

Mon Dec 21 16:22:13 2015

Multicast Stats - GlobalRouter

GroupAddress	SourceCounter	IngressPackets	IngressBytes	AverageSize
225.0.0.1	1	9582672	1226582016	128
225.0.0.2	1	9582674	1226582144	127
225.0.0.3	1	9582675	1226582272	127
225.0.0.4	1	9582676	1226582400	127
225.0.0.5	1	9582677	1226582656	128
225.0.0.6	1	9582678	1226582784	128
225.0.0.7	1	9582679	1226582912	128
225.0.0.8	1	9582681	1226583040	127
225.0.0.9	1	9582681	1226583040	127
225.0.0.10	1	9582681	1226583168	128

MULTICAST STATISTIC

Monitor Interval: 5sec | Monitor Duration: 300sec

Mon Dec 21 16:22:19 2015

Multicast Stats - GlobalRouter

GroupAddress	SourceCounter	IngressPackets	IngressBytes	AverageSize
225.0.0.1	1	9625009	1232001152	128
225.0.0.2	1	9625011	1232001280	127
225.0.0.3	1	9625012	1232001408	127
225.0.0.4	1	9625013	1232001536	127
225.0.0.5	1	9625014	1232001792	128
225.0.0.6	1	9625015	1232001920	128
225.0.0.7	1	9625016	1232002048	128
225.0.0.8	1	9625018	1232002176	127
225.0.0.9	1	9625019	1232002304	127
225.0.0.10	1	9625018	1232002304	128

...

...

Switch:1#

Variable definitions

The following table defines parameters for the **monitor ip mroute stats** command.

Variable	Definition
<code>vrf WORD<1-32></code>	Specifies a VRF name.
<code>vrfid <0-255></code>	Specifies a VRF ID.
<code>WORD<7-160></code> <code>{A.B.C.D[,E.F.G.H][,...]}</code>	Specifies the multicast group IP address for which to monitor statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses, up to a maximum of 10.

Enable IPv6 Multicast Route Statistics

Enable the collection of IPv6 multicast route statistics.

These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on multicast group classification. By default, collection of multicast route statistics is disabled.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable the collection of IPv6 multicast route statistics.
`ipv6 mroute stats enable`
3. (Optional) Set the IPv6 multicast route statistics to default:
`default ipv6 mroute stats`
4. (Optional) Disable the collection of IPv6 multicast route statistics.
`no ipv6 mroute stats`
5. View the IPv6 multicast route statistics.
`show ipv6 mroute stats [WORD<7-400> {Ipv6address[,Ipv6address][,...]}]`



Note

The maximum number of multicast group IP addresses is 10.

Example:

Enable collection of IPv6 multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ipv6 mroute stats enable
```

View the IPv6 multicast route statistics:

```
Switch:1#show ipv6 mroute stats
```

```
=====
Multicast Stats
```



```
=====
-----
Statistics : Enabled
```

View the statistics for the multicast group IP address FF05::1:

```
Switch#show ipv6 mroute stats FF05::1

=====
Multicast Stats - GlobalRouter
=====
GroupAddress          SourceCounter    IngressPackets    IngressBytes
AverageSize
-----
ff05:0:0:0:0:0:0:1    1                1962750            2355300000        1200
```

View the statistics for multiple group IP addresses (up to a maximum of 10).

```
Switch#show ipv6 mroute stats
FF05::1,FF05::2,FF05::3,FF05::4,FF05::5,FF05::6,FF05::7,FF05::8,FF05::9,FF05::a

=====
Multicast Stats - GlobalRouter
=====
GroupAddress          SourceCounter    IngressPackets    IngressBytes
AverageSize
-----
ff05:0:0:0:0:0:0:1    1                2027508            2433009600        1200
ff05:0:0:0:0:0:0:2    1                2027507            2433008400        1200
ff05:0:0:0:0:0:0:3    1                2027507            2433008400        1200
ff05:0:0:0:0:0:0:4    1                2027507            2433008400        1200
ff05:0:0:0:0:0:0:5    1                2027507            2433008400        1200
ff05:0:0:0:0:0:0:6    1                2027505            2433006000        1200
ff05:0:0:0:0:0:0:7    1                2027505            2433006000        1200
ff05:0:0:0:0:0:0:8    1                2027505            2433006000        1200
ff05:0:0:0:0:0:0:9    1                2027505            2433006000        1200
ff05:0:0:0:0:0:0:a    1                2027505            2433006000        1200
```

Variable definitions

Use the data in the following table to use the **show ipv6 mroute stats** command

Variable	Definition
WORD<7-400> {Ipv6address[, Ipv6address] [,...]}	Specifies the multicast group IP address for which to display statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses. You can specify a maximum of 10 groups.

Clear IPv6 Multicast Route Statistics

Use this procedure to clear the IPv6 multicast route statistics. This resets the IP multicast statistics counters.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Clear the IPv6 multicast route statistics:
clear ipv6 mroute stats

Example:

Clear the IPv6 multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clear ipv6 mroute stats
```

Monitor IPv6 Multicast Route Statistics

Use this procedure to monitor IPv6 multicast route statistics at regular intervals.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Monitor IPv6 multicast route statistics:
monitor ipv6 mroute stats WORD<7-400> {Ipv6address[,Ipv6address]
[,...]}



Note

You can monitor a maximum of 10 group IP addresses.

Examples

Monitor the IPv6 multicast route statistics for the group IPv6 address FF05::1. In this example, the statistics are monitored at intervals of 5 seconds for a duration of 300 seconds.

The following example displays monitoring output from three consecutive intervals.

```
Switch:1>enable
Switch:1#monitor IPv6 mroute stats FF05::1
```

```

                                MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec          Tue Dec 22 16:54:25 2015

=====
                                Multicast Stats - GlobalRouter
=====
GroupAddress                SourceCounter    IngressPackets    IngressBytes    AverageSize
-----
ff05:0:0:0:0:0:0:1          1                2446250           2935500000      1200
                                MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec          Tue Dec 22 16:54:31 2015
```

```

=====
Multicast Stats - GlobalRouter
=====
GroupAddress          SourceCounter    IngressPackets    IngressBytes      AverageSize
-----
ff05:0:0:0:0:0:1      1                2448947           2938736400        1200
MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Tue Dec 22 16:54:37 2015

=====
Multicast Stats - GlobalRouter
=====
GroupAddress          SourceCounter    IngressPackets    IngressBytes      AverageSize
-----
ff05:0:0:0:0:0:1      1                2452185           2942622000        1200
...
...
Switch:1#

```

Monitor the IPv6 multicast route statistics for a maximum of 10 group IPv6 addresses. The statistics are monitored at intervals of 5 seconds for a duration of 300 seconds.

The following example displays monitoring output from three consecutive intervals.

```

Switch:1#monitor IPv6 mroute stats
FF05::1,FF05::2,FF05::3,FF05::4,FF05::5,FF05::6,FF05::7,FF05::8,FF05::9,FF05::a
MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Tue Dec 22 17:04:55 2015

=====
Multicast Stats - GlobalRouter
=====
GroupAddress          SourceCounter    IngressPackets    IngressBytes      AverageSize
-----
ff05:0:0:0:0:0:1      1                2768926           3322711200        1200
ff05:0:0:0:0:0:2      1                2768925           3322710000        1200
ff05:0:0:0:0:0:3      1                2768925           3322710000        1200
ff05:0:0:0:0:0:4      1                2768925           3322710000        1200
ff05:0:0:0:0:0:5      1                2768925           3322710000        1200
ff05:0:0:0:0:0:6      1                2768923           3322707600        1200
ff05:0:0:0:0:0:7      1                2768923           3322707600        1200
ff05:0:0:0:0:0:8      1                2768923           3322707600        1200
ff05:0:0:0:0:0:9      1                2768923           3322707600        1200
ff05:0:0:0:0:0:a      1                2768923           3322707600        1200
MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Tue Dec 22 17:05:01 2015

=====
Multicast Stats - GlobalRouter
=====
GroupAddress          SourceCounter    IngressPackets    IngressBytes      AverageSize
-----
ff05:0:0:0:0:0:1      1                2771625           3325950000        1200
ff05:0:0:0:0:0:2      1                2771625           3325950000        1200
ff05:0:0:0:0:0:3      1                2771625           3325950000        1200

```

```

ff05:0:0:0:0:0:0:4      1      2771624      3325948800    1200
ff05:0:0:0:0:0:0:5      1      2771624      3325948800    1200
ff05:0:0:0:0:0:0:6      1      2771622      3325946400    1200
ff05:0:0:0:0:0:0:7      1      2771622      3325946400    1200
ff05:0:0:0:0:0:0:8      1      2771622      3325946400    1200
ff05:0:0:0:0:0:0:9      1      2771622      3325946400    1200
ff05:0:0:0:0:0:0:a      1      2771622      3325946400    1200

MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec      Tue Dec 22 17:05:07 2015

=====
Multicast Stats - GlobalRouter

=====
GroupAddress      SourceCounter      IngressPackets      IngressBytes      AverageSize
-----
ff05:0:0:0:0:0:0:1      1      2774864      3329836800    1200
ff05:0:0:0:0:0:0:2      1      2774863      3329835600    1200
ff05:0:0:0:0:0:0:3      1      2774863      3329835600    1200
ff05:0:0:0:0:0:0:4      1      2774863      3329835600    1200
ff05:0:0:0:0:0:0:5      1      2774863      3329835600    1200
ff05:0:0:0:0:0:0:6      1      2774861      3329833200    1200
ff05:0:0:0:0:0:0:7      1      2774861      3329833200    1200
ff05:0:0:0:0:0:0:8      1      2774861      3329833200    1200
ff05:0:0:0:0:0:0:9      1      2774861      3329833200    1200
ff05:0:0:0:0:0:0:a      1      2774861      3329833200    1200

...
...

Switch:1#

```

Variable definitions

Use the data in the following table to use the **monitor ipv6 mroute stats** command:

Variable	Definition
WORD<7-400> {Ipv6address[, Ipv6address] [, ...]}	Specifies the multicast group IP address for which to monitor statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses, up to a maximum of 10.

Multicast route statistics configuration using EDM

The following sections provide procedural information you can use to configure multicast route statistics using the Enterprise Device Manager (EDM).

Enable IP Multicast Route Statistics

About This Task

Use this procedure to enable IP multicast route statistics.



Note

When you enable or clear IP multicast route statistics on the Controller node of a DVR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

Procedure

1. In the navigation pane, expand **Configuration > IP**
2. Select **Multicast**.
3. Select the **Globals** tab.
4. In **StatsEnabled**, select the option to enable or disable the collection of statistics.
5. (Optional) To clear the statistics, select **StatsClear**.
6. Select **Apply**.

Globals Field Definitions

Use the data in the following table to use the **Globals** tab.

Field	Description
StatsEnabled	Displays whether the multicast route statistics is enabled.
StatsClear	Clears the multicast route statistics.

View IP Multicast Route Statistics

Use this procedure to view IP multicast route statistics.

Before You Begin

- You must enable the collection of multicast statistics.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Stats** tab to view the statistics.

Stats Field Definitions

Use the data in the following table to use the **Stats** tab.

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.
SourceCount	Displays the number of sources associated with the multicast route record.
Pkts	Displays the number of packets received for the associated IP address.
Bytes	Displays the number of bytes received for the associated IP address.
AverageSizePerPkt	Displays the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the following formula: ingress packet/ingress byte.

Enable IPv6 Multicast Route Statistics

About This Task

Enable the collection of IPv6 multicast route statistics.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 Mroute**.
3. Select the **Globals** tab.
4. In **StatsEnabled**, select the option to enable or disable the collection of statistics.
5. (Optional) To clear the statistics, select **StatsClear**.
6. Select **Apply**.

Globals Field Definitions

Use the data in the following table to use the **Globals** tab.

Field	Description
StatsEnabled	Displays whether the multicast route statistics is enabled.
StatsClear	Clears the multicast route statistics.

View IPv6 Multicast Route Statistics

Use this procedure to view IPv6 multicast route statistics.

Before You Begin

- You must enable the collection of multicast statistics.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 Mroute**.
3. Select the **Stats** tab to view the statistics.

Stats Field Definitions

Use the data in the following table to use the **Stats** tab.

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.
SourceCount	Displays the number of sources associated with the multicast route record.
Pkts	Displays the number of packets received for the associated IP address.
Bytes	Displays the number of bytes received for the associated IP address.
AverageSizePerPkt	Displays the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the following formula: ingress packet/ingress byte.



IP Multicast over Fabric Connect

[IP Multicast over Fabric Connect Basic Configuration](#) on page 1947

[Layer 2 VSN Configuration for IP Multicast over Fabric Connect](#) on page 1978

[IP Shortcuts Configuration for IP Multicast over Fabric Connect](#) on page 1997

[Layer 3 VSN Fundamentals for IP Multicast over Fabric Connect](#) on page 2011

[IP SPB Multicast Policy Configuration](#) on page 2034

The topics in this section provide conceptual and procedural information about IP Multicast over Fabric Connect.

IP Multicast over Fabric Connect Basic Configuration

Table 154: IP Multicast over Fabric Connect product support

Feature	Product	Release introduced
IP Multicast over Fabric Connect	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
IP Multicast config-lite for Fabric Connect Note: There is no limit on VRF membership.	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8

Table 154: IP Multicast over Fabric Connect product support (continued)

Feature	Product	Release introduced
IP SPB Multicast Policy	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8
TTL handling of bridged traffic with routed SPB IP Multicast	4220 Series	Fabric Engine 9.4
	5320 Series	Fabric Engine 9.4
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.4
	5720 Series	Fabric Engine 9.4
	7520 Series	Fabric Engine 9.4
	7720 Series	Fabric Engine 9.4
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 9.4
	VSP 7400 Series	VOSS 9.4
Universal Plug and Play (uPnP) Filtering	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

IP Multicast over Fabric Connect Fundamentals

This section provides information about IP Multicast over Fabric Connect.

IP Multicast over Fabric Connect

Extreme Networks is leading the industry with a new approach to transporting IP multicast using IP Multicast over Fabric Connect. IP Multicast over Fabric Connect greatly simplifies multicast deployment, with no need for any multicast routing protocols such as Protocol Independent Multicast-Sparse Mode (PIM-SM) or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM). A BEB can forward a multicast stream anywhere in an SPBM network where IS-IS advertises the stream to the rest of the fabric.

The advantage of this solution over traditional approaches is the simplicity in provisioning and deploying IP multicast bridging and routing. Also, due to the fact that only one control plane protocol (IS-IS) exists, convergence times in the event of a network failure, are typically sub second.

You can compare the quick convergence times for IP Multicast over Fabric Connect to Interior Gateway Protocols like Open Shortest Path First (OSPF) combined with PIM-SM or PIM-SSM. OSPF combined with PIM-SM or PIM-SSM can have recovery times that are sub optimal with convergence times that take tens of seconds. PIM experiences longer convergence times, in part, because unicast IP routing protocols must converge before PIM can converge. PIM also maintains the network state for every multicast group and uses a mechanism based on each hop to update the network about state changes, which affects scalability.

IP Multicast over Fabric Connect is extremely scalable because you only apply the multicast bridging and routing functionality at the SPBM fabric edge, with the streams mapped to SPBM multicast trees in the fabric.

IP Multicast over Fabric Connect introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. IP Multicast over Fabric Connect uses these extensions, along with the Internet Group Management Protocol (IGMP) Snooping and Querier functions at the edge of the SPBM cloud, to create sub-trees of the VSN SPB for each multicast group to transport IP multicast data.

With IP Multicast over Fabric Connect, the switch supports the following:

- Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network (Layer 2 VSN with IP Multicast over Fabric Connect). Example application: Multicast in data centers.
- IP multicast routing support for IP Shortcuts using SPBM in the core and IGMP on the access (IP Shortcuts with IP Multicast over Fabric Connect). Example applications: Video surveillance, TV/Video/Ticker/Image distribution, VX-LAN.
- Layer 3 Virtual Services Network with VRF based routing support for IP Multicast over Fabric Connect in the core and IGMP on the access (Layer 3 VSN with IP Multicast over Fabric Connect). Example applications: Video surveillance, TV/Video/Ticker/Image Distribution, VX-LAN, Multi-tenant IP multicast.

IP Multicast over Fabric Connect and Universal Plug and Play Filtering

When multicast packets are received on IGMP-enabled interfaces and the multicast group matches the range of groups to be filtered, Universal Plug and Play (uPnP) Filtering drops them.

For more information, see [Universal Plug and Play Filtering](#) on page 1957.

How IP Multicast over Fabric Connect Works

The BEBs act as the boundary between the multicast domain (currently only IGMP dynamic or static) and the SPBM domain. Multicast senders (sources) and receivers connect directly or indirectly (using Layer 2 switches) to the BEBs. You can enable IP Multicast over Fabric Connect services at the Layer 2 VSN level or the Layer 3 VSN level (including the GRT).

The following figure shows how multicast senders and receivers connect to the SPBM cloud using BEBs.




Figure 142: IP Multicast over Fabric Connect streams

The following list describes how multicast senders and receivers connect to the SPBM cloud using BEBs in the preceding diagram:

1. The sender transmits multicast traffic with group IP address 233.252.0.1.
2. After the BEB receives the IP multicast stream from the sender, the BEB allocates data I-SID 16000001 for the S,G multicast stream. The BEB sends an LSP with the TLV 185 (for Layer 2 VSN multicast and Layer 3 VSN multicast) or TLV 186 (for IP Shortcuts multicast) with the transmit bit set. The BEB also sends an IS-IS service identifier and

- unicast address sub-TLV (where the unicast address has the multicast bit set and the I-SID is the Data I-SID).
3. The receiver sends a join request to Group 233.252.0.1.
 4. The BEB (acting as the IGMP Querier) queries the IS-IS database to find all senders for group 233.252.0.1. If the group exists, the BEB sends an LSP with the IS-IS service identifier and unicast address sub-TLV (where the unicast address has the multicast bit set and the nickname is the stream transmitter BEB and the I-SID is the data I-SID).
 5. The multicast tree is calculated for the data I-SID and the data starts flowing from the sender.

Scope Level

IP Multicast over Fabric Connect constrains all multicast streams within the level in which they originate, which is called the scope level. In other words, if a sender transmits a multicast stream to a BEB on a C-VLAN (a VLAN that is mapped to an I-SID, for instance, a Layer 2 VSN) with IP Multicast over Fabric Connect enabled, only receivers that are part of the same Layer 2 VSN can receive that stream. Similarly, if a sender transmits a multicast stream to a BEB on a VLAN that is part of the GRT or a Layer 3 VSN with IP Multicast over Fabric Connect enabled, only receivers that are part of the same Layer 3 instance (GRT or L3 VSN) can receive that stream.



Note

In the context of IP Multicast over Fabric Connect, scope is either the Global Routing Table or the I-SID value of the Layer 2 or Layer 3 VSN associated with the local VLAN on which the IP multicast data was received.

Data I-SID

After the BEB receives the IP multicast stream from the sender, a BEB allocates a data Service Identifier (I-SID) in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the S,G, V tuple, which is the source IP address, the group IP address, and the local VLAN the multicast stream is received on.

The BEB propagates this information through the SPBM cloud by using IS-IS TLV updates in LSPs, which results in the creation of a multicast tree for that stream. All BEBs now use the data I-SID for that stream and its scope. The data I-SID is a child of the scope or VSN I-SID. If no receiver requests the IP multicast stream, the ingress BEB does not forward the multicast stream.

IGMP

After a BEB receives an IGMP join message from a receiver, a BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the requested stream exists, the BEB sends an IS-IS TLV update to its neighbors to inform them of the presence of a receiver, and this information is propagated through the SPBM cloud.

IS-IS acts dynamically using the TLV information it receives from BEBs that connect to the sender and the receivers to create a multicast tree between them. IS-IS creates very

efficient multicast trees for the data I-SID allocated at the sender edge of the SPBM cloud to transport data between the sender and the receivers. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID. After IS-IS creates the multicast tree, the sender transports data to the receiver across the SPBM cloud using the data I-SID.

The trigger to send IS-IS updates to announce a multicast stream into the SPBM cloud is the multicast traffic arriving at the BEB. Because the BEB only interacts with IGMP and not PIM, all multicast traffic must be drawn towards the BEB for the stream to be announced, which SPBM accomplishes by making the BEB an IGMP Querier. In a VLAN, the IGMP Querier sends out periodic IGMP queries.

**Note**

The BEB must be the only IGMP Querier in the VLAN. If the BEB receives an IGMP query from any other device, it causes unexpected behavior, including traffic loss.

IP Multicast config-lite for Fabric Connect

With IP Multicast config-lite for Fabric Connect you can configure Layer 3 IP Multicast routing over Fabric Connect using a Layer 2 edge node without associating an IP address on the VLAN. When you configure an MVPN I-SID on a VLAN, a dynamic VRF named **13vsn** is created, and **ipvpn** is configured.

**Note**

ipvpn enable is not configured on this VRF.

Configure **ip spb-multicast enable** on a VLAN interface without an IP address to avoid depleting the IP address pool of the network.

**Note**

Manually configuring Layer 3 VSNs is not available.

You can convert a dynamic VRF to a static VRF by using typical VRF creation with the VRF name. For more information, see [Create a VRF Instance](#) on page 4081.

Restrictions

The following list provides restrictions related to this feature:

- If a VRF with the same I-SID already exists on the device, the VLAN is associated with the existing VRF after you enter **mvpn-isid** on the VLAN.
- You cannot have an IP address and **mvpn-isid** configured on the same VLAN at the same time.
- Dynamic VRFs are not present in the running configuration and are not saved.

BEB as IGMP Querier

The BEB acts as the IGMP Querier and creates tables for links that need IP multicast streams. IGMP and IGMP Snooping cannot work without an IGMP Querier that sends out periodic IGMP queries.

The BEB only interacts with IGMP messages and not PIM. All multicast traffic must enter the BEB for the data stream to be announced.

The BEB must be the only IGMP Querier in the VLAN. If the BEB receives an IGMP query from any other device, unexpected behavior results, including traffic loss.

The IGMP query message is an IP packet and requires a source IP address. However, Layer 2 IGMP Snooping with SPBM by default turns on the service without the configuration of an IP address on the VLAN. By default, the BEB sends an IGMP query message with an IP source address of 0.0.0.0. If there are interoperability issues with third party vendors as a result of the 0.0.0.0 IP address, then you can configure the querier address under IGMP, without having to configure an IP address for the Layer 2 VSN VLAN.

IGMP Snooping, operating on the Layer 2 VSN, listens to conversations between hosts and routers, and maintains a table for links that need IP multicast streams.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

For more conceptual and configuration information on IGMP, see [IP Multicast](#) on page 1696.

Switch Clustering at the Edge of the SPBM Network

Typical customer deployments require redundancy all the way to the access side of the network. IP Multicast over Fabric Connect supports switch clustering, Split Multilink Trunking (SMLT) technology, at the edge of the SPBM fabric, providing redundancy to the access Layer 2 switch where you can attach multicast senders and receivers. Typical SPBM fabric deployments use two or more B-VLANs for Equal Cost Multipath (ECMP) and resiliency. For simplicity in understanding how the SPBM network works, assume that there are two B-VLANs (primary and secondary).

The following figure shows how multicast senders and receivers connect to the SPBM cloud using BEBs.




Figure 143: IP Multicast over Fabric Connect streams in an SMLT configuration

The following list describes the preceding diagram:

1. The edge switch hashes the sender multicast data to a specific MLT link.
2. A multicast stream received at the edge of the SPBM fabric is mapped to a dedicated multicast data I-SID.
3. For the non-SMLT attached sender 2, the stream is hashed to the primary or secondary B-VLAN based on whether the data I-SID is even or odd numbered. For the SMLT attached to sender 1, IS-IS advertises the stream to the rest of the fabric on the primary B-VLAN and synchronizes information to the vIST peer.
4. The edge switch hashes the receiver IGMP join to a specific MLT link.
5. Both BEBs on both B-VIDs advertise the IGMP join.
6. The multicast tree is built for (S1,G1), which is rooted in the primary sender BEB. The multicast tree is built for (S1,G1), which is rooted in the secondary sender BEB.

IGMP Snooping is widely used on Layer 2 access switches to prune multicast traffic. In IP Multicast over Fabric Connect, BEBs are the IGMP Queriers, therefore access switches forward multicast data from the senders as well as IGMP control messages from receivers to the BEBs.

When a sender transmits multicast data to the Layer 2 access switch that has an MLT to the switch cluster, it is hashed towards one or the other BEBs in the switch cluster. The receiving BEB allocates a data I-SID and sends a TLV update on either the primary B-VLAN or the secondary B-VLAN, depending on whether the BEB is the primary or secondary switch. The primary switch uses the primary B-VLAN, whereas, the secondary switch uses the secondary B-VLAN. This information is propagated through the SPBM fabric so all BEBs are aware of this stream availability.

The sender information is also synchronized over the vIST to the peer switch. Then the peer switch allocates a data I-SID for the multicast stream and sends a TLV update on the appropriate B-VLAN to announce the availability of the stream. The data I-SIDs allocated by the primary and secondary switch cluster peers may be the same or different, as they are allocated independently by each switch.

**Note**

If a sender attaches to only one BEB in a switch cluster, the sender information is not synchronized over the vIST because it is not SMLT attached. The sender information is advertised, and data is sent on either the primary or secondary B-VLAN. The odd-numbered data I-SIDs use the primary B-VLAN, and the even-numbered data I-SIDs use the secondary B-VLAN. The same hashing rules apply to the forwarding of multicast data.

When a receiver sends an IGMP join message to the Layer 2 access switch that has an MLT to the switch cluster, it is hashed towards one or the other BEBs in the switch cluster. The receiving BEB queries the IS-IS Link State Database (LSDB) to check if a sender exists for the requested stream within the scope of the receiver.

If the requested stream does not exist, the BEB keeps the IGMP information but no further action is taken. If the requested stream exists, the BEB sends an IS-IS Link State Packet (LSP), with TLV update information, for both primary and secondary B-VLANs to its neighbors to inform them of the presence of a receiver. The BEB propagates this information through LSPs through the SPBM cloud. The receiver information is also synchronized over the vIST to the peer switch. The peer switch then queries its IS-IS Link State Database (LSDB) and, if the requested stream exists, it sends an IS-IS LSP, with a TLV update, for both primary and secondary B-VLANs to its neighbors to inform them of the presence of the receiver.

IS-IS uses these TLV updates in LSPs to create multicast shortest path first trees in the SPBM fabric. IS-IS creates a shortest path first tree for the primary and secondary B-VLANs, but only one of the B-VLANs transports multicast data with the other in active standby in case of failures at the SPBM edge. After IS-IS creates the trees, multicast data flows between senders and receivers.

IP Multicast over Fabric Connect and SMLT

The following section summarizes the IP Multicast over Fabric Connect actions in an SMLT environment. The BEBs on the sender side function as follows:

- Primary SMLT peer BEB always advertises the streams it receives, and sends data for them on the primary B-VLAN.
- Secondary SMLT peer BEB always advertises the streams it receives, and sends data for them on the secondary B-VLAN.
- Non-SMLT BEBs or SMLT BEBs with single attached senders advertise streams, and send data on the primary or secondary B-VLAN based on hash criteria (odd-numbered data I-SIDs use primary B-VLAN; even-numbered data I-SIDs use secondary B-VLAN).

The BEBs on the receiver side function as follows:

- The primary SMLT peer BEB that receives multicast data on the primary B-VLAN sends it to both SMLT and non-SMLT SPBM access (UNI) links.
- The primary SMLT peer BEB that receives multicast data on the secondary B-VLAN sends it to non-SMLT SPBM access (UNI) links only.
- The secondary SMLT peer BEB that receives multicast data on primary B-VLAN sends it to non-SMLT SPBM access (UNI) links only.
- The secondary SMLT peer BEB that receives multicast data on secondary B-VLAN sends data to both SMLT and non-SMLT SPBM access (UNI) links.
- The non-SMLT BEB that receives multicast data on primary or secondary B-VLAN sends data to all SPBM access (UNI) links.

Layer 2 Querier behavior for a switch cluster

For C-VLANs in an SMLT environment, the vIST ports are not part of the VLAN.

IGMP on a C-VLAN functions as follows to account for the fact that vIST peers do not see the membership queries of each other:

- The vIST peer with the higher IP address sends the queries out all SMLT and non-SMLT ports on SPBM access links.
- The vIST peer with the lower IP address only sends out queries on its non-SMLT ports. This includes SMLT ports whose remote ports are down (SMLT state of 'norm').
- With the existence of an vIST peer with a higher IP address and an vIST peer with a lower IP address, it means two queriers exist within the C-VLAN. Having two queriers poses no problems in this SPB environment, as all SMLT access devices see the vIST peer with the higher IP address as the querier, and non-SMLT access devices see the directly connected vIST peer as the querier. Non-SMLT access devices that connect on either side of the vIST peers can talk to each other using the SPBM cloud.


SPBM Multicast FIB

In SPBM, BMAC addresses are carried within the IS-IS link-state database. To do this, SPBM supports an IS-IS TLV that advertises the I-SID and BMAC information across the network. Each node has a System ID, which also serves as Backbone MAC address (BMAC) of the switch. These Backbone MAC addresses are populated into the SPBM VLAN Forwarding Information Base (FIB).

SPBM runs all pair Dijkstras to produce the multicast FIB. The computing node loops through each node to run Dijkstra using that node as the root, and then prunes paths to only keep the shortest paths. The computing node then computes the intersection of the set of I-SIDs for which the root node transmits, with the set of I-SIDs for which the path endpoints receive.

When the network topology is discovered and stored in the IS-IS link-state database, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes.

The multicast addresses are built out of two pieces: the SPBM Node Nickname and the I-SID ID converted to hexadecimal format to form the multicast MAC address.



For example, if the nickname is 0.00.10 and the I-SID is 100 (0x64), the multicast address is 03:00:10:00:00:64.

The multicast FIB is not produced until virtual services are configured and learned.

TLVs

This information is propagated through the SPBM cloud using IS-IS Link State Packets (LSPs), which carry TLV updates, that result in the multicast tree creation for that stream. The LSPs carry I-SID information and information about where IP multicast stream senders and receivers exist using the following TLVs:

- For Layer 2 VSNs and Layer 3 VSNs, using TLV 144 and TLV 185.
- For IP Multicast over Fabric Connect within the GRT, using TLV 144 and TLV 186.
- For bridged multicast forwarding based on Layer 2 VSN assignments, using TLV 188.

TLV 144 on the BEB bridge, where the sender is located, has the Tx bit set. All BEB bridges, where a receiver exists, have the Rx bit set.

TLV 185 on the BEB where the source is located, displays the multicast source and group addresses and has the Tx bit set.

TLV 186 on the BEB where the source is located displays the multicast source and group addresses and have the Tx bit set.

TLV 188 on the BEB where the sender is located advertises a mapping between the sender Layer 2 VSN I-SID and the SPB IP Multicast data I-SID for each multicast stream. Receiving BEBs use this information to determine whether to forward the multicast traffic using bridged or routed forwarding.

Each multicast group has its own unique data I-SID with a value between 16,000,000 to 16,512,000.

IS-IS acts dynamically using the TLV information received from BEBs that connect to the sender and the receivers to create a multicast tree between them.

Universal Plug and Play Filtering

The switch can filter multicast packets destined for the Universal Plug and Play (uPnP) multicast IP address with a uPnP Filtering option. uPnP Filtering drops all incoming multicast packets received by a switch on an IGMP-enabled interface if the destination multicast IP address matches the configured range.

uPnP Filtering applies to both multicast receivers and multicast senders. If you want to use the uPnP Filtering address range for actual multicast streaming, you must disable uPnP Filtering on the IGMP interface.

uPnP Filtering is disabled by default. If you create a new IGMP interface, uPnP Filtering is enabled automatically on the interface for the destination multicast IP address range 239.255.255.250/32. If you enable uPnP Filtering on an existing IGMP-enabled interface with senders and receivers already present, the filter does not delete the existing senders or receivers; the filter begins to drop packets from that point forward. Existing senders and receivers eventually expire and senders are not relearned.

You can use CLI or EDM to configure a different destination multicast IP address range.

uPnP Filtering functions in the following scenarios:

- IGMP snooping is enabled on a VLAN and IP Multicast over Fabric Connect on Layer 2 VSN is configured.
- IP Multicast over Fabric Connect within the Global Routing Table (GRT) is configured.
- IP Multicast over Fabric Connect on a Layer 3 VSN is configured.

Considerations when you connect an IP Multicast over Fabric Connect network to a PIM network

IP Multicast over Fabric Connect does not integrate PIM functionality. Apply the following considerations when you connect to a PIM network:

- You must configure static IGMP receivers on the BEB access interface that faces the PIM network when the sender is on the SPBM access network and the receiver is on the PIM network.



Note

The PIM router must have a configuration option to accept streams with non-local sources or the router drops the packets. The switch does not support a configuration option to accept streams with non-local sources.

You must configure static IGMP receivers on the PIM interface that face the IP Multicast over Fabric Connect network when the sender is on the PIM network and the receiver is on the SPBM access network.



Note

For security reasons and to limit unnecessary multicast streams from being injected into the SPBM domain, you should configure ACLs on the BEB facing the PIM network.

IP Multicast over Fabric Connect Restrictions

Review the following restrictions for the IP Multicast over Fabric Connect feature.

DvR

You cannot configure IP Multicast over Fabric Connect on DvR Leaf nodes. You cannot deploy IP Multicast over Fabric Connect, statically or dynamically, using EAP on a DvR Leaf. This configuration is pushed from the DvR Controllers in the domain.

IGMP

The BEB must be the only IGMP querier in the network. If the BEB receives an IGMP query from any other device, it causes unpredictable behavior, including traffic loss.

SPBM supports IGMP Snooping on a C-VLAN, but it does not support PIM on a C-VLAN. If you enable IGMP Snooping on a C-VLAN, then its operating mode is Layer 2 VSN with IP Multicast over Fabric Connect.

SPBM supports Network Load Balancing (NLB) unicast and multicast modes. SPBM does not support NLB Multicast operation with IGMP.

**Note**

The NLB Multicast operation feature is not supported on all hardware platforms. For more information about feature support, see [VLAN Feature Support](#) on page 4005.

You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is either the same as the IGMP version configured on the IGMP Snooping VLAN, or that compatibility mode is enabled.

SSM

If you delete any **ssm-map** in a static range group, the switch deletes the entire static range group. For example, create an ssm-map for 232.122.122.122 to 232.122.122.128 and after that configure this same range in a static group. If you delete any ssm-map between 232.122.122.122. to 232.122.122.128, the switch deletes the entire static range group.

PIM

There can be no interaction with PIM and multicast routers on the access.

The BEB only interacts with IGMP messages and not PIM, so all multicast traffic must be drawn towards the BEB, which acts as the IGMP querier, for the stream to be announced.

IP Multicast over Fabric Connect does not integrate PIM functionality so the following considerations apply when connecting to a PIM network:

- You must configure static IGMP receivers on the BEB access interface facing the PIM network when the sender is on the SPBM access network and the receiver is on the PIM network. Static IGMP receivers make the PIM router accept streams and avoid a Reverse Path Forwarding (RPF) check that can change the source of the stream.
- You must configure static IGMP receivers on the PIM interface facing the IP Multicast over Fabric Connect network when the sender is on the PIM network and the receiver is on the SPBM access network.
- You must configure Access Control Lists (ACLs) on the BEB facing the PIM network for security.

Data I-SID

The BEB matches a single multicast stream to a particular data I-SID. As a result there is a one-to-one mapping between the S,G to data I-SID for each BEB.

Supported Services

The switch does not support IP Multicast over Fabric Connect routing on inter-VSN routing interfaces.

The switch supports the following modes of IP Multicast over Fabric Connect:

- Layer 2 VSN multicast service — Multicast traffic remains within the same Layer 2 VSN across the SPBM cloud.
- Layer 3 VSN multicast service — Multicast traffic remains within the same Layer 3 VSN across the SPBM cloud.
- IP Shortcuts multicast service — Multicast traffic can cross VLAN boundaries but remains confined to the subset of VLANs with the Global Routing Table that have IP Multicast over Fabric Connect enabled.

IP Multicast over Fabric Connect Configuration using the CLI

This section provides procedures to configure IP Multicast over Fabric Connect using the command line interface (CLI).

Enable IP Multicast over Fabric Connect Globally

Before You Begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.

About This Task

Use this procedure to enable IP Multicast over Fabric Connect globally on the Backbone Edge Bridges (BEBs) that directly, or indirectly (using Layer 2 switches), connect to IP multicast senders or receivers. By default, IP Multicast over Fabric Connect is disabled. You do not need to enable IP Multicast over Fabric Connect on the Backbone Core Bridges (BCBs).

You must configure IP Multicast over Fabric Connect globally, and then enable it on the service option or options you choose.

You must enable IP multicast over Fabric Connect globally on all DvR enabled nodes (Controllers and Leaf nodes) in a DvR domain.



Note

IP Multicast over Fabric Connect uses I-SIDs starting at 16,000,000 and above. If Layer 2 or Layer 3 I-SIDs are in this range, the system displays an error message and the switch does not enable IP Multicast over Fabric Connect.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Verify no I-SIDs exist in the default reserved range:

- a. For Layer 2 use the following command:

```
show vlan i-sid
```

- b. For Layer 3 use the following command:

```
show ip ipvpn vrf WORD<1-16>
```

3. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

4. Enable IP Multicast over Fabric Connect globally:

```
spbm <1-100> multicast enable
```

**Note**

The switch only supports one SPBM instance.

Example

Enable IP Multicast over Fabric Connect globally:

```
Switch:1>show vlan i-sid
=====
Vlan I-SID
=====
VLAN_ID      I-SID      VLAN_NAME      I-SID NAME
-----
1
10           100        VLAN-10         Hospital-Server-10
90           1000       VLAN-90         ISID-1000
3 out of 3 Total Num of Vlans displayed
```

Variable definitions

The following table defines parameters for the **spbm** command.

Variable	Value
<1-100>	Enables IP Multicast over Fabric Connect globally. The default is disabled. Specifies the SPBM instance. The switch only supports one instance.

Display IP Multicast over Fabric Connect information

Use this procedure to display IP Multicast over Fabric Connect summary information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the status of the global IP Multicast over Fabric Connect configuration:

```
show isis spbm multicast
```

3. Display IP Multicast over Fabric Connect summary information for each S, G, V tuple:

```
show isis spb-mcast-summary [count] [host-name WORD<0-255>] [lspid
<xxxx.xxxx.xxxx.xx-xx>]
```

4. Display information about the multicast routes on the switch:

```
show ip mroute route [vrf WORD<1-32>] [vrfrids WORD<0-255>]
```

Example

Display IP Multicast over Fabric Connect global configuration information:

```
Switch:1>show isis spbm multicast
```

```
multicast : enable
fwd-cache-timeout(seconds) : 210
routed-multicast-ttl1-bridging : disable
```

```
Switch:1#show isis spb-mcast-summary
```

```
=====
                        SPB multicast - Summary
=====
SCOPE      SOURCE      GROUP      DATA      LSP  HOST
I-SID      ADDRESS      ADDRESS      I-SID  BVID  FRAG  NAME
-----
GRT        192.0.2.102    233.252.0.1    16000001  63   0x0   DIST5A
Switch:1>show ip mroute route

=====
                        Mroute Route - GlobalRouter
=====
GROUP      SOURCE      SRCMASK      UPSTREAM_NBR  IF      EXPIR  PROT
-----
233.252.0.1    0.0.0.0      0.0.0.0      0.0.0.0      V3      30    spb-access
233.252.0.1    192.0.2.102  255.255.255.0  0.0.0.0      -       0     spb-network
233.252.0.2    0.0.0.0      0.0.0.0      0.0.0.0      V2      30    pimsm
225.1.1.1      198.51.100.99 255.255.255.0  0.0.0.0      V3      173   spb-pim-gw

Total 4
```

Variable Definitions

The following table defines parameters for the **show isis spb-mcast-summary** command.

Variable	Value
<i>count</i>	Displays the total number of SPB multicast entries.
<i>host-name WORD<0-255></i>	Displays the IP Multicast over Fabric Connect summary information for a specific host-name.
<i>lspid <xxx.xxx.xxx.xx-xx></i>	Displays the IP Multicast over Fabric Connect summary information for the specified LSP ID that you enter in xxx.xxx.xxx.xx-xx — 8 byte format.

The following table defines parameters for the **show ip mroute route** command.

Variable	Value
<code>vrf WORD<1-32></code>	Specifies a VRF.
<code>vrfids WORD<0-255></code>	Specifies the VRF ID

Configure IP Multicast config-lite for Fabric Connect

Before You Begin

Create a VLAN. For more information see, [VLAN Configuration](#) on page 3981

About This Task

Use **mvpn-isid** to dynamically create a VRF on one of the reserved Layer 3 VSN VLANs. The dynamically created VRF, **l3vsn-<I-SID number>**, obtains the I-SID number from the **mvpn-isid** command. You can use the **ip spb-multicast enable** command on that VLAN without configuring an IP address.



Note

If you are using a DVR leaf instance, the VRF name is **l3vsn-<l3isid>**.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```
2. Configure an MVPN I-SID on the VLAN:

```
mvpn-isid <0-15999999>[0-511]
```
3. (Optional) Configure a routed SPB querier source IP address:

```
ip igmp routed-spb-querier-addr {A.B.C.D}
```
4. Enable SPB Multicast on the VLAN:

```
ip spb-multicast enable
```

Example

Enable SPB Multicast for a VLAN:

```
Switch:1>enable
Switch:1#config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 10
Switch:1(config-if)#mvpn-isid 100
Switch:1(config-if)#ip igmp routed-spb-querier-addr 192.0.2.1
Switch:1(config-if)#ip spb-multicast enable
```

```
Switch:1#show ip igmp interface vrf l3vsn-100
=====
                        Igmp Interface - VRF l3vsn-100
=====
  QUERY      OPER      QUERY  WRONG      LASTMEM      SCOPE
IF  INTVL STATUS VERS.  VERS  QUERIER  MAXRSPT QUERY JOINS ROBUST QUERY  MODE      L2ISID  ISID
-----
```

```
V10    125    inact  2      2      0.0.0.0    100    0      0      2      10      routed-spb 0      100
```

Variable Definitions

The following table defines parameters for the **mvpn-isid** command.

Variable	Definition
<0-15999999>[0-511]	Specifies the I-SID. The MAC offset value is optional. Note: Use <code>mvpn-isid 0</code> to enable spb-multicast for a VLAN in the Global Routing Table.

The following table defines parameters for the **ip igmp routed-spb-querier-addr** command.

Variable	Definition
{A.B.C.D}	The SPB querier address. The default is 0.0.0.0.

Display the Multicast FIB

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the SPBM multicast FIB:
show isis spbm multicast-fib [vlan <1-4059>] [i-sid <1-16777215>]
[nick-name <x.xx.xx>] [summary]

Example

```
Switch:1#show isis spbm multicast-fib detail
=====
SPBM MULTICAST FIB ENTRY DETAIL INFO
=====
```

MCAST DA	ISID	BVLAN	SYSID	HOST-NAME	OUTGOING-INTERFACES	INCOMING INTERFACE	CVLAN	AREA	AREA-NAME
03:00:20:00:00:65	101	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	101	HOME	area-10.0010
03:00:20:00:00:67	103	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	103	HOME	area-10.0010
03:00:20:00:00:69	105	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	105	HOME	area-10.0010
03:00:20:00:00:6b	107	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	107	HOME	area-10.0010
03:00:20:00:00:6d	109	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	109	HOME	area-10.0010
03:00:20:00:00:6f	111	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	111	HOME	area-10.0010
03:00:20:00:00:71	113	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	113	HOME	area-10.0010
03:00:20:00:00:73	115	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	115	HOME	area-10.0010
03:00:20:00:00:75	117	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	117	HOME	area-10.0010
03:00:20:00:00:77	119	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	119	HOME	area-10.0010
03:00:20:00:00:79	121	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	121	HOME	area-10.0010
03:00:20:00:00:7b	123	4051	0200.20ff.fff0	vn-area-20.0020	1/32	Virtual	123	HOME	area-10.0010
13:00:11:00:00:65	101	4051	a425.1b53.7084	BEB1-7200	1/32	MLT-2	101	HOME	area-10.0010

```
-----
13 out of 2354 Total Num of SPBM MULTICAST FIB entries displayed
Remote: 0 out of 0 Total Num of SPBM MULTICAST FIB entries displayed
-----
```

Variable Definitions

The following table defines parameters for the **show isis spbm multicast-fib** command.

Variable	Value
<i>vlan</i> <1-4059>	Displays the FIB for the specified SPBM VLAN.
<i>i-sid</i> <1-16777215>	Displays the FIB for the specified I-SID.
<i>nick-name</i> <x.xx.xx>	Displays the FIB for the specified nickname.
<i>summary</i>	Displays a summary of the FIB.

View TLV Information for IP Multicast over Fabric Connect

About This Task

Use the following commands to check TLV information.

The LSPs carry I-SID information and information about where IP multicast stream senders and receivers exist using the following TLVs:

- For Layer 2 and Layer 3 VSNs, using TLV 144 and TLV 185.
- For IP Multicast over Fabric Connect within the GRT, using TLV 144 and TLV 186.
- For bridged multicast forwarding based on Layer 2 VSN assignments, using TLV 188.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IS-IS Link State Database information by TLV:


```
show isis lsdb tlv <1-236> [sub-tlv <1-3>] [detail] [home|remote]
```
3. Display IS-IS Link State Database information by Link State Protocol ID:


```
show isis lsdb lspid <xxxx.xxxx.xxxx.xx-xx> tlv <1-236> [sub-tlv <1-3>] [detail] [home|remote]
```

Examples

Display TLV information for Layer 2 or Layer 3 VSNs:

```
Switch:1# show isis lsdb tlv 185 detail
=====
ISIS LSDB (DETAIL)
=====
-----
Level-1LspID: 000c.f803.83df.00-00 SeqNum: 0x000001ae Lifetime: 898
Chksum: 0xcebe PDU Length: 522
Host_name: Switch
Attributes: IS-Type 1
TLV:185 SPBM IPVPN :
VSN ISID:5010
BVID :10
    Metric:0
    IP Source Address: 192.0.2.1
    Group Address : 233.252.0.1
    Data ISID : 16300001
    TX : 1
```

```
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.3
Data ISID : 16300003
TX : 1
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.5
Data ISID : 16300005
TX : 1
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.7
Data ISID : 16300007
TX : 1
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.9
Data ISID : 16300009
TX : 1
LSVSN Mapping: false
VSN ISID:5010
BVID :20
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.2
Data ISID : 16300002
TX : 1
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.4
Data ISID : 16300004
TX : 1
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.6
Data ISID : 16300006
TX : 1
LSVSN Mapping: false
Metric:0
IP Source Address: 192.0.2.1
Group Address : 233.252.0.8
Data ISID : 16300008
TX : 1
Metric:0
LSVSN Mapping: false
IP Source Address: 192.0.2.1
Group Address : 233.252.0.10
Data ISID : 16300010
TX : 1
LSVSN Mapping: false

Switch:1# show isis lsdb lspid 000c.f803.83df.00-05 tlv 144 detail
=====
ISIS LSDB (DETAIL)
=====
-----
Level-1 LspID: 000c.f803.83df.00-00 SeqNum: 0x00000477 Lifetime: 903
```

```

Chksum: 0x200b PDU Length: 522
Host_name: Switch
Attributes: IS-Type 1
    Instance: 0
    Metric: 0
    B-MAC: 03-00-00-00-00-00
    BVID:10
    Number of ISID's:5
        16000001(Tx),16000003(Tx),16000005(Tx),16000007(Tx),16000009(Tx)
    Instance: 0
    Metric: 0
    B-MAC: 03-00-00-00-00-00
    BVID:20
    Number of ISID's:5
        16000002(Tx),16000004(Tx),16000006(Tx),16000008(Tx),16000010(Tx)

```

Display TLV information for IP Multicast over Fabric Connect within the GRT:

```

Switch:1# show isis lsdb tlv 186 detail
=====
ISIS LSDB (DETAIL)
=====
-----
Level-1 LspID: 000c.f803.83df.00-06 SeqNum: 0x000002eb Lifetime: 1113
Chksum: 0x7e3b PDU Length: 556
Host_name: Switch
Attributes: IS-Type 1
TLV:186 SPBM IP Multicast:
    GRT ISID
    Metric:0
    IP Source Address: 192.2.0.10
    Group Address : 233.252.0.1
    Data ISID : 16300012
    BVID : 20
    TX : 1
    LSVSN Mapping: false
    Route Type : Internal
    GRT ISID
    Metric:0
    IP Source Address: 192.2.0.10
    Group Address : 233.252.0.2
    Data ISID : 16300013
    BVID : 10
    TX : 1
    LSVSN Mapping: false
    Route Type : Internal
    GRT ISID
    Metric:0
    IP Source Address: 192.2.0.10
    Group Address : 233.252.0.3
    Data ISID : 16300014
    BVID : 20
    TX : 1
    LSVSN Mapping: false
    Route Type : Internal
    GRT ISID
    Metric:0
    IP Source Address: 192.2.0.10
    Group Address : 233.252.0.4
    Data ISID : 16300015
    BVID : 10
    TX : 1
    LSVSN Mapping: false
    Route Type : Internal

```

```

GRT ISID
Metric:0
IP Source Address: 192.2.0.10
Group Address : 233.252.0.5
Data ISID : 16300016
BVID : 20
TX : 1
LSVSN Mapping: false
Route Type : Internal
GRT ISID
Metric:0
IP Source Address: 192.2.0.10
Group Address : 233.252.0.6
Data ISID : 16300017
BVID : 10
TX : 1
LSVSN Mapping: false
Route Type : Internal
GRT ISID
Metric:0
IP Source Address: 192.2.0.10
Group Address : 233.252.0.7
Data ISID : 16300018
BVID : 20
TX : 1
LSVSN Mapping: false
Route Type : Internal

```

Variable Definitions

The following table defines parameters for the **show isis lsdb** command.

Variable	Value
<i>detail</i>	Displays detailed information about the IS-IS Link State database.
<i>home</i>	Displays the IS-IS LSDB information that the system configures in the home area.
<i>level {11, 12, 112}</i>	Displays information on the IS-IS level. The IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1. Level 2 function is disabled.
<i>local</i>	Displays information on the local LSDB.
<i>lspid <xxxx.xxxx.xxxx.xx-xx></i>	Specifies information about the IS-IS Link State database by LSP ID.
<i>remote</i>	Displays the IS-IS LSDB information that the system configures in the remote area.
<i>sub-tlv <1-3></i>	Specifies information about the IS-IS Link State database by sub-TLV.
<i>sysid <xxxx.xxxx.xxxx></i>	Specifies information about the IS-IS Link State database by System ID.
<i>tlv <1-236></i>	Specifies information about the IS-IS Link State database by TLV.

Configure Universal Plug and Play (uPnP) Filtering

Before You Begin

Create a port-based VLAN.

About This Task

Use the following procedure to enable Universal Plug and Play (uPnP) Filtering on an IGMP-enabled interface. uPnP Filtering is disabled by default.

The default uPnP Filtering multicast group address range is 239.255.255.250/32. If you do not configure the multicast group range, uPnP Filtering filters multicast packets destined for the default multicast group range.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable uPnP Filtering:

```
ip igmp upnp-filter [ip {A.B.C.D/X}]
```

Example

Enable uPnP Filtering on a VLAN using the default multicast group address range:

```
Switch:1(config-if)#ip igmp upnp-filter
```

Enable uPnP Filtering on a VLAN with configured multicast group address range:

```
Switch:1(config-if)#ip igmp upnp-filter ip 233.252.0.0/24
```

Variable Definitions

The following table defines parameters for the **ip igmp upnp-filter** command.

Variable	Value
<code>gigabitethernet {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>ip {A.B.C.D/X}</code>	Configures the multicast destination IP address range to filter for an IGMP interface. The default multicast group address is 239.255.255.250/32.
<code>vlan <1-4059></code>	Specifies the VLAN.

View uPnP Filtering information on an IGMP-enabled interface

Use the following command to display uPnP Filtering information on an IGMP-enabled interface.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about the interfaces where IGMP is enabled:
show ip igmp interface [gigabitethernet {slot/port[-slot/port] [, ...] [slot/all] [all]}] [vlan <1-4059> [vrf WORD<1-16>] [vrfsids WORD<0-512>]]

Example

View the default uPnP Filtering information on an IGMP-enabled interface:

```
Switch:1(config)#show ip igmp interface vlan 20
=====
Vlan Ip Igmp
=====
VLAN QUERY QUERY ROBUST VERSION VERSION LAST PROXY SNOOP SNOOP SSM UPnP FAST FAST FAST
ID INTVL MAX RESP ORIGIN MEMB SNOOP ENABLE ORIGIN SNOOP FILTER LEAVE LEAVE LEAVE
                QUERY ENABLE ENABLE ENABLE ORIGIN ENABLE ENABLE ENABLE ORIGIN PORTS
-----
20 125 100 2 2 CONFIG 10 true true RADIUS false true false RADIUS

VLAN SNOOP SNOOP DYNAMIC COMPATIBILITY EXPLICIT HOSTS UPnP
ID QUERIER QUERIER DOWNGRADE MODE TRACKING FILTER
  ENABLE ORIGIN ADDRESS ORIGIN VERSION ENABLE ORIGIN ENABLE ORIGIN ADDRESS
-----
20 true RADIUS 192.0.2.1 RADIUS enable disable RADIUS disable RADIUS 198.51.100.0/24
```

View uPnP Filtering information on an IGMP-enabled interface when uPnP Filtering is enabled and a non-default multicast group address is configured:

```
Switch:1(config)#show ip igmp interface vlan 2
=====
Vlan Ip Igmp
=====
VLAN QUERY QUERY ROBUST VERSION LAST PROXY SNOOP SNOOP SSM UPnP FAST FAST
ID INTVL MAX MEMB SNOOP ENABLE ORIGIN SNOOP FILTER LEAVE LEAVE
                QUERY ENABLE ENABLE ENABLE ORIGIN ENABLE ENABLE ENABLE PORTS
                RESP
-----
20 true RADIUS 192.0.2.1 RADIUS enable disable RADIUS disable RADIUS 198.51.100.0/24
```


2	125	100	2	2	10	false	false	RADIUS	false	true	false
VLAN ID	SNOOP QUERIER ENABLE	SNOOP QUERIER ADDRESS		DYNAMIC DOWNGRADE VERSION	COMPATIBILITY MODE		EXPLICIT HOSTS TRACKING	UPnP FILTER ADDRESS			
2	false	0.0.0.0		enable	disable		disable	233.252.0.0/24			

Variable Definitions

The following table defines parameters for the **show ip igmp interface** command.

Variable	Value
<i>gigabitethernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>vlan <1-4059></i>	Specifies the VLAN.
<i>vrf WORD<1-16></i>	Specifies the VRF by name.
<i>vrfids WORD<0-512></i>	Specifies the VRF by VRF ID.

Determine the Data Stream Learned with IP Multicast over Fabric Connect on the VLAN

Use this procedure to determine the data stream learned when IP Multicast over Fabric Connect is configured on the VLAN.

About This Task

The following section shows sample output for the **show ip mroute route** command.

In this table, every stream uses one (*,G) entry and x (S,G) entries, depending on how many servers forward traffic to the same group.

The 0.0.0.0 mask is always tied to a (*,G) entry.

If you do not specify a VRF name or range of VRF IDs, the results display information for the Global Router. If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Determine the data stream learned:

```
show ip mroute route [vrf WORD <0-32>] [vrfids <0-255>]
```

Example

Determine the data stream learned:

```
Switch:1>show ip mroute route
```

```
=====
Mroute Route - GlobalRouter
=====
GROUP          SOURCE          SRCMASK          UPSTREAM_NBR    IF      EXPIR  PROT
-----
233.252.0.1    0.0.0.0          0.0.0.0          0.0.0.0         V3      30    spb-access
233.252.0.1    192.0.2.102     255.255.255.0    0.0.0.0         -       0     spb-network
233.252.0.2    0.0.0.0          0.0.0.0          0.0.0.0         V2      30    pims
225.1.1.1      198.51.100.99   255.255.255.0    0.0.0.0         V3      173   spb-pim-gw

Total 4
```

Variable Definitions

The following table defines parameters for the **show ip mroute route** command.

Variable	Value
<i>vrf WORD<0-32></i>	Specifies the VRF name.
<i>vrfids <0-255></i>	Specifies the VRF ID.

IP Multicast over Fabric Connect configuration using EDM

This section provides procedures to configure IP Multicast over Fabric Connect using Enterprise Device Manager (EDM).

Configure IP Multicast over Fabric Connect Globally

Before You Begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- You must create the customer VLANs (C-VLANs) and add slots/ports.
- You must add IST slot/ports to the C-VLAN for an SMLT topology.

About This Task

Use this procedure to enable IP Multicast over Fabric Connect globally on the Backbone Edge Bridges (BEBs) that directly, or indirectly (using Layer 2 switches), connect to IP multicast senders or receivers. By default, IP Multicast over Fabric Connect is disabled. You do not need to enable IP Multicast over Fabric Connect on the Backbone Core Bridges (BCBs).

You must configure IP Multicast over Fabric Connect globally, and then enable it on the service option or options you choose.

You must enable IP multicast over Fabric Connect globally on all DvR enabled nodes (Controllers and Leaf nodes) in a DvR domain.



Note

IP Multicast over Fabric Connect uses I-SIDs starting at 16,000,000 and above. If Layer 2 or Layer 3 I-SIDs are in this range, the system displays an error message and the switch does not enable IP Multicast over Fabric Connect.

Procedure

1. In the navigation pane, expand **Configuration > Fabric > SPBM**.
2. Select the **I-SID** tab to determine if the I-SIDs are within the default range reserved for multicast.
3. Select the **SPBM** tab.
4. To enable multicast on an SPBM instance that already exists, in the **Mcast** column in the table, select **enable**.
5. To enable multicast on a new SPBM instance, select **Insert**.
6. In **Mcast**, select **enable** to enable IP Multicast over Fabric Connect globally.
7. Select **Insert**.
8. Select **Apply**.

SPBM Field Descriptions

Use the data in the following table to use the **SPBM** tab.

Name	Description
Id	Specifies the SPBM instance ID. Only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Specifies the primary SPBM B-VLANs to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.
LsdbTrap	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.
IpShortcut	Enables or disables SPBM IP shortcut state. The default is disable.
SmltSplitBEB	Specifies whether the switch is the primary or secondary vIST peer. The default is primary.
SmltVirtualBmac	Specifies a virtual MAC address that can be used by both peers.
SmltPeerSysId	Specifies the system ID of the SPBM SMLT for this SPBM instance.

Name	Description
Mcast	Specifies if IP multicast over SPBM is enabled. The default is disabled.
McastFwdCacheTimeout	Specifies the global forward cache timeout in seconds. The default is 210 seconds.
Ipv6Shortcut	Enables or disables SPBM IPv6 shortcut state. The default is disable.
McastSpbPimGwControllerEnable	Enables or disables IS-IS PLSB Multicast SPB PIM Gateway controller. Disabled by default.
McastSpbPimGwGatewayEnable	Enables or disables IS-IS PLSB Multicast SPB PIM Gateway. Disabled by default.
StpMultiHoming	Enables or disables MSTP-Fabric Connect Multi Homing. The default is disabled (false).
BVlanOrigin	Shows how the B-VLAN was created. The values can be config for manual configuration using CLI or SNMP, or dynamic through Zero Touch Fabric Configuration and Auto-sense. The default is dynamic.
RoutedMulticastTtl1Bridging	Enables or disables routed SPB IP Multicast TTL 1 bridging. The default is disabled.

View IP Multicast over Fabric Connect Routes

Use this procedure to display the IP Multicast over Fabric Connect routes.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **IpMcastRoutes** tab.

IpMcastRoutes Field Descriptions

Use the data in the following table to use the **IpMcastRoutes** tab.

Name	Description
VsnIsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Source	Specifies the IP address where the IP Multicast over Fabric Connect route originated.

Name	Description
NickName	Specifies the nickname used to filter criteria.
SourceBeb	Specifies the source BEB for the IP multicast route.
VlanID	Specifies the ID for the C-VLAN.
VrfName	Specifies the VRF name.
Datalsid	Specifies the data I-SID for the IP Multicast over Fabric Connect route. A a BEB receives IP multicast data from a sender, a BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
Type	Specifies the type for the IP Multicast over Fabric Connect route.
Bvlan	Specifies the B-VLAN for the IP Multicast over Fabric Connect route.
NniInterfaces	Specifies the network-to-network interface (NNI) ports for the IP multicast route. SPBM runs in the core on the ports that connect to the core. These ports are NNI ports. Ports that face a customer VLAN are user-to-network interface (UNI) ports.

Configure IP Multicast config-lite for Fabric Connect

Before You Begin

Create a VLAN. For more information see, [VLAN Configuration using EDM](#) on page 4043

About This Task

Use this procedure to dynamically create a VRF on one of the reserved Layer 3 VSN VLANs. The dynamically created VRF, **13vsn-*<I-SID number>***, obtains the I-SID number from **mvpnlsidValue**. You can enable SPB Multicast on the VLAN without configuring an IP address.



Note

If you are using a DvR leaf instance, the VRF name is **13vsn-*<13isid>***.

For additional information about the tabs referenced in this procedure, see the following sections:

- **VLAN > VLANs > Advanced** tab: [Advanced Field Descriptions](#) on page 4056
- **IP > IGMP > Interface** tab: [Interface Field Descriptions](#) on page 1991

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. In the appropriate VLAN row, for **mvpnlsidStatus**, select **Static**.
5. In **mvpnlsidValue**, type the I-SID number.
6. (Optional) In **mvpnlsidOffset**, type the MAC offset value.
7. Select **Apply**.
8. (Optional) Configure a routed SPB querier source IP address:
 - a. Expand **Configuration > IP**.
 - b. Select **IGMP**.
 - c. Select the **Interface** tab.
 - d. In **ExtnRoutedSpbQuerierAddr** for the appropriate interface, type the IP address.
 - e. Select **Apply**.
9. Return to the **Configuration > VLAN > VLANs > Advanced** tab.
10. In the appropriate row, for **SpbMulticast**, select **enable**.
11. Select **Apply**.

Display the UNI Ports for IP Multicast Routes

Use this procedure to display UNI ports associated with particular IP multicast routes.

Procedure

1. In the navigation pane, expand **Configuration > Fabric > SPBM**.
2. Select the **IpMcastRoutes** tab.
3. Select the desired row, and then select the **UNI Ports** tab to display the UNI ports associated with a particular stream.

IpMcastRoutes Uni Ports Field Descriptions

Use the data in the following table to use the **IpMcastRoutes Uni Ports** tab.

Name	Description
Group	Specifies the group IP address for the IP Multicast over Fabric Connect route.
Source	Specifies the IP address where the IP Multicast over Fabric Connect route originated.
Vsnlsid	Specifies the VSN I-SID. Layer 2 VSN and Layer 3 VSN each require a VSN I-SID.

Name	Description
Datalsid	Specifies the data I-SID for the IP multicast route. After a BEB receives the IP multicast data from a sender, a BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the source IP address, group IP address, and the local VLAN the stream is received on. The data I-SID is a child of the scope or VSN I-SID.
SourceBeb	Specifies the source BEB for the IP multicast route.
VlanId	Specifies the ID for the C-VLAN.
VrfName	Specifies the VRF name.
NniPorts	Specifies the network-to-network interface (NNI) ports for the IP multicast route. SPBM runs in the core on the ports that connect to the core. These ports are NNI ports. Ports facing a customer VLAN are user-to-network interface (UNI) ports.
Type	Specifies the type for the IP multicast route.
Bvlan	Specifies the B-VLANs for the IP multicast route.

Display the Multicast FIB

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **Multicast FIB** tab.

Multicast FIB Field Descriptions

Use the data in the following table to use the **Multicast FIB** tab.

Name	Description
SysId	System ID of the node where the multicast FIB entry originated.
Vlan	VLAN of the multicast FIB entry.
McastDestMacAddr	Multicast destination MAC Address of the multicast FIB entry
Isid	I-SID of the multicast FIB entry.
Isid Name	Name assigned to the I-SID.
HostName	Host name of the node where the multicast FIB entry originated.

Name	Description
OutgoingInterfaces	Specifies the switched UNI port outgoing interface of multicast FIB entry.
IncomingInterface	Specifies the incoming interface (port or MLT) of the multicast FIB entry.

IP Multicast over Fabric Connect configuration examples

This section provides configuration examples to configure IP Multicast over Fabric Connect.

IP multicast over Fabric Connect global configuration

The following sections show the steps required to configure IP multicast over Fabric Connect at a global level

SwitchC

```
enable
configure terminal
prompt SwitchC

ISIS SPBM CONFIGURATION
router isis
spbm 1 multicast enable
exit
```

SwitchG

```
enable
configure terminal
prompt SwitchG

ISIS SPBM CONFIGURATION
router isis
spbm 1 multicast enable
exit
```

SwitchD

```
enable
configure terminal
prompt SwitchD

ISIS SPBM CONFIGURATION
router isis
spbm 1 multicast enable
exit
```

Layer 2 VSN Configuration for IP Multicast over Fabric Connect

The topics in this section provide conceptual and procedural information about Layer 2 VSN for IP Multicast over Fabric Connect. For information on Layer 2 VSN basic configuration, see [Layer 2 VSN Configuration](#) on page 1514.

Layer 2 VSN IP Multicast over Fabric Connect

IP Multicast over Fabric Connect supports Layer 2 VSN functionality where multicast traffic is bridged over the SPBM core infrastructure. An application for Layer 2 VSNs using IP Multicast over Fabric Connect is multicast traffic in data centers.

For more information on Layer 2 VSN configuration, see [Layer 2 VSN Configuration](#) on page 1514.

After you configure **ip igmp snooping** on a VLAN that has an I-SID configured (a C-VLAN), that VLAN is automatically enabled for IP Multicast over Fabric Connect services. No explicit configuration exists separate from that to enable Layer 2 VSN IP Multicast over Fabric Connect.

Multicast traffic remains in the same Layer 2 VSN across the SPBM cloud for Layer 2 VSN IP Multicast over Fabric Connect. IP Multicast over Fabric Connect constrains all multicast streams within the scope level in which they originate. If a sender transmits a multicast stream to a BEB on a Layer 2 VSN with IP Multicast over Fabric Connect enabled, only receivers that are part of the same Layer 2 VSN can receive that stream.

I-SIDs

After a BEB receives IP multicast data from a sender, the BEB allocates a data service instance identifier (I-SID) in the range of 16,000,000 to 16,512,000 for the multicast stream. The stream is identified by the S, G, V tuple, which is the source IP address, the group IP address and the local VLAN the multicast stream is received on. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID.

In the context of Layer 2 VSNs with IP Multicast over Fabric Connect, the scope is the I-SID value of the Layer 2 VSN associated with the local VLAN on which the IP multicast data was received.

IP Multicast Forwarding Behavior within a Layer 2 VSN

Layer 2 VSN multicast forwarding allows the Fabric to bridge multicast traffic when the sender and receiver belong to the same Layer 2 VSN. When an IP Multicast sender is processed on a VLAN that has a Layer 2 VSN I-SID assigned, the ingress BEB advertises TLV 188 for each SPB IP Multicast stream. TLV 188 maps the sender Layer 2 VSN I-SID to the corresponding SPB IP Multicast data I-SID. If the sender processes on a VLAN without an Layer 2 VSN assignment, the ingress BEB does not advertise TLV 188.

Receiving BEBs use TLV 188 to determine the following forwarding behavior on a per-stream basis:

- If the receiver belongs to the same Layer 2 VSN I-SID as the sender, the switch bridges the multicast traffic. The IP TTL value is not decremented and the source C-MAC address is preserved.
- If the multicast stream does not map to the receiver Layer 2 VSN I-SID through TLV 188, or if the receiver processes on a VLAN without an Layer 2 VSN assignment, the switch routes the traffic. The IP TTL value is decremented and the source C-MAC address is modified.

TLVs

See [TLVs](#) on page 1957.

IGMP

After a BEB receives an IGMP join message from a receiver, a BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the request stream exists, the BEB sends an IS-IS TLV update to its neighbors to inform them of the presence of a receiver and this information is propagated through the SPBM cloud.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

NEW! TTL Handling for Bridged Traffic with Routed SPB IP Multicast

This feature enhances SPB IP Multicast forwarding by preserving the IP TTL when a sender and receiver operate within the same Layer 2 VSN. When the ingress BEB processes a multicast sender in a Layer 2 VSN, it advertises TLV 188 to map the Layer 2 VSN I-SID to the SPBM IP Multicast data I-SID. This mapping enables the switch to bridge the multicast stream and maintain the original TTL. TTL sensitive applications such as PTPv1 can therefore run over a Fabric that uses IGMP Snooping or SPBM Multicast-lite.

Configuration Considerations

Consider the following when you configure TTL handling on Multi-area SPB deployments:

- You cannot configure TTL handling for bridged traffic with routed SPB IP Multicast on a Multi-area boundary node.
- TTL handling and **router isis remote** are mutually exclusive:
 - You cannot enable TTL handling if **router isis remote** is administratively enabled.
 - You cannot enable **router isis remote** if TTL handling is already configured.

Despite this configuration dependency, inter-area availability remains unaffected, and edge nodes in different SPB areas can still bridge traffic without decrementing the TTL.

The following list describes how TLV 188 controls bridged and routed forwarding for SPB IP Multicast traffic:

- The ingress BEB advertises TLV 188 when it processes a sender on a VLAN with an assigned Layer 2 VSN I-SID. TLV 188 maps the sender Layer 2 VSN I-SID to the SPB IP Multicast data I-SID.
- The ingress BEB suppresses TLV 188 when it processes a sender on a VLAN without a Layer 2 VSN assignment.

- A receiver in a VLAN that belongs to the same Layer 2 VSN I-SID advertised in TLV 188 bridges the multicast stream. The switch does not decrement the IP TTL.
- A receiver handling a multicast stream that does not map to the same Layer 2 VSN I-SID through TLV 188 routes the traffic and decrements the TTL.
- A receiver on a VLAN without a Layer 2 VSN assignment also routes the multicast streams and decrements the TTL.
- Multi-area boundary nodes redistribute TLV 188 so the forwarding behavior remains consistent across different SPB areas.

Layer 2 VSN Configuration using the CLI

This section provides procedures to configure Layer 2 VSN using the command line interface (CLI).

Configure Layer 2 VSN IP Multicast over Fabric Connect

Use this procedure to configure IP Multicast over Fabric Connect for Layer 2 VSN functionality. With Layer 2 VSN IP Multicast over Fabric Connect, multicast traffic remains in the same Layer 2 VSN across the SPBM cloud.

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the customer VLANs (C-VLANs) and add slots/ports.
- Assign the same I-SID to the C-VLANs on all the BEBs where you configure the C-VLAN.
- Enable IP Multicast over Fabric Connect globally.

About This Task

Traffic is only delivered to UNIs on the Layer 2 VSN where the switch receives IGMP joins and reports. Traffic does not cross the Layer 2 VSN boundary.

Configuring **ip igmp snooping** on a VLAN that has an I-SID configured (a C-VLAN) automatically enables that VLAN for IP Multicast over Fabric Connect services. No explicit configuration exists separate from that to enable Layer 2 VSN IP Multicast over Fabric Connect.

SPBM supports enabling IGMP Snooping on a C-VLAN, but it does not support enabling Protocol Independent Multicast (PIM) on a C-VLAN. If you enable IGMP snooping on a C-VLAN, then its operating mode is Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network.

The switch only supports IPv4 multicast traffic.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```

2. Enable proxy snoop:

```
ip igmp proxy
```

3. Enable IGMP snooping:

```
ip igmp snooping
```

4. (Optional) If you want to configure an address for the IGMP queries, enter the following command:

```
ip igmp snoop-querier-addr <A.B.C.D>
```

This step is not always required. The IGMP Querier on the BEB uses a source address 0.0.0.0 by default. When you do not configure this, a BEB sends IGMP queries on the UNI ports with 0.0.0.0 as the source IP address. Some Layer 2 edge switches do not support a 0.0.0.0 querier. You can use a fictitious IP address as the querier address, and use the same address on all BEBs in the network.

5. (Optional) Enable IGMPv3 at a VLAN level by enabling SSM-snooping and IGMPv3:

```
ip igmp ssm-snoop

ip igmp version 3
```

You must enable SSM snoop before you configure IGMP version 3 and both ssm-snoop and snooping must be enabled for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

Examples

Enable IGMPv2 at a VLAN level:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config-if)#interface vlan 501
Switch:1(config-if)#ip igmp proxy
Switch:1(config-if)#ip igmp snooping
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
```

Enable IGMPv3 at a VLAN level:

```
Switch:>enable
Switch:#configure terminal
Switch:1(config)#interface vlan 2256
Switch:1(config-if)#ip igmp proxy
Switch:1(config-if)#ip igmp snooping
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
Switch:1(config-if)#ip igmp version 3
Switch:1(config-if)#ip igmp ssm-snoop
```

NEW! Enable Routed SPB IP Multicast TTL1 Bridging

Before You Begin

- Disable IP Multicast.

About This Task

Perform the following procedure to enable routed SPB IP Multicast TTL1 bridging. The default is disabled.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

2. Enable TTL1 bridged traffic with routed SPB IP Multicast:

```
spbm <1-100> routed-multicast-ttl1-bridging enable
```

3. Display the configuration status:

```
show isis spbm multicast
```

Example

```
Switch:1#show isis spbm multicast

multicast : enable
fwd-cache-timeout(seconds) : 210
routed-multicast-ttl1-bridging : disable
```

View Layer 2 VSN IP Multicast over Fabric Connect information

Use the following options to display Layer 2 VSN information to confirm proper configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all IP Multicast over Fabric Connect route information:

```
show isis spbm ip-multicast-route [all]
```
3. Display detailed IP Multicast over Fabric Connect route information:

```
show isis spbm ip-multicast-route [detail]
```
4. Display IP multicast route information by VLAN:

```
show isis spbm ip-multicast-route [vlan <1-4059>]
```
5. Display IP Multicast over Fabric Connect route information by VSN I-SID:

```
show isis spbm ip-multicast-route [vsid <1-16777215>]
```
6. Display IP Multicast over Fabric Connect route information by group address:

```
show isis spbm ip-multicast-route [group {A.B.C.D}]
```

7. Display IP Multicast over Fabric Connect route information by source address:

```
show isis spbm ip-multicast-route [source {A.B.C.D}]
```

**Important**

When you use the command `show isis spbm ip-multicast-route` without parameters, or use the detail or group optional parameters without specifying a VLAN ID or VSN I-SID, the command output displays Layer 3 context only. No Layer 2 context is displayed.

8. Display summary information for each S, G, V tuple with the corresponding scope, data I-SID, and the host name of the source:

```
show isis spb-mcast-summary [count] [host-name WORD<0-255>] [lspid <xxxx.xxxx.xxxx.xx-xx>]
```

Examples

```
Switch:1#show isis spbm ip-multicast-route all
```

SPBM IP-MULTICAST ROUTE INFO ALL										
Type	VrfName	Vlan Id	Source	Group	VSN-ISID	Data ISID	BVLAN	Source-BEB	AREA	AREA-NAME
snoop	GRT	501	192.0.2.1	233.252.0.1	5010	16300001	10	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.2	5010	16300002	20	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.3	5010	16300003	10	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.4	5010	16300004	20	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.5	5010	16300005	10	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.6	5010	16300006	20	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.7	5010	16300007	10	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.8	5010	16300008	20	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.9	5010	16300009	10	e12	HOME	area-10.0010
snoop	GRT	501	192.0.2.1	233.252.0.10	5010	16300010	20	e12	HOME	area-10.0010

```
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
```

```
Switch:1>show isis spbm multicast
```

```
multicast : enable
fwd-cache-timeout(seconds) : 210
routed-multicast-ttl1-bridging : disable
```

```
Switch:1#show isis spbm ip-multicast-route vlan 100
```

SPBM IP-MULTICAST ROUTE INFO - VLAN ID : 100, VSN-ISID : 100						
Source	Group	Data ISID	BVLAN	Source-BEB	AREA	AREA-NAME
100.100.0.105	225.1.2.1	16003014	4052	vn-area-20.0020	HOME	area-10.0010
100.100.0.105	225.1.2.1	16153041	4051	vn-area-30.0030	REMOTE	area-20.0020
100.100.0.105	225.1.2.2	16003011	4051	vn-area-20.0020	HOME	area-10.0010
100.100.0.105	225.1.2.2	16153049	4051	vn-area-30.0030	REMOTE	area-20.0020
100.100.0.105	225.1.2.3	16003008	4052	vn-area-20.0020	HOME	area-10.0010
100.100.0.105	225.1.2.3	16153057	4051	vn-area-30.0030	REMOTE	area-20.0020
100.100.0.105	225.1.2.4	16003005	4051	vn-area-20.0020	HOME	area-10.0010
100.100.0.105	225.1.2.4	16153065	4051	vn-area-30.0030	REMOTE	area-20.0020

```

100.100.0.105 225.1.2.5 16003002 4052 vn-area-20.0020 HOME area-10.0010
100.100.0.105 225.1.2.5 16153073 4051 vn-area-30.0030 REMOTE area-20.0020
100.100.0.105 225.1.3.1 16002999 4051 vn-area-20.0020 HOME area-10.0010
100.100.0.105 225.1.3.1 16153081 4051 vn-area-30.0030 REMOTE area-20.0020
100.100.0.105 225.1.3.2 16002997 4051 vn-area-20.0020 HOME area-10.0010
100.100.0.105 225.1.3.2 16153087 4051 vn-area-30.0030 REMOTE area-20.0020

```

```
--More-- (q = quit)
```

```
Switch:1# show isis spbm ip-multicast-route vsn-isisd 5010
```

```
=====
SPBM IP-MULTICAST ROUTE INFO - VLAN ID : 501, VSN-ISID : 5010
=====
```

Source	Group	Data ISID	BVLAN	Source-BEB	AREA	AREA-NAME
192.0.2.1	233.252.0.1	16300001	10	e12	HOME	area-10.0010
192.0.2.1	233.252.0.2	16300002	20	e12	HOME	area-10.0010
192.0.2.1	233.252.0.3	16300003	10	e12	HOME	area-10.0010
192.0.2.1	233.252.0.4	16300004	20	e12	HOME	area-10.0010
192.0.2.1	233.252.0.5	16300005	10	e12	HOME	area-10.0010
192.0.2.1	233.252.0.6	16300006	20	e12	HOME	area-10.0010
192.0.2.1	233.252.0.7	16300007	10	e12	HOME	area-10.0010
192.0.2.1	233.252.0.8	16300008	20	e12	HOME	area-10.0010
192.0.2.1	233.252.0.9	16300009	10	e12	HOME	area-10.0010
192.0.2.1	233.252.0.10	16300010	20	e12	HOME	area-10.0010

```
-----
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
-----
```

```
Switch:1# show isis spb-mcast-summary
```

```
=====
SPB Multicast - Summary
=====
```

SCOPE	SOURCE	GROUP	DATA	LSP	HOST	AREA	AREA NAME
I-SID	ADDRESS	ADDRESS	I-SID	BVID	FRAG NAME		
5010	192.0.2.1	233.252.0.1	16300001	10	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.3	16300003	10	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.5	16300005	10	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.7	16300007	10	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.9	16300009	10	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.2	16300002	20	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.4	16300004	20	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.6	16300006	20	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.8	16300008	20	0x0 e12	HOME	area-10.0010
5010	192.0.2.1	233.252.0.10	16300010	20	0x0 e12	HOME	area-10.0010

Variable definitions

The following table defines parameters for the **show isis spbm ip-multicast-route** command.

Variable	Value
<i>all</i>	Displays all IP Multicast over Fabric Connect route information.
<i>detail</i>	Displays detailed IP Multicast over Fabric Connect route information.

Variable	Value
<i>group {A.B.C.D}</i> <i>source {A.B.C.D}</i>	Displays information on the group IP address for the IP Multicast over Fabric Connect route. If you select source it will also display the source IP address.
<i>vlan <0-4084></i>	Displays IP Multicast over Fabric Connect route information by VLAN.
<i>vrf WORD<1-16></i>	Displays IP Multicast over Fabric Connect route information by VRF.
<i>vsn-isid <1-16777215></i>	Displays IP Multicast over Fabric Connect route information by I-SID.

The following table defines parameters for the **show isis spb-mcast-summary** command.

Variable	Value
<i>count</i>	Displays the total number of SPB multicast entries.
<i>host-name WORD<0-255></i>	Displays the IP Multicast over Fabric Connect summary for a given host-name.
<i>lspid <xxxx.xxxx.xxxx.xx-xx></i>	Displays the IP Multicast over Fabric Connect summary for a given LSP ID.

View IGMP Information for Layer 2 VSN Multicast

Use the following commands to display IGMP information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display information about the interfaces where IGMP is enabled:

```
show ip igmp interface [gigabitethernet {slot/port[-slot/port]}[,...]  
[slot/all][all]][vlan <1-4059>[vrf WORD<1-16>][vrfids WORD<0-512>]
```

Ensure that the output displays snoop-spb under MODE.

3. Display information about the IGMP cache:

```
show ip igmp cache [vrf WORD<1-16>][vrfids WORD<0-512>]
```

4. Display information about the IGMP group:

```
show ip igmp group [count][group {A.B.C.D}][member-subnet {A.B.C.D/X}]  
[vrf WORD<1-16>][vrfids WORD<0-512>]
```

5. Display information about the IGMP sender:

```
show ip igmp sender [count][group {A.B.C.D}][member-subnet {A.B.C.D/  
X}][vrf WORD<1-16>][vrfids WORD<0-512>]
```

6. Display information about IGMP snoop-trace information:

```
show ip igmp snoop-trace [group {A.B.C.D}][source {A.B.C.D}][vrf  
WORD<1-16>][vrfids WORD<0-512>]
```


Example

```
Switch:1#show ip igmp interface vrf l3vsn-100
```

Igmp Interface - VRF l3vsn-100												
IF	QUERY INTVL	STATUS	VERS.	OPER VERS	QUERIER	QUERY MAXRSPT	WRONG QUERY	JOINS	ROBUST	LASTMEM QUERY	MODE	SCOPE L2ISID
V10	125	inact	2	2	0.0.0.0	100	0	0	2	10	routed-spb	0 100

```
Switch:1(config)#show ip igmp interface vlan 20
```

Vlan Ip Igmp														
VLAN ID	QUERY INTVL	QUERY MAX RESP	ROBUST	VERSION	VERSION ORIGIN	LAST MEMB QUERY	PROXY SNOOP ENABLE	SNOOP ENABLE	SNOOP ORIGIN	SSM SNOOP ENABLE	UPnP FILTER ENABLE	FAST LEAVE ENABLE	FAST LEAVE ORIGIN	FAST LEAVE PORTS
20	125	100	2	2	CONFIG	10	true	true	RADIUS	false	true	false	RADIUS	
VLAN ID	SNOOP QUERIER	SNOOP ORIGIN	SNOOP QUERIER ADDRESS	DYNAMIC ORIGIN	COMPATIBILITY DOWNGRADE VERSION	COMPATIBILITY MODE	EXPLICIT ORIGIN	HOSTS TRACKING	UPnP FILTER ADDRESS					
20	true	RADIUS	192.0.2.1	RADIUS	enable	disable	RADIUS	disable	RADIUS 198.51.100.0/24					

```
Switch:1#show ip igmp group
```

Igmp Group - GlobalRouter					
GRPADDR	INPORT	MEMBER	EXPIRATION	TYPE	L2ISID
224.5.2.1	V701-1/4	62.0.1.1	214	Dynamic	40400
224.5.2.2	V702-1/4	62.0.2.1	221	Dynamic	40400
224.5.2.3	V703-1/4	62.0.3.1	217	Dynamic	40400
224.5.2.4	V704-1/4	62.0.4.1	223	Dynamic	40400

4 out of 4 group Receivers displayed

Total number of unique groups 2

```
Switch:1# show ip igmp sender
```

IGMP Sender - GlobalRouter							
PORT/							
GRPADDR	IFINDEX	MEMBER	MLT	STATE	L2ISID	SCOPEISID	DATAISID
224.1.1.4	Vlan 10	192.2.0.1	1/18/1	NOTFILTERED	0	10010	16000001
224.1.1.4	Vlan 10	192.2.0.1	1/18/1	NOTFILTERED	0	10010	16000004
224.1.1.6	Vlan 10	192.2.0.1	1/18/1	FILTERED	0	10010	16000010
224.1.1.7	Vlan 10	192.2.0.1	1/18/1	FILTERED	0	10010	16000009
224.1.1.8	Vlan 10	192.2.0.1	1/18/1	NOTFILTERED	0	10010	210

5 out of 5 entries displayed

```
Switch:1# show ip igmp snoop-trace
```

Snoop Trace - GlobalRouter						
GROUP ADDRESS	SOURCE ADDRESS	IN VLAN	IN PORT	OUT VLAN	OUT PORT	TYPE
233.252.0.1	192.0.2.6	500	1/1	500	1/5	ACCESS
233.252.0.10	192.0.2.7	500	1/1	500	1/10	ACCESS

Variable Definitions

The following table defines parameters for the **show ip igmp interface** command.

Variable	Value
<i>gigabitethernet {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>vlan <1-4059></i>	Specifies the VLAN.
<i>vrf WORD<1-16></i>	Specifies the VRF by name.
<i>vrfids WORD<0-512></i>	Specifies the VRF by VRF ID.

The following table defines parameters for the **show ip igmp cache** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF by name.
<i>vrfids WORD<0-512></i>	Specifies the VRF by VRF ID.

The following table defines parameters for the **show ip igmp group** command.

Variable	Value
<i>count</i>	Specifies the number of entries.
<i>group {A.B.C.D}</i>	Specifies the group address.
<i>member-subnet {A.B.C.D/X}</i>	Specifies the IP address and network mask.
<i>vrf WORD<1-16></i>	Displays the multicast route configuration for a particular VRF by name.
<i>vrfids WORD<0-512></i>	Displays the multicast route configuration for a particular VRF by VRF ID.

The following table defines parameters for the **show ip igmp sender** command.

Variable	Value
<i>count</i>	Specifies the number of entries.
<i>group {A.B.C.D}</i>	Specifies the group address.
<i>member-subnet {A.B.C.D/X}</i>	Specifies the IP address and network mask.

Variable	Value
<code>vrf WORD<1-16></code>	Displays the multicast route configuration for a particular VRF by name.
<code>vrfids WORD<0-512></code>	Displays the multicast route configuration for a particular VRF by VRF ID.

The following table defines parameters for the **show ip igmp snoop-trace** command.

Variable	Value
<code>group {A.B.C.D}</code>	Specifies the group address.
<code>source {A.B.C.D}</code>	Specifies the source address.
<code>vrf WORD<1-16></code>	Displays the multicast route configuration for a particular VRF by name.
<code>vrfids WORD<0-512></code>	Displays the multicast route configuration for a particular VRF by VRF ID.

Troubleshoot IP Multicast over Fabric Connect for Layer 2 VSNs

If traffic is not moving properly, use the following procedure to determine the issue.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Ensure that all switch nodes in the network operate with the most recent software release to support IP Multicast over Fabric Connect:
`show software`
3. Ensure that you create and enable SPBM infrastructure globally.
 - a. Ensure that SPBM is enabled globally:
`show spbm`
 - b. Ensure that IS-IS is enabled globally:
`show isis`
 - c. Ensure an SPBM instance exists with at least one Backbone VLAN (B-VID). Also ensure multicast is enabled:
`show isis spbm`
4. Ensure that you enable the CFM configuration.
 - a. Ensure a CFM maintenance-association exists:
`show cfm maintenance-association`
 - b. Ensure a CFM maintenance-domain exists:
`show cfm maintenance-domain`
 - c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:
`show cfm maintenance-endpoint`

5. Ensure a Customer VLAN (C-VLAN) exists and ensure you add UNI ports to the C-VLAN.
 - a. Display C-VLAN information:
`show vlan i-sid`
 - b. Display ports for the C-VLAN:
`show vlan members port {slot/port[-slot/port]}[,...][slot/all][all]}`
 - c. Display NNI and UNI receivers:
`show isis spbm ip-multicast-route detail`
6. Ensure that you assign the same I-SID to the C-VLAN on all of the BEBs where you configure the C-VLAN:
`show vlan i-sid`
7. Ensure that you enable IP Multicast over Fabric Connect globally:
`show isis spbm`
8. Ensure that you enable IGMP Snooping on the C-VLAN on all of the Backbone Edge Bridges (BEBs). Ensure the protocol configured on the VLAN added is snoop-spb in the MODE column, which indicates IGMP is enabled on a VLAN with an associated I-SID (IP Multicast over Fabric Connect for a Layer 2 VSN):
`show ip igmp interface`
9. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:
`show ip igmp snoop-trace`

`show ip igmp interface`
10. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:
`show ip igmp interface`

Layer 2 VSN Configuration using EDM

This section provides procedures to configure Layer 2 VSN using Enterprise Device Manager (EDM).

View the IGMP Interface Table

Use the Interface tab to view the IGMP interface table. When an interface does not use an IP address, the system does not display it in the IGMP table.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IGMP**.
3. Select the **Interface** tab.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The default is 125.
Status	Shows the IGMP row status. If an interface uses an IP address and PIM-SM is enabled, the status is active. Otherwise, it is notInService.
Version	Configures the version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	Shows the version of IGMP that currently runs on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface attaches.
QueryMaxResponseTime	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds.)</p> <p>Important: You must configure this value lower than the QueryInterval.</p>
WrongVersionQueries	Shows the number of queries received with an IGMP version that does not match the interface. You must configure all routers on a LAN to run the same version of IGMP. If the interface receives queries with the wrong version, this value indicates a version mismatch.
Joins	Shows the number of times this interface added a group membership, which is the same as the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	<p>Tunes for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, increase the robustness value.</p> <p>The default value of 2 means that the switch drops one query for each query interval without the querier aging out.</p>

Name	Description
LastMembQueryIntvl	<p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of second. As a best practice, configure this parameter to values greater than 3. If you do not need a fast leave process, you can configure values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)</p>
OtherQuerierPresent Timeout	Shows the length of time that must pass before a multicast router determines that no other querier exists. If the local router is the querier, the value is 0.
FlushAction	<p>Configures the flush action to one of the following:</p> <ul style="list-style-type: none"> • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	<p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p>Important: To maximize network performance, configure this parameter according to the version of IGMP currently in use.</p> <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop.
SnoopQuerierEnable	Enables IGMP Layer 2 Querier.
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts for each channel or group. The default is disabled. You must select this parameter to use fast leave for IGMPv3.

Name	Description
McastMode	Specifies the protocol configured on the VLAN. <ul style="list-style-type: none"> snoop — Indicates IGMP snooping is enabled on a VLAN. snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over Fabric Connect for a Layer 2 VSN). pim — Indicates PIM is enabled. routed-spb — Indicates IP multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts.
ExtnUpnpFilterEnable	Enables Universal Plug and Play (uPnP) Filtering to filter multicast packets destined for a specific range. The default is disabled.
ExtnUpnpFilterAddress	Specifies the multicast destination IP address to filter on an IGMP-enabled interface. The default is 239.255.255.250/32.
ExtnUpnpFilterAddressMask	Specifies the IGMP uPnP Filtering IP subnet to which this interface is attached.
ExtnSnoopOrigin	Specifies the origin of IGMP Snooping configuration on the port. The supported values are: <ul style="list-style-type: none"> config - Configured by the user. radius - Configured by the Remote Authentication Dial-In User Service (RADIUS) attribute.
ExtnRoutedSpbQuerierAddr	Specifies the IP address of the SPB querier.

Configure IP Multicast over Fabric Connect on a Layer 2 VSN

Use this procedure to enable IP Multicast over Fabric Connect for a Layer 2 VSN. With Layer 2 VSN IP Multicast over Fabric Connect, multicast traffic remains in the same Layer 2 VSN across the SPBM cloud.

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the customer VLANs (C-VLANs) and add slots/ports.
- Assign the same I-SID to the C-VLANs on all the BEBs where you configure the C-VLAN.
- Enable IP Multicast over Fabric Connect globally.

About This Task

Traffic is only delivered to UNIs on the Layer 2 VSN where the switch receives IGMP joins and reports. Traffic does not cross the Layer 2 VSN boundary.

No explicit configuration exists for a Layer 2 VSN. After you configure IP IGMP snooping on a VLAN that has an I-SID configured, the device enables that VLAN for IP Multicast over Fabric Connect services.

SPBM supports enabling IGMP snooping on a C-VLAN, but it does not support enabling PIM on a C-VLAN. If you enable IGMP snooping on a C-VLAN, then its operating mode is Layer 2 VSN with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network.

This switch only supports IPv4 multicast traffic.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **IGMP** tab.
7. Select the **SnoopEnable** check box.
8. (Optional) Select the **SsmSnoopEnable** check box, if you use IGMP version 3.
9. (Optional) Select the **ProxySnoopEnable** check box.
10. (Optional) If you want to enable IGMP version 3, select version3 in the **Version** check box.

You must enable SSM snoop before you configure IGMP version 3 and both ssm-snoop and snooping must be enabled for IGMPv3.

11. If you want to enable IGMP version 2, select version2 in the **Version** check box.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

12. (Optional) If you want to enable snoop querier, select **SnoopQuerierEnable**.
13. (Optional) If you want to configure an address for IGMP queries, enter the IP address in **SnoopQuerierAddr**.



Note

This step is not always required. The IGMP Querier on the BEB uses a source address 0.0.0.0 by default. When you do not configure this, a BEB sends IGMP queries on the UNI ports with 0.0.0.0 as the source IP address. Some Layer 2 edge switches do not support a 0.0.0.0 querier. You can use a fictitious IP address as the querier address, and use the same address on all BEBs in the network.

14. Select **Apply**.

NEW! *Enable TTL Handling for Bridged Traffic with Routed SPB IP Multicast*

Before You Begin

- Disable IP Multicast.

About This Task

Use the following procedure to enable TTL Handling for Bridged Traffic with Routed SPB IP Multicast.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **SPBM** tab.
4. Select **Insert**.
5. For **RoutedMulticastTtlBridging**, select **enable**.
6. Select **Insert**.

SPBM Field Descriptions

Use the data in the following table to use the **SPBM** tab.

Name	Description
Id	Specifies the SPBM instance ID. Only one SPBM instance is supported.
NodeNickName	Specifies a nickname for the SPBM instance globally. Valid value is 2.5 bytes in the format <x.xx.xx>.
PrimaryVlan	Specifies the primary SPBM B-VLANs to add to the SPBM instance.
Vlans	Specifies the SPBM B-VLANs to add to the SPBM instance.
LsdbTrap	Configures whether to enable or disable a trap when the SPBM LSDB changes. The default is disable.
IpShortcut	Enables or disables SPBM IP shortcut state. The default is disable.
SmltSplitBEB	Specifies whether the switch is the primary or secondary VIST peer. The default is primary.
SmltVirtualBmac	Specifies a virtual MAC address that can be used by both peers.
SmltPeerSysId	Specifies the system ID of the SPBM SMLT for this SPBM instance.
Mcast	Specifies if IP multicast over SPBM is enabled. The default is disabled.
McastFwdCacheTimeout	Specifies the global forward cache timeout in seconds. The default is 210 seconds.
Ipv6Shortcut	Enables or disables SPBM IPv6 shortcut state. The default is disable.
McastSspbPimGwControllerEnable	Enables or disables IS-IS PLSB Multicast SPB PIM Gateway controller. Disabled by default.
McastSspbPimGwGatewayEnable	Enables or disables IS-IS PLSB Multicast SPB PIM Gateway. Disabled by default.

Name	Description
StpMultiHoming	Enables or disables MSTP-Fabric Connect Multi Homing. The default is disabled (false).
BVlanOrigin	Shows how the B-VLAN was created. The values can be config for manual configuration using CLI or SNMP, or dynamic through Zero Touch Fabric Configuration and Auto-sense. The default is dynamic.
RoutedMulticastTtl1Bridging	Enables or disables routed SPB IP Multicast TTL 1 bridging. The default is disabled.

Layer 2 VSN with IP Multicast over Fabric Connect Configuration Example

The following example shows the configuration steps to enable IP Multicast over Fabric Connect support on C-VLAN 1001 that is part of a Layer 2 VSN, including the querier address.

```
enable
configure terminal

ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VLAN CONFIGURATION

interface vlan 9
ip igmp snooping
ip igmp snoop-querier-addr 192.0.2.201
exit
```

When using IGMPv3, the configuration is:

```
enable
configure terminal

ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VLAN CONFIGURATION

interface vlan 19
ip igmp snooping
ip igmp version 3
ip igmp ssm-snoop
```

```
ip igmp snoop-querier-addr 192.0.2.201
exit
```

**Note**

You must enable SSM snoop before you configure IGMP version to version 3, and you must enable both **ssm-snoop** and **snooping** for IGMPv3.

**Note**

You must configure basic SPBM and IS-IS infrastructure.

IP Shortcuts Configuration for IP Multicast over Fabric Connect

This section provides concepts and procedures for IP Shortcuts configuration for IP Multicast over Fabric Connect. For information on IP Shortcuts basic configuration, see [IP Shortcuts Configuration](#) on page 1580.

IP Multicast over Fabric Connect within the GRT

Use IP Multicast over Fabric Connect within the GRT to exchange IP multicast traffic with all or a subset of VLANs that are in the Global Routing Table (GRT). This restriction is called the scope level, which IP Multicast over Fabric Connect uses to constrain the multicast streams within the level in which they originate. For example, if a sender transmits a multicast stream to a BEB on a VLAN that is part of the GRT with IP Multicast over Fabric Connect enabled, only receivers that are part of the same GRT can receive that stream.

Applications that can use IP Multicast over Fabric Connect within the GRT include:

- Video surveillance
- TV/Video/Ticker/Image distribution
- VX-LAN.

Both IP Shortcuts and IP Multicast over Fabric Connect within the GRT use the GRT for the scope level to constrain multicast streams. However, they are separate features that work independently from each other.

**Important**

You do not have to enable IP Shortcuts to support IP Multicast over Fabric Connect within the GRT.

With IP Multicast over Fabric Connect within the GRT, routing of IP multicast traffic is allowed within the subset of VLANs in the GRT that have IP Multicast over Fabric Connect enabled. When you enable IP Multicast over Fabric Connect on a VLAN, the VLAN automatically becomes a multicast routing interface.

You must enable **ip spb-multicast** on each of the VLANs within the GRT that need to support IP multicast traffic. Enable IP Multicast over Fabric Connect on all VLANs

to which IP multicast senders and receivers attach. IP Multicast over Fabric Connect is typically configured only on BEBs.

**Note**

If no IP interface exists on the VLAN, then you must create one. The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN.

If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect routing does not depend on unicast routing, which allows for you to more easily migrate from a PIM environment to IP Multicast over Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first, and then migrate unicast separately or not at all.

I-SIDs

Unlike IP Shortcuts with unicast, a data I-SID (for mac-in-mac encapsulation of the multicast traffic) is required for IP Multicast over Fabric Connect within the GRT. When the multicast stream reaches the BEB, the BEB assigns a data I-SID to the stream. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID.

Unlike Layer 2 VSNs and Layer 3 VSNs, IP Multicast over Fabric Connect within the GRT does not have a scope I-SID to determine the scope of the multicast traffic. Instead the scope is the Global Routing Table.

TLVs

See [TLVs](#) on page 1957.

IGMP

After you configure **ip spb-multicast enable**, you cannot enable IGMP, IGMP Snooping, or IGMP proxy on the interface. If you try to enable IGMP Snooping or proxy on an interface where IP Multicast over Fabric Connect is enabled, the system displays an error message.

After you configure **ip spb-multicast enable** on each of the VLANs within the GRT that need to support IP multicast traffic, IGMP functions required for IP Multicast over Fabric Connect within the GRT are automatically enabled. You do not need to configure anything IGMP related.

DvR

When you enable **ip spb-multicast** on the Controller nodes, the configuration is automatically pushed to all the Leaf nodes within the domain.

For more information on DvR, see [Distributed Virtual Routing](#) on page 915.

IP Shortcuts Configuration using the CLI

This section provides procedures to configure IP Shortcuts using the command line interface (CLI).

Configure IP Multicast over Fabric Connect on a VLAN Within the GRT

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.
- If no IP interface exists on the VLAN, create one. The IP interface must be the same subnet as the IGMP hosts that connect to the VLAN.

About This Task

Use this procedure to configure IP Multicast over Fabric Connect within the GRT. The default is disabled.

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

The switch only supports IPv4 addresses with IP Multicast over Fabric Connect.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Create an IP interface on the VLAN:

```
ip address <A.B.C.D/X>
```

3. Enable IP Multicast over Fabric Connect:

```
ip spb-multicast enable
```



Note

After you configure **ip spb-multicast enable**, you cannot enable IGMP, IGMP Snooping, or IGMP proxy on the interface. If you try to enable IGMP Snooping or proxy on any interface where IP Multicast over Fabric Connect is enabled, an error message displays.

When you configure **ip spb-multicast enable** on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

4. Ensure IP Multicast over Fabric Connect within the GRT is configured properly:

```
show ip igmp interface
```

If `routed-spb` displays under `MODE`, IP Multicast over Fabric Connect within the GRT is properly enabled on the VLAN.

Example

Enable IP Multicast over Fabric Connect within the GRT:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 500
Switch:1(config-if)#ip address 192.0.2.1 255.255.255.0
Switch:1(config-if)#ip spb-multicast enable
Switch:1#show ip igmp interface
```

```
=====
                        Igmp Interface - GlobalRouter
=====
```

IF	QUERY INTVL	STATUS	VERS.	OPER VERS	QUERIER	QUERY MAXRSPT	WRONG QUERY	JOINS	ROBUST	LASTMEM QUERY	MODE

ISID											

V500	125	active	2	2	0.0.0.0	100	0	0	2	10	routed-spb
V2000	125	inact	2	2	0.0.0.0	100	0	0	2	10	

```
1 out of 1 entries displayed
```

Variable Definitions

The following table defines parameters for the **interface vlan** command.

Variable	Value
<1-4059>	Specifies the VLAN ID.

The following table defines parameters for the **interface GigabitEthernet** command.

Variable	Value
<i>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

The following table defines parameters for the **ip address** command.

Variable	Value
<A.B.C.D/X>	Specifies the address and mask.

Configure the VRF Timeout To Age Out a Multicast Sender

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.

About This Task

Use this procedure to configure the VRF timeout value. The timeout value ages out the sender when there is no multicast stream on the VRF. The default is 210 seconds.



Note

You can use this procedure for Layer 3 VSN with IP Multicast over Fabric Connect services and IP Multicast over Fabric Connect for IP Shortcuts.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the timeout value on the VRF:

```
mvpn fwd-cache-timeout <10-432000>
```

Example

Configure the timeout value on the VRF to 500 seconds:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#mvpn fwd-cache-timeout 500
```

Variable definitions

The following table defines parameters for the **router vrf** command.

Variable	Value
<i>WORD</i> <1-16>	Specifies the VRF name.

The following table defines parameters for the **mvpn fwd-cache-timeout (seconds)** command.

Variable	Value
<10-432000>	Specifies the timeout value. The default is 210 seconds.

Configure the GRT Timeout to Age Out a Multicast Sender

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.

About This Task

Use this procedure to configure the timeout value in the GRT. The timeout value ages out the sender when there are no multicast streams coming from the sender for a specified period of time in seconds. The default timeout value is 210 seconds.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable

configure terminal

router isis
```

2. Configure the IP Multicast over Fabric Connect forward-cache timeout:

```
spbm <1-100> multicast fwd-cache-timeout <10-432000>
```

Example

Configure the IP Multicast over Fabric Connect forward-cache timeout to 300:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 multicast 1 fwd-cache-timeout 300
```

Variable definitions

The following table defines parameters for the **spbm** command.

Variable	Value
<1-100>	Specifies the SPBM instance. The switch only supports one instance.
<10-432000>	Specifies the IP Multicast over Fabric Connect forward-cache timeout in seconds. The default is 210 seconds.

View IP Multicast over Fabric Connect within the GRT information

About This Task

Use the following options to display IP Multicast over Fabric Connect within the GRT information to confirm proper configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all IP Multicast over Fabric Connect route information:

```
show isis spbm ip-multicast-route [all]
```
3. Display detailed IP Multicast over Fabric Connect route information:

```
show isis spbm ip-multicast-route [detail]
```
4. Display the IP Multicast over Fabric Connect multicast group and source address information:

```
show isis spbm ip-multicast-route [group {A.B.C.D}] [source {A.B.C.D}] [source-beb WORD<0-255>]
```
5. Display summary information for each S, G, V tuple with the corresponding scope, data I-SID, and the host name of the source:

```
show isis spb-mcast-summary [count] [host-name WORD<0-255>] [lspid <xxxx.xxxx.xxxx.xx-xx>]
```

Variable definitions

The following table defines parameters for the **show isis spbm ip-multicast-route** command.

Variable	Value
<i>all</i>	Displays all IP Multicast over Fabric Connect route information.
<i>detail</i>	Displays detailed IP Multicast over Fabric Connect route information.
<i>group {A.B.C.D}</i> <i>source {A.B.C.D}</i> [source-beb WORD<0-255>]	Displays information on the group IP address for the IP Multicast over Fabric Connect route. If you select source it will also display the source IP address. Specifies the source BEB name.
<i>vlan</i>	Displays IP Multicast over Fabric Connect route information by VLAN.

Variable	Value
<i>vrf</i>	Displays IP Multicast over Fabric Connect route information by VRF.
<i>vsn-isid</i>	Displays IP Multicast over Fabric Connect route information by I-SID.

The following table defines parameters for the **show isis spb-mcast-summary** command.

Variable	Value
<i>count</i>	Displays the total number of SPB multicast entries.
<i>host-name</i>	Displays the IP Multicast over Fabric Connect summary for a given host-name.
<i>lspid</i> <xxxx.xxxx.xxxx.xx-xx>	Displays the IP Multicast over Fabric Connect summary for a given LSP ID.

View IGMP Information for IP Multicast over Fabric Connect within the GRT

Use the following commands to display IGMP information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display information about the interfaces where IGMP is enabled:

```
show ip igmp interface [gigabitethernet {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]][vlan <1-4059>][vrf WORD<1-16>]
[vrfids WORD<0-512>]
```

Ensure that the output displays `routed-spb` under `MODE`.

3. Display information about the IGMP cache:

```
show ip igmp cache [vrf WORD<1-16>][vrfids WORD<0-512>]
```

4. Display information about the IGMP group:

```
show ip igmp group [count][group {A.B.C.D}][member-subnet default|
{A.B.C.D/X}][vrf WORD<1-16>][vrfids WORD<0-512>]
```

5. Display information about the IGMP sender:

```
show ip igmp sender [count][group {A.B.C.D}][member-subnet default|
{A.B.C.D/X}][vrf WORD<1-16>][vrfids WORD<0-512>]
```

Example

Display IGMP information for IP multicast over Fabric Connect within the GRT:

```
Switch:1#show ip igmp interface vrf l3vsn-100
=====
                        Igmp Interface - VRF l3vsn-100
=====
      QUERY      OPER      QUERY  WRONG      LASTMEM      SCOPE
IF  INTVL STATUS VERS.  VERS  QUERIER  MAXRSPT  QUERY  JOINS  ROBUST  QUERY  MODE    L2ISID  ISID
=====
```

```

V10    125    inact  2      2      0.0.0.0    100    0      0      2      10      routed-spb 0      100

Switch:1(config)#show ip igmp interface vlan 20
=====
                        Vlan Ip Igmp
=====
VLAN QUERY QUERY ROBUST VERSION VERSION LAST  PROXY  SNOOP  SNOOP          SSM  UPnP  FAST  FAST  FAST
ID  INTVL MAX  RESP          ORIGIN  MEMB  SNOOP  ENABLE ORIGIN          ENABLE  ENABLE  ENABLE  LEAVE  LEAVE  LEAVE
-----
20  125    100  2      2      CONFIG 10    true  true  RADIUS          false  true  false  RADIUS

VLAN SNOOP          SNOOP          DYNAMIC  COMPATIBILITY  EXPLICIT HOSTS  UPnP
ID  QUERIER          QUERIER          DOWNGRADE MODE  TRACKING        FILTER
   ENABLE ORIGIN  ADDRESS  ORIGIN  VERSION  ENABLE  ORIGIN  ENABLE  ORIGIN  ADDRESS
-----
20  true  RADIUS  192.0.2.1 RADIUS  enable  disable RADIUS  disable RADIUS  198.51.100.0/24

Switch:1# show ip igmp sender
=====
                        IGMP Sender - GlobalRouter
=====
                        PORT/
GRPADDR  IFINDEX  MEMBER  MLT    STATE      L2ISID SCOPEISID  DATAISID
-----
224.1.1.4  Vlan 10  192.2.0.1  1/18/1  NOTFILTERED  0  10010  16000001
224.1.1.4  Vlan 10  192.2.0.1  1/18/1  NOTFILTERED  0  10010  16000004
224.1.1.6  Vlan 10  192.2.0.1  1/18/1  FILTERED     0  10010  16000010
224.1.1.7  Vlan 10  192.2.0.1  1/18/1  FILTERED     0  10010  16000009
224.1.1.8  Vlan 10  192.2.0.1  1/18/1  NOTFILTERED  0  10010  210

5 out of 5 entries displayed
Switch:1#show ip igmp group
=====
                        Igmp Group - GlobalRouter
=====
GRPADDR      INPORT      MEMBER      EXPIRATION  TYPE      L2ISID
-----
224.5.2.1      V701-1/4      62.0.1.1      214      Dynamic  40400
224.5.2.2      V702-1/4      62.0.2.1      221      Dynamic  40400
224.5.2.3      V703-1/4      62.0.3.1      217      Dynamic  40400
224.5.2.4      V704-1/4      62.0.4.1      223      Dynamic  40400

4 out of 4 group Receivers displayed

Total number of unique groups 2

```

Variable definitions

The following table defines parameters for the **show ip igmp interface** command.

Variable	Value
<i>gigabitethernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>vlan <1-4059></i>	Specifies the VLAN.
<i>vrf WORD<1-16></i>	Specifies the VRF by name.
<i>vrfids WORD<0-512></i>	Specifies the VRF by VRF ID.

The following table defines parameters for the **show ip igmp cache** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF by name.
<i>vrfids WORD<0-512></i>	Specifies the VRF by VRF ID.

The following table defines parameters for the **show ip igmp group** command.

Variable	Value
<i>count</i>	Specifies the number of entries.
<i>group {A.B.C.D}</i>	Specifies the group address.
<i>member-subnet {A.B.C.D/X}</i>	Specifies the IP address and network mask.
<i>vrf WORD<1-16></i>	Displays the multicast route configuration for a particular VRF by name.
<i>vrfids WORD<0-512></i>	Displays the multicast route configuration for a particular VRF by VRF ID.

The following table defines parameters for the **show ip igmp sender** command.

Variable	Value
<i>count</i>	Specifies the number of entries.
<i>group {A.B.C.D}</i>	Specifies the group address.
<i>member-subnet {A.B.C.D/X}</i>	Specifies the IP address and network mask.

Variable	Value
<code>vrf WORD<1-16></code>	Displays the multicast route configuration for a particular VRF by name.
<code>vrfids WORD<0-512></code>	Displays the multicast route configuration for a particular VRF by VRF ID.

Troubleshoot IP Multicast over Fabric Connect for IP Shortcuts

If traffic is not moving properly, use the following procedure to determine the issue.

Procedure

1. Ensure that all switch nodes in the network operate with the most recent software release to support IP Multicast over Fabric Connect:

```
show software
```
2. Ensure that you create and enable SPBM infrastructure globally.
 - a. Ensure that SPBM is enabled globally:

```
show spbm
```
 - b. Ensure that IS-IS is enabled globally:

```
show isis
```
 - c. Ensure an SPBM instance exists and at least one Backbone VLAN exists (B-VID). Also ensure multicast is enabled:

```
show isis spbm
```
3. Ensure that you enable the CFM configuration.
 - a. Ensure a CFM maintenance-association exists:

```
show cfm maintenance-association
```
 - b. Ensure a CFM maintenance-domain exists:

```
show cfm maintenance-domain
```
 - c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

```
show cfm maintenance-endpoint
```
4. Ensure the following on all BEBs where you want IP Multicast over Fabric Connect. Ensure that you enable IP Multicast over Fabric Connect globally:

```
show isis spbm
```
5. On the VLANs that need Layer 3 VSN IP Multicast over Fabric Connect routing, create an IP interface on the VLAN if one does not exist. The address should be on the same subnet as the IGMP hosts connected to the VLAN. Also, ensure that you enable IP Multicast over Fabric Connect. Create an IP interface on the VLAN and enable IP Multicast over Fabric Connect.
6. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

7. Create an IP interface on the VLAN and enable IP Multicast over Fabric Connect:

```
ip address <A.B.C.D>
```

```
ip spb-multicast enable
```

8. Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

```
show ip igmp snoop-trace
```

```
show ip igmp interface
```

9. Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

```
show ip igmp interface
```

IP Shortcuts configuration using EDM

This section provides procedures to configure IP Shortcuts using EDM.

Configure IP Multicast over Fabric Connect on a VLAN Within the GRT

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.
- If no IP interface exists on the VLAN, create one. The IP interface must be the same subnet as the IGMP hosts that connect to the VLAN.

About This Task

Use this procedure to configure IP Multicast over Fabric Connect within the GRT. The default is disabled.

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

The switch only supports IPv4 addresses with IP Multicast over Fabric Connect.

To configure a VRF with IP Multicast over Fabric Connect, see [Configure IP Multicast over Fabric Connect on a VLAN for Layer 3 VSN](#) on page 2027.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Choose a VLAN, and then select **IP**.

4. Select the **SPB Multicast** tab.

**Note**

After you enable IP Multicast over Fabric Connect, you cannot enable IGMP, IGMP Snooping, or IGMP proxy on the interface. If you try to enable IGMP Snooping or proxy on any interface where SPBM multicast is enabled, the system displays an error message.

When you enable IP Multicast over Fabric Connect on a Controller switch in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

5. Select **Enable**.
6. Select **Apply**.

Configure IP Multicast over Fabric Connect on a Brouter Port within the GRT

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.
- If no IP interface exists on the VLAN, create one. The IP interface must be the same subnet as the IGMP hosts that connect to the VLAN.

About This Task

Use this procedure to enable IP Multicast over Fabric Connect on a brouter port IP interface. The default is enabled.

To configure a brouter port for a VRF with IP Multicast over Fabric Connect, see [Configure IP Multicast over Fabric Connect on a Brouter Port for a Layer 3 VSN](#) on page 2028.

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

The switch only supports IPv4 addresses with IP Multicast over Fabric Connect.

Procedure

1. Select an enabled port on the Physical Device View.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **SPB Multicast** tab.
5. Select **Enable**.

**Note**

When you enable IP Multicast over Fabric Connect on a DvR Controller switch in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

6. Select **Apply**.

SPB Multicast Field Description

Use the data in the following table to use the **SPB Multicast** tab.

Name	Description
Enable	Enables or disables SPB Multicast. The default is disable.

Configure the GRT Timeout To Age Out a Multicast Sender

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.

About This Task

Use this procedure to configure the timeout value in the GRT. The timeout value ages out the sender when there are no multicast streams coming from the sender for a specified period of time. The default timeout value is 210 seconds.

Procedure

1. In the navigation pane, expand **Configuration > Fabric > SPBM**.
2. Select the **SPBM** tab.
3. Modify the **McastFwdCacheTimeout** value.
4. Select **Apply**.

Configure the VRF Timeout To Age Out a Multicast Sender

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.

About This Task

Use this procedure to configure the VRF timeout value. The timeout value ages out the sender when there is no multicast stream on the VRF. The default is 210 seconds.



Note

You can use this procedure for Layer 3 VSN with IP Multicast over Fabric Connect services and IP Multicast over Fabric Connect for IP Shortcuts.

Procedure

1. In the navigation pane, expand **Configuration > IP**

2. Select **IP-MVPN**.
3. Select the **MVPN** tab.
4. In the row for the appropriate VRF, type a value in **FwdCacheTimeout**.
5. Select **Apply**.

IP multicast over SPBM within the GRT configuration example

The following example shows the configuration steps to enable IP multicast over SPBM support on VLANs 10 and 11 that are part of the GRT:

```
ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VLAN CONFIGURATION - PHASE I

interface vlan 500
ip address 192.0.2.1 255.255.255.0
ip spb-multicast enable
exit

interface vlan 501
ip address 192.0.2.2 255.255.255.0
ip spb-multicast enable
exit
```

Layer 3 VSN Fundamentals for IP Multicast over Fabric Connect

This section provides fundamentals concepts for Layer 3 VSN configuration for IP Multicast over Fabric Connect. For information on Layer 3 VSN basic configuration, see [Layer 3 VSN Configuration](#) on page 1654.

Layer 3 VSN with IP Multicast over Fabric Connect

IP Multicast over Fabric Connect supports Layer 3 VSN functionality where multicast traffic is bridged over the SPBM core infrastructure. Layer 3 VSN using IP Multicast over Fabric Connect is helpful when you need complete security and total isolation of data. No one outside of the Layer 3 VSN can join or even see the Layer 3 VSN. Applications that can use Layer 3 VSN with IP Multicast over Fabric Connect include:

- Video surveillance
- TV/Video/Ticker/Image Distribution
- VX-LAN
- Multi-tenant IP multicast

Configure the Layer 3 VSN (VRF) as a multicast VPN, and then enable IP Multicast over Fabric Connect on VRF VLANs to which IP multicast senders and receivers attach. This configuration automatically enables IGMP snooping and proxy on those VLANs. IGMPv2 at the VLAN level is the default setting, with no other configuration required. If you want to use IGMPv3, you must configure IGMPv3.

IP Multicast over Fabric Connect is only configured on BEBs.

**Note**

- You do not need to enable IP Shortcuts to support multicast routing in the Layer 3 VSN using SPBM. IPVPN creation and I-SID assignment for the IPVPN is required, but you do not need to enable IPVPN.
- If you only want to use IP Multicast over Fabric Connect, you do not need to enable the Layer 3 VSN or redistribute unicast routes into or out of IS-IS. IP Multicast over Fabric Connect routing does not depend on unicast routing for Layer 3 VSNs using VRFs, which allows you to more easily migrate from a PIM environment to IP Multicast over Fabric Connect. You can migrate a PIM environment to IP Multicast over Fabric Connect first and then migrate unicast separately or not at all.
- If no IP interface exists on the VLAN, then you create one. (The IP interface must be the same subnet as the IGMP hosts that connect to the VLAN).

With Layer 3 VSN with IP Multicast over Fabric Connect, multicast traffic remains in the same Layer 3 VSN across the SPBM cloud. For a Layer 3 VSN, traffic can cross VLAN boundaries but remains confined to the subset of VLANs within the VRF that has IP Multicast over Fabric Connect enabled. If a sender transmits a multicast stream to a BEB on a Layer 3 VSN with IP Multicast over Fabric Connect enabled, only receivers that are part of the same Layer 3 VSN can receive that stream.

All or a subset of VLANs within a Layer 3 VSN can exchange multicast traffic. The BEB only sends out traffic for a multicast stream on which IGMP joins and reports are received.

I-SIDs

After a BEB receives IP multicast data from a sender, the BEB allocates a data service instance identifier (I-SID) in the range of 16,000,000 to 16,512,000 for the multicast stream. The stream is identified by the S, G, V tuple, which is the source IP address, the group IP address and the local VLAN the multicast stream is received on. The data I-SID uses Tx/Rx bits to signify whether the BEB uses the I-SID to transmit, receive, or both transmit and receive data on that I-SID.

In the context of Layer 3 VSNs with IP Multicast over Fabric Connect, the scope is the I-SID value of the Layer 3 VSN associated with the local VLAN that the IP multicast data was received on.

TLVs

See [TLVs](#) on page 1957.

IGMP

After a BEB receives an IGMP join message from a receiver, the BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the requested stream exists, the BEB sends an IS-IS

TLV update to its neighbors to inform them of the presence of a receiver and this information is propagated through the SPBM cloud.

DvR

On DvR Controllers in a DvR domain, you must manually configure IP multicast over Fabric Connect on Layer 3 VSNs (VRFs). This configuration is then automatically pushed to the Leaf nodes in the DvR domain.

For more information on DvR, see [Distributed Virtual Routing](#) on page 915.

Layer 3 VSN Configuration using the CLI

This section provides procedures to configure Layer 3 VSN using the command line interface (CLI).

Configure Layer 3 VSN with IP Multicast over Fabric Connect

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect globally.
- Assign an I-SID for the IPVPN.

About This Task

Use this procedure to configure IP Multicast over Fabric Connect for a Layer 3 VSN.

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

The IP VPN does not need to be enabled for Layer 3 VSN multicast to function.

The switch only supports IPv4 multicast traffic.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable Layer 3 VSN IP Multicast over Fabric Connect for a particular VRF:

```
mvpn enable
```

The default is disabled.

3. Exit to Global Configuration mode:

```
exit
```

4. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

5. Enable Layer 3 VSN IP Multicast over Fabric Connect for a particular VRF:

```
ip spb-multicast enable
```

6. (Optional) Enable IGMP version 3:

```
ip igmp snooping
```

```
ip igmp ssm-snoop
```

```
ip igmp compatibility-mode
```

```
ip igmp version 3
```

**Note**

IGMPv2 at the VLAN level is the default setting, with no other configuration required. You only need to use these commands if you use IGMPv3. You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

7. (Optional) Enable the IGMP Layer 2 Querier address:

```
ip igmp snoop-querier-addr {A.B.C.D}
```

**Note**

If the SPBM bridge connects to an edge switch, it can be necessary to add an IGMP query address. If you omit adding a query address, the SPB bridge sends IGMP queries with a source address of 0.0.0.0. Some edge switch models do not accept a query with a source address of 0.0.0.0.

Example

Configure IP Multicast over Fabric Connect for a Layer 3 VSN:

```
Switch:>enable
Switch:#configure terminal
Switch:(config)# router vrf green
Switch:(config-vrf)#mvpn enable
Switch:(config)#exit
Switch:(config)#interface vlan 500
Switch:(config-if)#ip spb-multicast enable
```

Variable Definitions

The following table defines parameters for the **router vrf** command.

Variable	Value
<i>WORD</i> <1-16>	Specifies the name of the VRF.

The following table defines parameters for the **interface vlan** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

The following table defines parameters for the **GigabitEthernet** command.

Variable	Value
<i>GigabitEthernet</i> { <i>slot/port</i> [/ <i>sub-port</i>][<i>-slot/port</i> [/ <i>sub-port</i>]][, <i>...</i>][<i>slot/all</i>][<i>all</i>]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

The following table defines parameters for the **ip igmp** command.

Variable	Value
<i>access-list WORD<1-64> {A.B.C.D/X} <eny-tx deny-rx deny-both allow-only-tx allow-only-rx allow-only-both></i>	<p>Specifies the name of the access list from 1–64 characters.</p> <p>Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.</p> <p>Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic.</p>
<i>compatibility-mode</i>	<p>Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command:</p> <pre>default ip igmp compatibility-mode</pre> <p>, or use the no option to disable compatibility mode:</p> <pre>no ip igmp compatibility-mode</pre>
<i>dynamic-downgrade-version</i>	<p>Configures if the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:</p> <pre>default ip igmp dynamic-downgrade-version</pre> <p>or use the no option to disable downgrade:</p> <pre>no ip igmp dynamic-downgrade-version</pre>
<i>igmpv3-explicit-host-tracking</i>	Enables explicit host tracking on IGMPv3. The default state is disabled.
<i>immediate-leave</i>	Enables fast leave on a VLAN.
<i>immediate-leave-members {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</i>	Configures IGMP fast leave members on a VLAN to specify fast-leave-capable ports.

Variable	Value
<code>last-member-query-interval <0-255></code>	<p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decreasing the value reduces the time to detect the loss of the last member of a group. The default is 10 tenths of a second. Configure this value between 3–10 (equal to 0.3 – 1.0 seconds).</p>
<code>mrdisc [maxadvertinterval <2-180>] [maxinitadvertinterval <2-180>] [maxinitadvertisements <2-15>] [minadvertinterval <3-180>] [neighdeadinterval <2-180>]</code>	<p>Configure the multicast router discovery options to enable the automatic discovery of multicast capable routers. The default parameter values are:</p> <ul style="list-style-type: none"> • maxadvertinterval: 20 seconds • maxinitadvertinterval: 2 seconds • maxinitadvertisements: 3 • minadvertinterval: 15 seconds • neighdeadinterval: 60 seconds
<code>mrouter {slot/port[/sub-port]}[-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	Adds multicast router ports.
<code>proxy</code>	Activates the proxy-snoop option globally for the VLAN.
<code>query-interval <1-65535></code>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
<code>query-max-response <0-255></code>	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds).</p> <p>Important: You must configure this value lower than the query-interval.</p>
<code>robust-value <2-255></code>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.

Variable	Value
<code>router-alert</code>	<p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p>Important: To maximize network performance, configure this parameter according to the version of IGMP currently in use:</p> <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
<code>snoop-querier</code>	Enables the IGMP Layer 2 Querier feature on the VLAN. The default is disabled.
<code>snoop-querier-addr {A.B.C.D}</code>	Specifies the IGMP Layer 2 Querier source IP address.
<code>snooping</code>	Activates the snoop option for the VLAN.
<code>ssm-snoop</code>	Activates support for SSM on the snoop interface.
<code>static-group {A.B.C.D}</code> <code>{A.B.C.D} [port]</code> <code>{slot/port[/sub-port]} [-slot/</code> <code>port[/sub-port]] [,...] [slot/</code> <code>all] [all]} [static blocked]</code>	<p>Configures IGMP static members to add members to a snoop group.</p> <p><code>{A.B.C.D} {A.B.C.D}</code> indicates the IP address range of the selected multicast group.</p> <p><code>[port] {slot/port[/sub-port]} [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code> adds ports to a static group entry.</p> <p><code>[static blocked]</code> configures the route to static or blocked.</p>
<code>stream-limit stream-limit-max-streams <0-65535></code>	Configures multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. The default is 4.
<code>stream-limit-group</code> <code>{slot/port[/sub-port]} [-slot/</code> <code>port[/sub-port]] [,...] [slot/</code> <code>all] [all]} enable max-</code> <code>streams <0-65535></code>	Configures multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN. The default <code>max-streams</code> value is 4.
<code>version <1-3></code>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

View Layer 3 VSN with IP Multicast over Fabric Connect Information

Use the following options to display Layer 3 VSN with IP Multicast over Fabric Connect information to confirm proper configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all the VRFs that have MVPN enabled and their corresponding forward cache timeout values:

```
show ip vrf mvpn
```
3. Display IP Multicast over Fabric Connect route information:

```
show isis spbm ip-multicast-route [all][detail]
```
4. Display IP Multicast over Fabric Connect by group and source address:

```
show isis spbm ip-multicast-route [group {A.B.C.D}][detail][source {A.B.C.D}]
```
5. Display IP Multicast over Fabric Connect route information by VRF:

```
show isis spbm ip-multicast-route [vrf WORD<1-16>] [group {A.B.C.D}]
```
6. Display IP Multicast over Fabric Connect route information by VLAN:

```
show isis spbm ip-multicast-route [vlan <1-4059>][detail][group {A.B.C.D}]
```
7. Display IP Multicast over Fabric Connect information by VSN I-SID:

```
show isis spbm ip-multicast-route [vsn-isid <1-16777215>][detail][group {A.B.C.D}]
```
8. Display summary information for each S, G, V tuple with the corresponding scope, Data I-SID, and the host name of the source:

```
show isis spb-mcast-summary [count][host-name WORD<0-255>][lspid <xxxx.xxxx.xxxx.xx-xx>]
```

Example

Display Layer 3 VSN with IP Multicast over Fabric Connect information:

```
Switch:1>enable
Switch:1#show ip vrf mvpn
```

```

      Vrf name : green
      mvpn : enable
      fwd-cache-timeout(seconds) : 210

      Vrf name : 4
      mvpn : enable
      fwd-cache-timeout(seconds) : 210

      Vrf name : blue
      mvpn : enable
      fwd-cache-timeout(seconds) : 210

```

```
Switch:1#show isis spbm ip-multicast-route all
```

```

=====
SPBM IP-multicast ROUTE INFO ALL
=====

```

Type	VrfName	Vlan Id	Source	Group	VSN-ISID	Data ISID	BVLAN	Source-BEB
routed	GRT	501	192.0.2.1	233.252.0.1	5010	16300001	10	e12
routed	GRT	501	192.0.2.1	233.252.0.2	5010	16300002	20	e12
routed	GRT	501	192.0.2.1	233.252.0.3	5010	16300003	10	e12
routed	GRT	501	192.0.2.1	233.252.0.4	5010	16300004	20	e12

```

routed GRT      501 192.0.2.1 233.252.0.5 5010      16300005 10    e12
routed GRT      501 192.0.2.1 233.252.0.6 5010      16300006 20    e12
routed GRT      501 192.0.2.1 233.252.0.7 5010      16300007 10    e12
routed GRT      501 192.0.2.1 233.252.0.8 5010      16300008 20    e12
routed GRT      501 192.0.2.1 233.252.0.9 5010      16300009 10    e12
routed GRT      501 192.0.2.1 233.252.0.10 5010     16300010 20    e12

```

```

-----
Total Number of SPBM IP multicast ROUTE Entries: 10
-----

```

```

Switch:1#show isis spbm ip-multicast-route vrf green

```

```

=====
                        SPBM IP-MULTICAST ROUTE INFO
=====
Source          Group          Data ISID  BVLAN  Source-BEB
-----
192.0.2.10      233.252.0.1      16300001  10     e12
192.0.2.10      233.252.0.2      16300002  20     e12
192.0.2.10      233.252.0.3      16300003  10     e12
192.0.2.10      233.252.0.4      16300004  20     e12
192.0.2.10      233.252.0.5      16300005  10     e12
192.0.2.10      233.252.0.6      16300006  20     e12
192.0.2.10      233.252.0.7      16300007  10     e12
192.0.2.10      233.252.0.8      16300008  20     e12
192.0.2.10      233.252.0.9      16300009  10     e12
192.0.2.10      233.252.0.10     16300010  20     e12

```

```

-----
Total Number of SPBM IP MULTICAST ROUTE Entries: 10
-----

```

```

Switch:1#show isis spbm ip-multicast-route vlan 501

```

```

=====
                        SPBM IP-multicast ROUTE INFO ALL
=====
Type VrfName Vlan  Source      Group          VSN-ISID Data ISID  BVLAN Source-BEB
-----
Id
-----
routed GRT   501 192.0.2.1 233.252.0.1 5010      16300001 10    e12
routed GRT   501 192.0.2.1 233.252.0.2 5010      16300002 20    e12
routed GRT   501 192.0.2.1 233.252.0.3 5010      16300003 10    e12
routed GRT   501 192.0.2.1 233.252.0.4 5010      16300004 20    e12
routed GRT   501 192.0.2.1 233.252.0.5 5010      16300005 10    e12
routed GRT   501 192.0.2.1 233.252.0.6 5010      16300006 20    e12
routed GRT   501 192.0.2.1 233.252.0.7 5010      16300007 10    e12
routed GRT   501 192.0.2.1 233.252.0.8 5010      16300008 20    e12
routed GRT   501 192.0.2.1 233.252.0.9 5010      16300009 10    e12
routed GRT   501 192.0.2.1 233.252.0.10 5010     16300010 20    e12

```

```

-----
Total Number of SPBM IP multicast ROUTE Entries: 10
-----

```

```

Switch:1# show isis spbm ip-multicast-route vsn-isid 5010

```

```

=====
SPBM IP-multicast ROUTE INFO - VLAN ID : 501, VSN-ISID : 5010
=====
Source          Group          Data ISID  BVLAN Source-BEB
-----
192.0.2.1      233.252.0.2      16300002  20    e12
192.0.2.1      233.252.0.3      16300003  10    e12
192.0.2.1      233.252.0.4      16300004  20    e12

```

```

192.0.2.1      233.252.0.5    16300005      10    e12
192.0.2.1      233.252.0.6    16300006      20    e12
192.0.2.1      233.252.0.7    16300007      10    e12
192.0.2.1      233.252.0.8    16300008      20    e12
192.0.2.1      233.252.0.9    16300009      10    e12
192.0.2.1      233.252.0.10   16300010      20    e12

```

```

-----
Total Number of SPBM IP multicast ROUTE Entries: 10
-----

```

```
Switch:1# show isis spb-mcast-summary
```

```

=====
                        SPB multicast - Summary
=====
SCOPE      SOURCE      GROUP      DATA      LSP   HOST
I-SID      ADDRESS     ADDRESS     I-SID      BVID   FRAG  NAME
-----
5010       192.0.2.1   233.252.0.1 16300001   10     0x0   e12
5010       192.0.2.1   233.252.0.3 16300003   10     0x0   e12
5010       192.0.2.1   233.252.0.5 16300005   10     0x0   e12
5010       192.0.2.1   233.252.0.7 16300007   10     0x0   e12
5010       192.0.2.1   233.252.0.9 16300009   10     0x0   e12
5010       192.0.2.1   233.252.0.2 16300002   20     0x0   e12
5010       192.0.2.1   233.252.0.4 16300004   20     0x0   e12
5010       192.0.2.1   233.252.0.6 16300006   20     0x0   e12
5010       192.0.2.1   233.252.0.8 16300008   20     0x0   e12
5010       192.0.2.1   233.252.0.10 16300010   20     0x0   e12

```

Variable Definitions

The following table defines parameters for the **show isis spbm ip-multicast-route** command.

Variable	Value
<i>all</i>	Displays all IP Multicast over Fabric Connect route information.
<i>detail</i>	Displays detailed IP Multicast over Fabric Connect route information.
<i>group</i> {A.B.C.D}	Displays information on the group IP address for the IP Multicast over Fabric Connect route.
<i>vlan</i> <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>vrf</i> WORD<1-16>	Displays IP Multicast over Fabric Connect route information by VRF.
<i>vsn-isid</i> <1-16777215>	Displays IP Multicast over Fabric Connect route information by I-SID.

The following table defines parameters for the **show isis spb-mcast-summary** command.

Variable	Value
count	Displays the total number of SPB multicast entries.
host-nameWORD<0-255>	Displays the IP Multicast over Fabric Connect summary information by host-name.
lspid<xxxx.xxxx.xxxx.xx-xx>	Displays the IP Multicast over Fabric Connect summary information by LSP ID.

View IGMP Information for Layer 3 VSN Multicast

Use the following commands to check IGMP information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display information about the interfaces where IGMP is enabled:

```
show ip igmp interface [gigabitethernet {slot/port[-slot/port]}[,...]  
[slot/all][all]] [vlan <1-4059>[vrf WORD<1-16>][vrfids WORD<0-512>]
```

Ensure that the output displays routed-spb under MODE.

3. Display information about the IGMP cache:

```
show ip igmp cache [vrf WORD<1-16>][vrfids WORD<0-512>]
```

4. Display information about the IGMP group:

```
show ip igmp group [count][member-subnet default|{A.B.C.D/X}][vrf  
WORD<1-16>][vrfids WORD<0-512>]
```

5. Display information about the IGMP sender:

```
show ip igmp sender [count][member-subnet default|{A.B.C.D/X}][vrf  
WORD<1-16>][vrfids WORD<0-512>]
```

Example

Display IGMP information for Layer 3 VSN with IP multicast over Fabric Connect:

```
Switch:#enable
Switch:1#show ip igmp interface vrf green

=====
                        Igmp Interface - GlobalRouter
=====
      QUERY      OPER      QUERY      WRONG      LASTMEM
IF    INTVL STATUS VERS.  VERS  QUERIER  MAXRSPT QUERY JOINS ROBUST QUERY MODE
-----
V100  125    activ  2     2    0.0.0.0   100    0     0     2     10   routed-spb

1 out of 1 entries displayed
```

```
Switch:1(config)#show ip igmp interface vlan 20
=====
                        Vlan Ip Igmp
=====
VLAN QUERY QUERY ROBUST VERSION VERSION LAST  PROXY  SNOOP  SNOOP      SSM    UPnP    FAST    FAST    FAST
ID   INTVL MAX          ORIGIN  MEMB   SNOOP  ENABLE ORIGIN  SNOOP  FILTER LEAVE  LEAVE  LEAVE
```

```

      RESP          QUERY ENABLE          ENABLE ENABLE ENABLE ORIGIN  PORTS
-----
20  125  100  2      2      CONFIG 10   true  true  RADIUS          false true  false  RADIUS
VLAN SNOOP          SNOOP          DYNAMIC  COMPATIBILITY  EXPLICIT HOSTS  UPnP
ID   QUERIER        QUERIER        DOWNGRADE  MODE          TRACKING        FILTER
      ENABLE ORIGIN  ADDRESS  ORIGIN  VERSION  ENABLE  ORIGIN  ENABLE  ORIGIN  ADDRESS
-----
20   true  RADIUS    192.0.2.1 RADIUS  enable  disable RADIUS  disable RADIUS  198.51.100.0/24

Switch:1# show ip igmp sender vrf vrf1
=====
                IGMP Sender - VRF vrf1
=====
                                PORT/
GRPADDR  IFINDEX  MEMBER  MLT   STATE      L2ISID SCOPEISID  DATAISID
-----
224.1.1.4  Vlan 10  192.2.0.1  1/18/1  NOTFILTERED  0  10010  16000001
224.1.1.4  Vlan 10  192.2.0.1  1/18/1  NOTFILTERED  0  10010  16000004
224.1.1.6  Vlan 10  192.2.0.1  1/18/1  FILTERED     0  10010  16000010
224.1.1.7  Vlan 10  192.2.0.1  1/18/1  FILTERED     0  10010  16000009
224.1.1.8  Vlan 10  192.2.0.1  1/18/1  NOTFILTERED  0  10010  210

5 out of 5 entries displayed
Switch:1# show ip igmp group vrf green
=====
                IGMP Group - GlobalRouter
=====
GRPADDR      INPORT      MEMBER      EXPIRATION TYPE
-----
233.252.0.1  V501-1/16   192.2.0.1   204      Dynamic
233.252.0.2  V501-1/16   192.2.0.1   206      Dynamic
233.252.0.3  V501-1/16   192.2.0.1   206      Dynamic
233.252.0.4  V501-1/16   192.2.0.1   207      Dynamic
233.252.0.5  V501-1/16   192.2.0.1   204      Dynamic
233.252.0.6  V501-1/16   192.2.0.1   209      Dynamic
233.252.0.7  V501-1/16   192.2.0.1   206      Dynamic
233.252.0.8  V501-1/16   192.2.0.1   206      Dynamic
233.252.0.9  V501-1/16   192.2.0.1   211      Dynamic
233.252.0.10 V501-1/16   192.2.0.1   207      Dynamic

10 out of 10 group Receivers displayed

Total number of unique groups 10

```

Variable Definitions

The following table defines parameters for the **show ip igmp interface** command.

Variable	Value
<i>gigabitethernet</i> <i>{slot/port[/sub-port]</i> <i>[-slot/port[/sub-port]][,...][slot/all][all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>vlan</i> <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>vrf</i> WORD<1-16>	Specifies the VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies the VRF by VRF ID.

The following table defines parameters for the **show ip igmp cache** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies the VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies the VRF by VRF ID.

The following table defines parameters for the **show ip igmp group** command.

Variable	Value
<i>count</i>	Specifies the number of entries.
<i>group</i> {A.B.C.D}	Specifies the group address.
<i>member-subnet</i> {A.B.C.D/X}	Specifies the IP address and network mask.
<i>vrf</i> WORD<1-16>	Displays the multicast route configuration for a particular VRF by name.
<i>vrfids</i> WORD<0-512>	Displays the multicast route configuration for a particular VRF by VRF ID.

The following table defines parameters for the **show ip igmp sender** command.

Variable	Value
<i>count</i>	Specifies the number of entries.
<i>group</i> {A.B.C.D}	Specifies the group address.

Variable	Value
<code>member-subnet {A.B.C.D/X}</code>	Specifies the IP address and network mask.
<code>vrf WORD<1-16></code>	Displays the multicast route configuration for a particular VRF by name.
<code>vrfids WORD<0-512></code>	Displays the multicast route configuration for a particular VRF by VRF ID.

Troubleshoot IP Multicast over Fabric Connect for Layer 3 VSNs

If traffic is not moving properly, use the following procedure to determine the issue.

Procedure

1. Ensure that all switch nodes in the network operate with the most recent software release to support IP Multicast over Fabric Connect:

```
show software
```

2. Ensure that you create and enable SPBM infrastructure globally.

- a. Ensure that SPBM is enabled globally:

```
show spbm
```

- b. Ensure that IS-IS is enabled globally:

```
show isis
```

- c. Ensure an SPBM instance exists with at least one Backbone VLAN (B-VID). Also ensure multicast is enabled:

```
show isis spbm
```

3. Ensure that you enable the CFM configuration.

- a. Ensure a CFM maintenance-association exists:

```
show cfm maintenance-association
```

- b. Ensure a CFM maintenance-domain exists:

```
show cfm maintenance-domain
```

- c. Ensure a maintenance-endpoint exists in the MEP ID column and is enabled in the ADMIN column:

```
show cfm maintenance-endpoint
```

4. Ensure the following on all the Backbone Edge Bridges (BEBs) where the Layer 3 VSN is present.

- a. Ensure that you enable IP multicast globally:

```
show isis spbm
```

- b. Ensure that you create an IPVPN for the VRF:

```
show ip ipvpn [vrf WORD<1-16>][vrfids WORD<0-512>]
```

- c. Ensure that you assign an I-SID to the VRF:

```
show isis spbm ip-multicast-route all
```

- d. Ensure that you enable the MVPN:

```
show ip vrf mvpn
```

- On the VLANs that need Layer 3 VSN IP Multicast over Fabric Connect routing, create an IP interface on the VLAN if one does not exist. The address should be on the same subnet as the IGMP hosts connected to the VLAN. Also, ensure that you enable IP Multicast over Fabric Connect.

- Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

- Create an IP interface on the VLAN and enable IP Multicast over Fabric Connect:

```
ip address <A.B.C.D>
```

```
ip spb-multicast enable
```

- Ensure that you enable IGMP Snooping on access Layer 2 switches to prevent flooding of multicast traffic to non-receiver ports:

```
show ip igmp snoop-trace
```

```
show ip igmp interface
```

- Ensure that the IGMP version used by the multicast hosts and the Layer 2 switches outside the SPBM network is the same as the IGMP version configured on the C-VLAN:

```
show ip igmp interface
```

Layer 3 VSN Configuration using EDM

This section provides procedures to configure Layer 3 VSN using Enterprise Device Manager (EDM).

Enable MVPN for a VRF

Use this procedure to enable MVPN for a particular VRF. IP Multicast over Fabric Connect, constrains multicast streams of senders to all receivers in the same Layer 3 VSN. MVPN functionality is disabled by default.



Note

VLAN level configuration is also required to turn on the service on each VLAN within the VRF on which this services is required. You can turn it on under the VLAN context or the brouter context.

Before You Begin

- You must enable IP Multicast over Fabric Connect globally.

Procedure

- In the navigation pane, expand **Configuration > IP**.
- Select **IP-MVPN**.
- Select the **MVPN** tab.
- Double-click in the **Enable** field in the table.

5. Select **Enable** from the drop down menu.
6. Double-click in the **FwdCacheTimeout** field in the table, and then type the VRF timeout value.
7. Select **Apply**.

MVPN Field Descriptions

Use the data in the following table to use the **MVPN** tab.

Name	Description
Vrflid	Specifies the VRF ID.
Enable	Enables Layer 3 VSN IP Multicast over Fabric Connect services for a particular VRF. The default is disabled.
FwdCacheTimeout	Specifies the VRF timeout value. The timeout value ages out the sender when there is no multicast stream on the VRF. The default is 210 seconds..

Configure IP Multicast over Fabric Connect on a VLAN for Layer 3 VSN

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Configure a VRF and an IP VPN instance with an I-SID configured under it on the switch. The IP VPN does not need to be enabled for Layer 3 VSN multicast to function.
- Enable IP Multicast over Fabric Connect globally.
- If there is no IP interface on the VLAN, then create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).
- Enable MVPN for the particular VRF.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Use this procedure to enable IP Multicast over Fabric Connect for a Layer 3 VSN. The default is disabled.

You must configure VLANs to turn on the service on each VLAN within the VRF on which the service is required. You can turn it on under the VLAN context or the router context.

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

The switch only supports IPv4 address with IP Multicast over Fabric Connect.

To configure a VLAN for IP Shortcuts with IP Multicast over Fabric Connect, see [Configure IP Multicast over Fabric Connect on a VLAN Within the GRT](#) on page 2008.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN, and then select **IP**.
5. Select the **SPB Multicast** tab.
6. Select **Enable**.
7. Select **Apply**.

SPB Multicast Field Descriptions

Use the data in the following table to use the **SPB Multicast** tab.

Name	Description
Enable	Enables SPB Multicast on the VLAN. The default is disabled.

Configure IP Multicast over Fabric Connect on a Brouter Port for a Layer 3 VSN

Before You Begin

- Configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Configure a VRF and an IP VPN instance with an I-SID configured under it on the switch. The IP VPN does not need to be enabled for Layer 3 VSN multicast to function.
- Enable IP Multicast over Fabric Connect globally.
- If there is no IP interface on the VLAN, then you create one. (The IP interface must be in the same subnet as the IGMP hosts that connect to the VLAN).
- Enable MVPN for the particular VRF.

About This Task

Use this procedure to enable IP Multicast over Fabric Connect on a brouter port. The default is disabled.

You must enable IP Multicast over Fabric Connect on each of the VLANs that need to support IP multicast traffic.

The switch only supports IPv4 address with IP Multicast over Fabric Connect.

To configure a brouter port for IP Shortcuts with IP Multicast over Fabric Connect, see [Configure IP Multicast over Fabric Connect on a Brouter Port within the GRT](#) on page 2009.

Procedure

1. In the navigation pane, expand **Configuration > VRF Context View**.
2. Select **Set VRF Context View**.
3. Choose a VRF name.
4. Select **Launch VRF Context View**.
5. Select an enabled port on the Physical Device View.
6. In the navigation pane, expand **Configuration > Edit > Port**.
7. Select **IP**.
8. Select the **SPB Multicast** tab.
9. Select **Enable**.
10. Select **Apply**.

Configure IGMP on a VLAN Interface for a Layer 3 VSN

Before You Begin

- Configure the required SPBM IS-IS infrastructure.
- Configure a VRF and IP VPN instance with an I-SID on the switch.
- Create the C-VLANs and add slots/ports.
- Enable IP Multicast over Fabric Connect for a Layer 3 VSN.
- Select and launch the VRF instance associated with the Layer 3 VSN. For more information about how to use EDM to configure a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100.

About This Task

Use this procedure to configure IGMP for each VLAN interface to enable the interface to perform multicast operations.

IGMPv2 at the VLAN level is the default setting, with no other configuration required. You only need to enable IGMPv3. You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is the same as the IGMP version configured on the IGMP Snooping VLAN, or that you enable compatibility mode.

You cannot enable IP PIM when IP Multicast over Fabric Connect is enabled on the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN.
5. Select **IP**.
6. Select **IGMP** tab.
7. (Optional) Select **SsmSnoopEnable**.

8. (Optional) Select **SnoopEnable**.

9. (Optional) In **Version**, select the correct IGMP version.

You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

10. (Optional) Select **SnoopQuerierEnable**.

Only select this option, if you want to configure an address for the IGMP queries.

11. (Optional) In **SnoopQuerierAddr**, type an IP address.



Note

If the SPBM bridge connects to an edge switch, it can be necessary to add an IGMP query address. If you omit adding a query address, the SPB bridge sends IGMP queries with a source address of 0.0.0. Some edge switch models do not accept a query with a source address of 0.0.0.0.

12. Select **Apply**.

IGMP Field Descriptions

Use the data in the following table to use the **IGMP** tab.

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the IGMP host query packets transmit on the interface. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	<p>Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values enable a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds.)</p> <p>Important: You must configure this value lower than the QueryInterval.</p>
Robustness	<p>Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value.</p> <p>The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.</p>

Name	Description
LastMembQueryIntvl	<p>Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.</p> <p>Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Configure this parameter to values greater than 3. If you do not require a fast leave process, use values greater than 10. (The value 3 is equal to 0.3 seconds, and 10 is equal to 1 second.)</p>
SnoopEnable	Enables snoop on the interface. The default is disabled.
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
ProxySnoopEnable	Enables proxy snoop on the interface. The default is disabled.
Version	Configures the version of IGMP (1, 2, or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables or disables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this VLAN.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this VLAN. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	<p>Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router.</p> <p>Important: Configure this field only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.</p>

Name	Description
RouterAlertEnable	<p>Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.</p> <p>To maximize network performance, configure this parameter according to the version of IGMP currently in use:</p> <ul style="list-style-type: none"> • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
DynamicDowngradeEnable	<p>Configures if the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.</p>
CompatibilityModeEnable	<p>Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.</p>
ExplicitHostTrackingEnable	<p>Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.</p>
SnoopQuerierEnable	<p>Enables snoop querier. The default is disabled.</p> <p>When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups.</p> <p>Enable Layer 2 Querier on only one node in the VLAN.</p>
SnoopQuerierAddr	<p>Specifies the pseudo IP address of the IGMP snoop querier. The default IP address is 0.0.0.0.</p>

Layer 3 VSN with IP Multicast over Fabric Connect Configuration Example

The following example shows the configuration to enable IP Multicast over Fabric Connect support on VLANs 500 and 501 that are part of VRF Green:

```

ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VRF CONFIGURATION

```

```
ip vrf green vrfid 2

VLAN CONFIGURATION - PHASE 1

vlan 110 i-sid 100
interface vlan 500
vrf green
ip address 192.0.2.1 255.255.255.0 1
ip spb-multicast enable
exit

vlan 111 i-sid 100
interface vlan 501
vrf green
ip address 192.0.2.2 255.255.0 0
ip spb-multicast enable
exit

ISIS SPBM IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 100
mvpn enable
exit
```

When using IGMPv3, the configuration is:

```
ISIS SPBM CONFIGURATION

router isis
spbm 1 multicast enable

VRF CONFIGURATION

ip vrf green vrfid 2

VLAN CONFIGURATION - PHASE 1

vlan 110 i-sid 100
interface vlan 500
vrf green
ip address 192.0.2.1 255.255.255.0 1
ip spb-multicast enable
ip igmp version 3
exit

vlan 111 i-sid 100
interface vlan 501
vrf green
ip address 192.0.2.2 255.255.0 0
ip spb-multicast enable
ip igmp version 3
exit

ISIS SPBM IPVPN CONFIGURATION

router vrf green
```

```
ipvpn
i-sid 100
mvpn enable
exit
```

IP SPB Multicast Policy Configuration

You can configure IP SPB Multicast Policy to permit only multicast senders or multicast receivers or deny both for specific IP multicast group addresses.

IP SPB Multicast Policy fundamentals

This section contains conceptual information about IP SPB Multicast Policy .

IP SPB Multicast Policy

The protocols that use IP multicast, for example, Universal Plug and Play (uPnP), cause performance issues in the network as they generate many changes in the Intermediate-System-to-Intermediate-System (IS-IS) Link State Database (LSDB). To address the network issues, the IP SPB Multicast Policy feature aggregates multiple multicast groups (senders or receivers) in a static data I-SID that you configure in GRT/VRF, by associating the IP SPB Multicast Policy with a route-map. You can also configure filtering on the multicast groups based on the traffic flow (senders only or receivers only) using the IP SPB Multicast Policy feature.

When IP SPB Multicast Policy matches the sender's multicast group IP address, the following list identifies potential IP multicast stream behavior:

- Dropped.
- Forwarded as one dynamic data I-SID per stream (pre-existing design).
- Forwarded as multiple IPMC streams mapped to the same static data-ISID.

Restrictions

The following list identifies restrictions of an IP SPB Multicast Policy :

- This feature is not supported on:
 - IGMP snooping
 - Distributed Virtual Routing (DvR) leaf
- Multi-area SPB boundary nodes with remote IS-IS enabled do not support IP SPB Multicast Policy .
- Multi-area SPB boundary nodes do not support static data I-SID configuration. However, it can redistribute the static data I-SID with MAC learning disabled.
- The system ignores the include and exclude lists for IGMPv3 membership reports and forwards all traffic that is received in the static data I-SID to the receiver.
- This feature is not compatible with multicast traffic processed on a PIM GW interface and will not influence its behavior.

Static Data I-SID

To reduce Link State Database (LSDB) changes and Intermediate-System-to-Intermediate-System (IS-IS) processing, the switch advertises or withdraws the static

data I-SIDs from the Shortest Path Bridging (SPB) fabric when you issue the apply command for the routed multicast policy. As a result, the switch triggers most of the IS-IS computation for the I-SIDs by user configuration and not by volatile network changes or UNI IGMP messages.

For Multi-area SPB boundary nodes, the static data I-SID multicast redistribution policy takes precedence over the Layer 2 multicast redistribution.

You can configure the static I-SID redistribution in the Multi-area SPB for any direction (home-to-remote or remote-to-home).

Multi-area SPB boundary nodes do not support static data I-SID configuration.

For more information, see the following sections:

- [Configure Multi-area SPB IPv4 Static I-SID Routed Multicast Redistribution](#) on page 1461
- [Configure Multi-area SPB Layer 3 Static I-SID Routed Multicast Redistribution Policy in Home Area](#) on page 1478
- [Configure Multi-area SPB Layer 3 Static I-SID Routed Multicast Redistribution Policy in Remote Area](#) on page 1479
- [Multi-area SPB](#) on page 1422

IP SPB Multicast Policy Configuration Using CLI

Perform the procedures in this section to configure the IP SPB Multicast Policy using the Command Line Interface (CLI).

Create an IP SPB Multicast Policy

Before You Begin

- Configure Multicast Label Switching Virtual Private Network (MVPN) on VRF.
- Enable IP Multicast over Fabric Connect globally.
- Configure IP prefix list in GRT/VRF.

About This Task

Perform this procedure to create the IP SPB Multicast Policy globally.



Note

- You can also configure the IP SPB Multicast Policy on a VRF instance in the same way as you configure on the GlobalRouter, except that you must use the VRF Router Configuration mode and create the route-map in VRF Router Configuration mode. Not all parameters are configurable on non0 VRFs.
- This configuration does not apply to IGMP snooping and DvR Leaf.
- Multi-area SPB boundary nodes with remote IS-IS enabled do not support IP SPB Multicast Policy .

Procedure

1. Enter Route-Map Configuration mode:

```
enable
```

```
configure terminal
```

```
route-map WORD<1-64> <1-65535>
```
2. Enable the route policy:

```
enable
```
3. Match the IP SPB Multicast Policy to the multicast group:

```
match multicast-group WORD<0-259>
```
4. (Optional) Set the data flow for the IP SPB Multicast Policy :

```
set {rx-only | tx-only}
```
5. Configure the data I-SID:

```
set data-isid <1-15999999>
```
6. Permit the route:

```
permit
```
7. Exit to Global Configuration mode:

```
exit
```
8. Associate the IP SPB Multicast Policy with the route-map:

```
ip spb-multicast-policy route-map WORD<0-64>
```
9. Enable the IP SPB Multicast Policy :

```
ip spb-multicast-policy enable
```
10. Verify the configuration:

```
show ip spb-multicast-policy
```
11. Apply the IP SPB Multicast Policy instance:

```
ip spb-multicast-policy apply
```

Example

Creating an IP SPB Multicast Policy globally.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#route-map rtest 210
Switch:1(route-map)#enable
Switch:1(route-map)#match multicast-group McastGroup
Switch:1(route-map)#set data-isid 210
Switch:1(route-map)#permit
Switch:1(route-map)#exit
Switch:1(config)#ip spb-multicast-policy route-map rtest
Switch:1(route-map)#ip spb-multicast-policy enable
Switch:1(config)#show ip spb-multicast-policy
=====
SPB ROUTED IP MULTICAST POLICY - GlobalRouter
=====
ENABLE      RPOLICY
=====
TRUE        rtest
Switch:1(route-map)#ip spb-multicast-policy apply
```

Creating an IP SPB Multicast Policy on VRF1.

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf VRF1
Switch:1(router-vrf)#route-map rtest 210
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#match multicast-group McastGroup
Switch:1(router-vrf-routemap)#set data i-sid 5679
Switch:1(router-vrf-routemap)#permit
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#ip spb-multicast-policy route-map rtest
Switch:1(router-vrf)#ip spb-multicast-policy enable
Switch:1(router-vrf)#exit
Switch:1(config)#show ip spb-multicast-policy vrf VRF1
=====
                SPB ROUTED IP MULTICAST POLICY - VRF VRF1
=====
ENABLE      RPOLICY
-----
TRUE        rtest
Switch:1(config)#ip spb-multicast-policy apply vrf VRF1

```

Variable Definitions

The following table defines parameters for the **ip spb-multicast-policy** command.

Variable	Value
<i>apply</i>	Applies the IP SPB Multicast Policy on route-map.
<i>enable</i>	Enables the IP SPB Multicast Policy on route-map.
<i>route-map WORD<0-64></i>	Specifies the route of the policy. <i>WORD<0-64></i> specifies the name of the policy.
<i>vrf WORD<1-16></i>	Specifies the VRF name.

The following table defines parameters for the **match** command.

Variable	Value
<i>multicast-groupWORD<0-259></i>	Specifies the multicast group for IP SPB Multicast Policy . <i>WORD<0-259></i> specifies the name of one or more prefix lists, separated by a comma.

The following table defines parameters for the **set** command.

Variable	Value
<i>data-isid</i> <1-15999999>	Specifies the data services instance identifier (I-SID) for static I-SID routed multicast policy.
<i>rx-only</i>	Specifies the received traffic only, applicable to static I-SID routed multicast policy.
<i>tx-only</i>	Specifies the transmitted traffic only, applicable to static I-SID routed multicast policy.

Display IS-IS SPBM Static I-SID IP Multicast Route Information

About This Task

Perform this procedure to display the Intermediate-System-to-Intermediate-System SPBM static I-SID IP multicast route information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the static I-SID IP multicast route information by group IP address:

```
show isis spbm static-isid-ipmc-route group {A.B.C.D} [vrf WORD<1-16>]
```
3. Display the local static I-SID IP multicast route information:

```
show isis spbm static-isid-ipmc-route local [group {A.B.C.D} [vrf WORD<1-16>] | vrf WORD<1-16>]]
```
4. Display the non-local static I-SID IP multicast route information:

```
show isis spbm static-isid-ipmc-route non-local [group {A.B.C.D} [vrf WORD<1-16>] | vrf WORD<1-16>]]
```
5. Display the static I-SID IP multicast route information on a VRF instance:

```
show isis spbm static-isid-ipmc-route vrf WORD<1-16>
```

Example

Displaying the IS-IS SPBM static I-SID IP multicast route information on a VRF instance.

```
Switch:1>show isis spbm static-isid-ipmc-route vrf test
```

```
=====
```

SPBM STATIC ISID IP-MULTICAST ROUTE INFO - VRF test						
Source	Group	Data ISID	BVLAN	NNI	Rcvrs	UNI Rcvrs
Source-BEB						
10.4.1.40	224.1.1.3	210	4051	1/9	-	LOCAL
10.4.1.41	224.1.1.3	210	4051	1/9	-	LOCAL
10.4.1.42	224.1.1.3	210	4051	1/9	-	LOCAL
0.0.0.0	224.1.1.3	210	-	-	V10:1/19/1	NON-LOCAL

```
=====
```

Variable Definitions

The following table defines parameters for the **show isis spbm static-isis-ipmc-route** command.

Variable	Value
<i>group {A.B.C.D}</i>	Displays the static I-SID IP multicast route information. <i>{A.B.C.D}</i> specifies the group IP address.
<i>vrf WORD<1-16></i>	Displays the static I-SID IP multicast route information. <i>WORD<1-16></i> specifies the name of the VRF.
<i>local</i>	Displays the local static I-SID IP multicast route information.
<i>non-local</i>	Displays the non-local static I-SID IP multicast route information.

IP SPB Multicast Policy Configuration using EDM

Perform the procedures in this section to configure the IP SPB Multicast Policy using the Enterprise Device Manager (EDM).

Configure IP SPB Multicast Policy

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Perform this procedure to configure the IP SPB Multicast Policy on the switch.



Note

This functionality is not supported on IGMP snooping and DvR Leaf. Multi-area SPB boundary nodes with remote IS-IS enabled do not support IP SPB Multicast Policy.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.
4. Select **SpbmMulticastPolicyEnable** to enable this functionality.
5. In **SpbmMulticastPolicyRmap**, type the route-map.
6. In **SpbmMulticastPolicyApply**, select **apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Forwarding	Configures the system for forwarding (routing) or for dropping. The default value is forwarding. You cannot configure this option on a DvR Leaf node.
DefaultTTL	Configures the default time-to-live (TTL) value for a routed packet. TTL indicates the maximum number of seconds elapsed before a packet is discarded. Enter an integer from 1 to 255. The default value of 255 is used if a value is not supplied in the datagram header.
ReasmTimeout	Specifies the maximum number of seconds that received fragments are held while they wait for reassembly. The default value is 30 seconds.
ICMPUnreachableMsgEnable	Enables the generation of Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this system. These messages help determine if the system is reachable over the network. The default is disabled. Important: As a best practice, enable icmp-unreach-msg only if it is absolutely required. If icmp-unreach-msg is enabled and a packet is received for which there is no route in the routing table, CPU utilization can dramatically increase.
ICMPRedirectMsgEnable	Enables or disables the system sending ICMP destination redirect messages.
IcmpEchoBroadcastRequestEnable	Enables or disables IP ICMP echo broadcast request feature. The default is enabled.
IcmpDropFragmentsEnable	Enables or disables IPv4 Fragmented ICMP packet filtering globally. The default is disabled.

Name	Description
AlternativeEnable	Globally enables or disables the Alternative Route feature. If the alternative-route parameter is disabled, all existing alternative routes are removed. After the parameter is enabled, all alternative routes are re-added. The default is enabled.
RouteDiscoveryEnable	Enables the ICMP Router Discovery feature. The default is disabled (not selected). Use ICMP Router Discovery to enable hosts attached to multicast or broadcast networks to discover the IP addresses of neighboring routers.
AllowMoreSpecificNonLocalRouteEnable	Enables or disables a more-specific nonlocal route. If enabled, the system can enter a more-specific nonlocal route into the routing table. The default is disabled.
SuperNetEnable	Enables or disables supernetting. If supernetting is globally enabled, the system can learn routes with a route mask less than 8 bits. Routes with a mask length less than 8 bits cannot have ECMP paths, even if you globally enable the ECMP feature. The default is disabled.
UdpChecksumEnable	Enables or disables the UDP checksum calculation. The default is enable.
SourceRouteEnable	Enables or disables IP Source Routing globally. It is disabled by default.
SpbMulticastPolicyEnable	Enables the IP SPB Multicast Policy .
SpbMulticastPolicyRmap	Configures the name of the route-map for IP SPB Multicast Policy .
SpbMulticastPolicyApply	Applies the IP SPB Multicast Policy . The default value is noAction.
ARPLifeTime	Specifies the lifetime of an ARP entry within the system, global to the switch. The default value is 360 minutes.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled. After ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.

Name	Description
EcmpMaxPath	Globally configures the maximum number of ECMP paths. You cannot configure this feature unless ECMP is enabled globally.
Ecmp1PathList	Selects a preconfigured ECMP path.
Ecmp2PathList	Selects a preconfigured ECMP path.
Ecmp3PathList	Selects a preconfigured ECMP path.
Ecmp4PathList	Selects a preconfigured ECMP path.
Ecmp5PathList	Selects a preconfigured ECMP path.
Ecmp6PathList	Selects a preconfigured ECMP path.
Ecmp7PathList	Selects a preconfigured ECMP path.
Ecmp8PathList	Selects a preconfigured ECMP path.
EcmpPathListApply	Applies changes in the ECMP pathlist configuration, or in the prefix lists configured as the pathlists.
TcpAdjustMssEnable	Adjusts the TCP maximum segment size (MSS) to improve the throughput for the TCP session over a Fabric Extend (FE) adjacency. The default is disabled.
TcpAdjustMssStatus	Displays the activation status of the MSS adjustment functionality.
TcpAdjustMssType	Displays if the MSS adjustment value is manually configured or auto-derived. The software does not support auto-derived values for this feature.
TcpAdjustMssValue	Configures the MSS adjustment value. <ul style="list-style-type: none"> The default value is 1300.

Display Static Data I-SID IP Multicast Routes Information

About This Task

Perform this procedure to display the static data I-SID IP multicast routes information.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **StaticIsidIpMcastRoutes** tab.

StaticIsidIpMcastRoutes Field Descriptions

Use the data in the following table to use the **StaticIsidIpMcastRoutes** tab.

Name	Descriptions
IsidType	Indicates the type of static data I-SID.
Group	Indicates the group address of the IP multicast route that contains multicast routing information.
Source	Indicates the IP address where the IP multicast route originated.
VsnIsid	Indicates the Virtual Routing Network (VSN) I-SID.
SourceBeb	Indicates the name of the source Backbone Edge Bridge for IP multicast route.
VrfName	Indicates the VRF name of the interface.
Datalsid	Indicates the static I-SID value for IP multicast route.
Bvlan	Indicates the Backbone VLANs (B-VLANs) for the IP multicast route.
NniInterfaces	Specifies the network-to-network interface (NNI) ports for the IP multicast route.SPBM runs in the core on the ports that connect to the core. These ports are NNI ports. Ports that face a customer VLAN are user-to-network interface (UNI) ports.



IPsec

[IPsec Fundamentals](#) on page 2044

[IPsec configuration using CLI](#) on page 2050

[IPsec configuration using EDM](#) on page 2066

[IPsec configuration examples](#) on page 2082

Table 155: IPsec product support

Feature	Product	Release introduced
IPsec transport mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

IPsec Fundamentals

Internet Protocol Security (IPsec) ensures the authenticity, integrity, and confidentiality of data at the network layer of the Open System Interconnection (OSI) stack.

The IPsec feature is a set of security protocols and cryptographic algorithms that protect communication in a network. Use IPsec in scenarios where you need to encrypt packets between two hosts, or two routers, or a router and a host.

You can only configure the IPsec policies for IPv4 addresses for UDP, TCP, and ICMPv4 protocols.

IPsec adds support for OSPF virtual link for the security protection of the communication between the end points. You can also use IPsec with OSPFv3 on a

router port or VLAN interface, for example, if you want to encrypt OSPFv3 control traffic on a broadcast network. You can also use IPsec with ICMPv6.



Note

- If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.
- You can only configure the IPsec policies for IPv4 addresses for UDP, TCP, and ICMPv4 protocols. You can continue to configure IPsec policies for IPv6 addresses for ICMPv6, OSPFv3, TCP, and UDP.



Note

- When an OSPFv3 virtual link between two end points is secured using IPsec, the IPsec status on the IPv6 interfaces is automatically updated. This is applicable only on those interfaces that have no IPSEC policies manually configured on them

The following figure displays the movement of traffic using IPsec.




Figure 144: Internet Protocol Security (IPsec)

The IPsec feature uses security ciphers and encryption algorithms like AES, DES, and 3DES to ensure confidentiality of data, and keyed MAC for authenticity of data. The encryption algorithms require shared keys to secure the communication. The device only supports manual keying and configuration for IPsec. The IPsec feature supports IPv4 and IPv6 interfaces.

To configure IPsec, you create an IPsec policy, and then link the IPsec policy to an interface. You also link each IPsec policy to an IPsec security association. The IPsec

policies define the amount of security applied to specific traffic on a specific interface. The IPsec feature supports the following security protocols:

- Encapsulating security payload (ESP)
- Authentication header (AH)

The device restricts IPsec encryption to control traffic through the CPU. The IPsec feature processes either the ingress, the egress, or both the egress and ingress control packets to and from the CPU.

The device checks every ingress or egress packet for the IPsec base protocol, either AH or ESP. The base protocol interacts with the security policy database (SPD) and security association database (SADB) to check the level of security to apply to the packet. The device consults the SPD for both ingress and egress traffic. For egress traffic, the device consults the SPD to determine if IPsec needs to apply security considerations. For ingress traffic, the device consults the SPD to determine whether the traffic received with IPsec encapsulation complies with the policies defined in the system.

For more information on IPsec, see [IPsec Support with OSPFv3](#) on page 2771, [View IPsec Statistics](#) on page 2064, and [Display IPsec Interface Statistics](#) on page 2073.

Authentication header

The authentication header (AH) authenticates IP traffic and ensures you connect with who you want to connect. The authentication header can detect if data is altered in transit and protect against replay attacks. The authentication header does not encrypt traffic.

The authentication header provides a small header that precedes the payload with the use of the security parameters index (SPI) and sequence number. The authentication header provides:

- IP datagram sender authentication by HMAC or MAC
- IP datagram integrity assurance by HMAC or MAC
- Replay detection and protection by sequence number

The IPsec feature inserts the AH header after the IP header in transport mode. Transport mode with AH authenticates only the payload of the IP packet.

Tunnel mode authenticates the entire IP packet, including the IP header and data, to provide a secure hop between two hosts, two routers, or a router and a host.

You can apply AH alone, or in combination with the Encapsulating Security Payload (ESP).

The following figures show an original IP packet and an IP packet with an AH header.



Figure 145: Original IP packet




Figure 146: AH in transport mode

Encapsulating security payload

The encapsulating security payload (ESP) encrypts traffic with use of encryption algorithms, such as 3DES, AES-CBC, and AES-CTR. The security association specifies the algorithm and key used in ESP.

The encapsulating security payload can protect origin authenticity, integrity, and confidentiality of packets. ESP supports encryption-only and authentication-only configurations. The IPsec feature inserts the ESP header after the IP header and before the next layer protocol header in transport mode. Transport mode with ESP encrypts or authenticates only the payload of the IP packet.

Tunnel mode encrypts or authenticates the entire IP packet, including the IP header and data, to provide a secure hop between two hosts, two routers, or a router and a host.

The following figures display the original IP packet and an IP packet with ESP.



Figure 147: Original IP packet




Figure 148: ESP in transport mode

IPsec Modes

The IPsec feature security protocols use two different modes to protect the entire IP payload or the upper layer protocols:

- Transport mode
- Tunnel mode

Transport mode IPsec protects the upper layer protocols. In transport mode, IPsec adds an IPsec header between the IP header and upper layer protocol header.

Tunnel mode IPsec protects the whole IP packet. In tunnel mode, IPsec inserts the IPsec header between another IP datagram IP header and inner IP header.



Note

The IPsec implementation on the switch only supports transport mode.

Security association

A security association (SA) is a group of algorithms and parameters used to encrypt and authenticate the flow of IP traffic in a particular direction. An SA contains the information IPsec needs to process an IP packet. IPsec identifies SAs by:

- Security Parameter Index (SPI)
- Protocol value (either AH or ESP)
- Destination address to which the SA applies

Creation of a security association

Typically SAs exist in pairs; one in each direction, either inbound or outbound.

You can create SAs manually or dynamically. After you create an SA manually, the SA has no defined lifetime and the SA exists until you manually delete the SA.

After the device creates the SA dynamically, the SA can have a lifetime value that IPsec peers negotiate through use of a key management protocol. If the device uses the key excessively unauthorized access can occur. You must define the IPsec lifetime and other configurable parameters manually.

Security associations reside in the Security Association Database (SADB), which maintains a list of active SAs. The IPsec feature uses outbound SAs to secure the outgoing traffic and inbound SAs to process the incoming traffic. The device checks every ingress or egress packet for the IPsec base protocol, either AH or ESP. The base protocol interacts with the security policy database (SPD) and security association database (SADB) to check the level of security to apply to that packet.

The IPsec feature restricts SAs to the source and destination address of the connected router.

Security policy

Use IPsec to create IPsec security policies that define the levels of security for different types of traffic. You can use IPsec security policies to create rules to filter traffic with IPsec. IPsec policies determine what IP traffic to secure. An IPsec security policy typically consists of:

- An IP filter
- Security algorithms for authentication and key exchange
- An action

Creation of a security policy

You can configure IPsec on IPv4/IPv6 interfaces. First, create and configure an IPsec policy, and then add and enable the policy on an interface.

After you enable IPsec, the device encrypts all control traffic on the interface based on the policy. You have to specify individual policies to target a particular interface address or multiple addresses. By default, this implementation does not work on a subnet.

The Security Policy Database (SPD) maintains the IPsec security policies. The device checks every ingress or egress packet for the IPsec base protocol, either AH or ESP. The base protocol interacts with the security policy database (SPD) and security association database (SADB) to check the level of security to apply to that packet.

The IPsec feature only adds policies if the source address in the policy specified matches an interface IP address.

The IPsec feature restricts the policy match source address to the interface address of the router and destination IPv6 address.

IPsec Limitations

- The IPsec feature implementation is available only in software. Hardware implementation is not available. Only control packets to and from the CPU are subject to IPsec. IPsec implements IPsec policies in the software on the control path.
- The device does not support address ranges facility for an IPsec policy.
- No fast-path support exists for IPsec.

IPsec configuration using CLI

The following section provides procedures to configure Internet Protocol Security (IPsec).

Create an IPsec Policy

Use the following procedure to configure an IPsec policy. An IPsec policy defines the level of security for different types of traffic.



Note

- You can only configure the IPsec policies for IPv4 addresses for UDP, TCP, and ICMPv4 protocols. You can continue to configure IPsec policies for IPv6 addresses for ICMPv6, OSPFv3, TCP, and UDP.
- If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create an IPsec policy:

```
ipsec policy WORD<1-32>
```

Example

Create an IPsec policy named newpolicy:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec policy newpolicy
```

Variable Definitions

The following table defines parameters for the **ipsec policy** command.

Variable	Value
WORD<1-32>	Specifies the IPsec policy name.

Enable an IPsec Policy

Use the following procedure to enable an IPsec policy. An IPsec policy defines the level of security for different types of traffic.



Note

If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.

Before You Begin

- Create an IPsec policy.

About This Task

The IPsec feature adds policies only if the admin status of the policy and the IPsec status on the interface are enabled.

If you disable the IPsec policy on an IPv4 or IPv6 interface, IPsec removes the policy-related information from the security policy database (SPD) and the security association database (SADB), but the information remains on the system. After you re-enable, the information reapplies on the interface.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable an IPsec policy:

```
ipsec policy WORD<1-32> admin enable
```

Example

Enable an IPsec policy named newpolicy:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec policy newpolicy admin enable
```

Variable Definitions

The following table defines parameters for the **ipsec policy** command.

Variable	Value
<i>admin enable</i>	Enables the policy.
<i>WORD<1-32></i>	Specifies the IPsec policy name.

Create an IPsec Security Association

Use the following procedure to create an IPsec security association. A security association (SA) is a group of algorithms and parameters used to encrypt and

authenticate the flow of IP traffic in a particular direction. An SA contains the information IPsec needs to process an IP packet.

About This Task

You cannot delete or modify a security association if the security association links to a policy. To modify a parameter in the security association or to delete the security association, you must first unlink the security association from a policy.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create an IPsec security association:
`ipsec security-association WORD<1-32>`

Example

Create an IPsec security association named `newsa`:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec security-association newsa
```

Variable Definitions

The following table defines parameters for the **ipsec security-association** command.

Variable	Value
<code>WORD<1-32></code>	Specifies the security association identifier.

Configure an IPsec Security Association

Use the following procedure to configure an IPsec security association (SA). An SA is a group of algorithms and parameters used to encrypt and authenticate the flow of IP traffic in a particular direction. An SA contains the information IPsec needs to process an IP packet.

Before You Begin

- Create an IPsec security association to configure.

About This Task

You cannot delete or modify a security association if the security association links to a policy. To modify a parameter in the security association, or to delete the security association, you must first unlink the security association from a policy. You can only unlink a security association from a policy if the policy does not link to an interface. If a policy links to an interface, you must first unlink the policy from the interface, and then unlink the policy from the security association.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the IPsec security association key-mode:

```
ipsec security-association WORD<1-32> key-mode <automatic|manual>
```
3. Configure the IPsec security association mode:

```
ipsec security-association WORD<1-32> mode transport
```
4. Configure the IPsec security association encapsulation protocol:

```
ipsec security-association WORD<1-32> encap-proto <AH|ESP>
```
5. Configure the IPsec security association security parameters index:

```
ipsec security-association WORD<1-32> spi <1-4294967295>
```

For IPsec to function, each peer must have the same SPI value configured on both peers for a particular policy.

6. Configure the IPsec security association encryption algorithm:

```
ipsec security-association WORD<1-32> Encript-algo <3DES|AES-CBC|AES-CTR|NULL> [EncriptKey WORD<1-256>] [KeyLength <1-256>]
```

The encryption algorithm parameters are only accessible if you configure the encapsulation protocol to ESP.

7. Configure the IPsec security association authentication algorithm:

```
ipsec security-association WORD<1-32> auth-algo <AES-XCBC-MAC|MD5|SHA1|SHA2> [auth-key WORD<1-256>] [KeyLength <1-256>]
```
8. Configure the IPsec security association lifetime value:

```
ipsec security-association WORD<1-32> lifetime <Bytes<1-4294967295>|seconds<1-4294967295>>
```

Example

Configure an IPsec security association named `newsa` to have a key-mode of `ASCII`, an SA mode of `transport`, and an encapsulation protocol of `ESP`. Configure the encryption algorithm to `3DES`, with an encryption key of `11111111111111111111111111111111`, and a key length of `24`. Configure the authorization algorithm to `SHA1`, the authorization key to `11111111111111111111111111111111`, and key length to `20`. Configure the SPI to `1` and the lifetime in seconds to `1000`.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec security-association newsa mode transport
Switch:1(config)#ipsec security-association newsa encap-proto ESP
Switch:1(config)#ipsec security-association newsa Encript-algo 3DES Encript-key
11111111111111111111111111111111 KeyLength 48
Switch:1(config)#ipsec security-association newsa auth-algo SHA1 auth-key
11111111111111111111 KeyLength 20
Switch:1(config)#ipsec security-association newsa key-mode manual
Switch:1(config)#ipsec security-association newsa spi 1
Switch:1(config)#ipsec security-association newsa lifetime seconds 1000
```

Variable Definitions

The following table defines parameters for the **ipsec security-association** command.

Variable	Value
<i>WORD</i> <1-32>	Specifies the security association.
<i>auth-algo</i> <AES-XCBC-MAC MD5 SHA1 SHA2> [<i>auth-key</i> <i>WORD</i> <1-256>] [<i>KeyLength</i> <1-256>]	<p>Specifies the authorization algorithm, which includes one of the following values:</p> <ul style="list-style-type: none"> • AES-XCBC-MAC • MD5 • SHA1 • SHA2 <p>The default authentication algorithm name is MD5.</p> <p>The parameter <i>auth-key</i> specifies the authentication key.</p> <p>The <i>KeyLength</i> parameter specifies a string value of 1 to 256 characters in length. The default <i>KeyLength</i> is 128. The <i>KeyLength</i> values are as follows: 3DES is 48, AES-CBC is 32, 48, or 64, AES-CTR is 32.</p>
<i>encap-proto</i> <AH ESP>	<p>Specifies the encapsulation protocol:</p> <ul style="list-style-type: none"> • AH—Specifies authentication header. • ESP—Specifies encapsulation security payload. <p>If you configure the encapsulation protocol as AH, you cannot configure the encryption algorithms and other encryption related attributes. You can only access the encryption algorithm parameters if you configure the encapsulation protocol to ESP.</p> <p>The default value is ESP.</p>
<i>Encrpt-algo</i> <3DES AES-CBC AES-CTR NULL> [<i>EncrptKey</i> <i>WORD</i> <1-256>] [<i>KeyLength</i> <1-256>]	<p>Specifies the encryption algorithm value as one of the following:</p> <ul style="list-style-type: none"> • 3DES-CBC • AES-CBC • AES-CTR • NULL—Only use the NULL parameter to debug. Do not use this parameter in other circumstances. <p>The default encryption algorithm is AES-CBC.</p> <p>You can only access the encryption algorithm parameters if you configure the encapsulation protocol to ESP.</p> <p>The <i>EncrptKey</i> specifies the encryption key.</p> <p>The <i>KeyLength</i> specifies the key length value in a string from 1 to 256 characters. The default <i>KeyLength</i> is 128. The <i>KeyLength</i> values are as follows: 3DES is 48, AES-CBC is 32, 48, or 64, AES-CTR is 32.</p>

Variable	Value
<code>key-mode <automatic manual></code>	Specifies the key-mode as one of the following: <ul style="list-style-type: none"> • automatic • manual The default is manual.
<code>lifetime <Bytes<1-4294967295> seconds<1-4294967295></code>	Specifies the lifetime value in seconds or bytes. The default lifetime value in seconds is 28800. The default lifetime value in bytes is 4294966272.
<code>mode transport</code>	Specifies the mode as transport, which encapsulates the IP payload and provides a secure connection between two end points. Note: The IPsec implementation on the switch only supports transport mode.
<code>spi<1-4294967295></code>	Specifies the security parameters index (SPI) value, which is a unique value. SPI is a tag IPsec adds to the IP header. The tag enables the system that receives the IP packet to determine under which security association to process the received packet. For IPsec to function, each peer must have the same SPI value configured on both peers for a particular policy. The default value is 0.

Configure an IPsec Policy

Use the following procedure to configure an IPsec policy. An IPsec policy defines the level of security for different types of traffic.

Before You Begin

- Create an IPsec policy.

About This Task

You can only configure the IPsec policies for IPv4 addresses for UDP, TCP, and ICMPv4 protocols. You can continue to configure IPsec policies for IPv6 addresses for ICMPv6, OSPFv3, TCP, and UDP.

If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.

You cannot delete or modify a policy if the policy links to a security association, or if the policy links to a port or VLAN interface. If you need to modify a policy you must first unlink the policy from the security association, and the port or VLAN interface.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the remote address:

```
ipsec policy WORD<1-32> raddr WORD<1-32>
```

3. (Optional) Configure the local address:

```
ipsec policy WORD<1-32> laddr WORD<1-32>
```

The `laddr` parameter is an optional parameter that you can configure to have multiple local addresses for each remote address.

4. Configure the protocol:

```
ipsec policy WORD<1-32>[protocol <icmp|icmpv6|ospfv3|tcp|udp>] [sport  
<1-65535|any>][dport <1-65535|any>]
```

5. Configure the policy action:

```
ipsec policy WORD<1-32> [action <drop|permit>]
```

Example

Configure the remote address to 2001:db8:0:0:0:0:0:1 and local address to 2001:db8:0:0:0:0:0:15. configure the protocol to TCP source port 4 and destination port 5. Configure the policy to permit.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec policy Ipv6policy raddr 2001:db8:0:0:0:0:0:1
Switch:1(config)#ipsec policy Ipv6policy laddr 2001:db8:0:0:0:0:0:15
Switch:1(config)#ipsec policy Ipv6policy protocol tcp sport 4 dport 5
Switch:1(config)#ipsec policy Ipv6policy action permit
```

Configure the remote address to 192.0.1.1 and local address to 192.0.1.2. configure the protocol to TCP source port 4 and destination port 5. Configure the policy to drop.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec policy Ipv4policy raddr 192.0.1.1
Switch:1(config)#ipsec policy Ipv4policy laddr 192.0.1.2
Switch:1(config)#ipsec policy Ipv4policy protocol tcp sport 4 dport 5
Switch:1(config)#ipsec policy Ipv4policy action drop
```

Variable Definitions

The following table defines parameters for the **ipsec policy** command.

Variable	Value
<i>action</i> <drop permit>	Specifies the action the policy takes. The default is permit.
<i>laddr</i> WORD<1-32>	Specifies the local address. The <i>laddr</i> parameter is an optional parameter that you can configure to have multiple local addresses for each remote address. The default is 0::0.
<i>protocol</i> <icmp icmpv6 ospfv3 tcp udp>] [<i>sport</i> <1-65535> any>] [<i>dport</i> <1- 65535> any>]	Specifies the protocol, as one of the following: <ul style="list-style-type: none"> • ICMP • ICMPv6 • OSPFv3 • TCP • UDP <i>sport</i> — Specifies the source port for TCP and UDP. You can specify any to configure any port as the source port. <i>dport</i> — Specifies the destination port for TCP and UDP. You can specify any to configure any port as the destination port. The default protocol is TCP any. IPv4 only supports ICMP, UDP, and TCP.
<i>raddr</i> WORD<1-32>	Specifies the remote address. The default is 0::0.
WORD<1-32>	Specifies the policy name.

Link the IPsec Security Association to an IPsec Policy

Use the following procedure to link the security association to an IPsec policy.

Before You Begin

- The IPsec security association and IPsec policy must exist.

About This Task

You cannot delete or modify a security association if the security association links to a policy. To modify a parameter in the security association, or to delete the security association, you must first unlink the security association from the policy. You can only unlink a security association from a policy if the policy does not link to an interface. If a policy links to an interface, you must first unlink the policy from the interface, and then unlink the policy from the security association.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Link the IPsec security association to the IPsec policy:

```
ipsec policy WORD<1-32> security-association WORD<1-32>
```

Example

Link the IPsec security association named `new_sa` to the IPsec policy named `newpolicy`:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipsec policy newpolicy security-association newsa
```

Variable Definitions

The following table defines parameters for the **ipsec policy** command.

Variable	Value
<code>WORD<1-32></code>	Specifies the policy ID.
<code>security-association WORD<1-32></code>	Specifies the security association ID.

Enable IPsec on an Interface

Use the following procedure to enable IPsec on an interface. You can configure IPsec on a port, VLAN, or loopback interface.



Note

Management interface does not support IPsec configuration.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

followed by one of the following:

- `interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][, ...][slot/all][all]}`
- `interface loopback <1-256>`
- `interface vlan <1-4059>`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format `slot/port/sub-port`.

2. Enable IPsec on an IPv6 interface:

```
ipv6 ipsec enable
```

3. Enable IPsec on an IPv4 interface:

```
ip ipsec enable
```


Example

Enable IPsec for IPv6 on VLAN 100:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 100
Switch:1(config-if)#ipv6 ipsec enable
```

Variable Definition

The following table defines parameters for the **ip ipsec** and **ipv6 ipsec** commands.

Variable	Value
<i>enable</i>	Enables IPsec on the interface.

Link an IPsec Policy to an Interface

Use the following procedure to link an IPsec policy to an interface, and configure a policy direction. By default, the direction is both.

Before You Begin

- You must enable IPsec on the interface first, and then you link the IPsec policy to the interface.

About This Task

You cannot delete or modify an IPsec policy if the policy links to a port or VLAN interface. If you need to modify the policy, first unlink the policy from the port or VLAN interface.

Procedure

- Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

followed by one of the following:

- `interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][, ...][slot/all][all]}`
- `interface loopback <1-256>`
- `interface vlan <1-4059>`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- Link the IPsec policy to an IPv4 interface:

```
ip ipsec policy WORD<1-32> dir <both|in|out>
```

3. Link the IPsec policy to an IPv6 interface:

```
ipv6 ipsec policy WORD<1-32> dir <both|in|out>
```

Example

Link the IPsec policy newpolicy to the IPv6 interface VLAN 100:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 100
Switch:1(config-if)#ipv6 ipsec policy newpolicy dir both
```

Variable Definitions

The following table defines parameters for the **ip ipsec policy** and **ipv6 ipsec policy** commands.

Variable	Value
<i>WORD<1-32></i>	Specifies the policy ID.
<i>dir <both in out></i>	Specifies the direction you want to protect with IPsec: <ul style="list-style-type: none"> • both—Specifies both ingress and egress traffic. • in—Specifies ingress traffic. • out—Specifies egress traffic. The default is both.

Display IPsec Information on an Interface

Use the following procedure to display IPsec information on an interface.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the IPsec status on an Ethernet interface:

```
show ipsec interface gigabitethernet {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]
```
3. Display the IPsec status on a VLAN interface:

```
show ipsec interface vlan <1-4059>
```
4. Display the IPsec status on a loopback interface:

```
show ipsec interface loopback <1-256>
```

Example

Display the IPsec status on a VLAN interface:

```
Switch:1>show ipsec interface vlan 22
```

```
=====
                        VLAN Interface Policy Table
=====
```

Vlan Interface	Policy Name	IPsec State	Direction
22	AAA	Enable	both
22	tcp	Enable	both

22	icmp	Enable	both
----	------	--------	------

Variable Definitions

The following table defines parameters for the **show ipsec interface** command.

Variable	Value
<i>gigabitethernet</i> <i>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>loopback</i> <1-256>	Specifies the loopback interface.
<i>vlan</i> <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Display Configured IPsec Policies

Use the following procedure to display IPsec policies.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all of the IPsec policies on the switch:
show ipsec policy all
3. Display a specific IPsec policy based on the policy name on the interface:
show ipsec policy interface WORD<1-32>
4. Display the IPsec policy based on the policy name:
show ipsec policy name WORD<1-32>

Example

Display IPsec policy information:

```
Switch:1>show ipsec policy all
=====
                        IPSEC Policy Table
=====
PolicyName           : ospf1
LocalAddress: 0::0
RemoteAddress: 0::0
Protocol: ospfv3
```

```

src-port: 0
dest-port: 0
Action: Permit
Admin: Enable

Switch:1>show ipsec policy interface ospf1

=====
                        IPsec Policy Interface Table
=====
-----
POLICY NAME      InterfaceIndex      Policy State      Direction
-----
ospf1            2/3                    Enable            both

Switch:1>show ipsec policy name ospf1
=====
                        IPSEC Policy Table
=====
PolicyName       : ospf1
LocalAddress: 0::0
RemoteAddress: 0::0
Protocol: ospfv3
src-port: 0
dest-port: 0
Action: Permit
Admin: Enable

```

Variable Definitions

The following table defines parameters for the **show ipsec policy** command.

Variable	Value
<i>all</i>	Displays all of the IPsec policies on the switch.
<i>interface WORD<1-32></i>	Displays a specific IPsec policy based on the policy name on the interface.
<i>name WORD<1-32></i>	Displays the IPsec policy based on the name of the policy.

Display IPsec Security Association Information

Use the following procedure to display IPsec security association information.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display all IPsec security associations:
show ipsec sa all
3. Display a specific IPsec security association:
show ipsec sa name WORD<1-32>
4. Display all security associations linked to a specific policy:
show ipsec sa-policy

Example

Display information on IPsec security association policies:

```
Switch:1>enable
Switch:1#show ipsec sa all
=====
                        IPSEC Security Association Table
=====
sa-name: ospf1
key-Mode: manual
Encap protocol: ESP
SPI Value: 9
Encrypt Algorithm: 3dec-cbc
Encrypt-key: 52fb29f723b0800870dc83e3
Encrypt-key-Len: 24
Auth Algorithm: hmac-md5
Auth-key: 123456789abcdef0
Auth-key-Len: 16
Mode: transport
Lifetime-Sec: 1000
Lifetime-Byte: 20000

Switch:1#show ipsec sa name ospf1
=====
                        IPSEC Security Association Table
=====
sa-name: ospf1
key-Mode: manual
Encap protocol: ESP
SPI Value: 9
Encrypt Algorithm: 3dec-cbc
Encrypt-key: 52fb29f723b0800870dc83e3
Encrypt-key-Len: 24
Auth Algorithm: hmac-md5
Auth-key: 123456789abcdef0
Auth-key-Len: 16
Mode: transport
Lifetime-Sec: 1000
Lifetime-Byte: 20000

Switch:1#show ipsec sa-policy
=====
                        SA POLICY TABLE
=====
Policy Name      Security Association
-----
ospf1            ospf1
-----
```

Variable Definitions

The following table defines parameters for the **show ipsec sa** command.

Variable	Value
<i>all</i>	Displays all security associations.
<i>name WORD<1-32></i>	Displays a specific security association based on name.

Use the data in the following table to use the **show ipsec** command.

Variable	Value
<i>sa-policy</i>	Displays all security associations linked to a specific policy.

View IPsec Statistics

Use the following procedure to clear Internet Protocol Security (IPsec) system statistics counters and view IPsec statistics on an interface. The device only clears system statistics counters on system reboot.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View IPsec statistics for the system:
`show ipsec statistics system`
3. View IPsec statistics for an Ethernet interface:
`show ipsec statistics gigabitethernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}`
4. View IPsec statistics for a VLAN interface:
`show ipsec statistics vlan <1-4059>`
5. View statistics for IPsec on the loopback interface:
`show ipsec statistics loopback <1-256>`
6. Clear IPsec system statistics counters:
`clear ipsec stats all`

Example

View IPsec statistics. Output is partial due to length.

```
Switch:1>show ipsec statistics system

=====
                        IPSEC Global Statistics
=====
InSuccesses             = 0
InSPViolations          = 0
InNotEnoughMemories     = 0
InAHESPReplays          = 0
InAHFailures            = 0
InESPFailures           = 0
OutSuccesses            = 0
OutSPViolations         = 0
OutNotEnoughMemories    = 0
generalError            = 0
InAHSuccesses           = 0
InESPSuccesses          = 0
OutAHSuccesses          = 0
OutESPSuccesses         = 0
OutKBytes               = 0
OutBytes                = 0
InKBytes                = 0
InBytes                 = 0
--More-- (q = quit)
```

```
Switch:1>show ipsec statistics gigabitethernet 1/13
```

```
=====
                        Isec  Port  Stats
=====
Ifindex                = 204
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories    = 0
InAHESPReplays         = 0
InAHFailures           = 0
InESPFailures          = 0
OutSuccesses           = 0
OutSPViolations        = 0
OutNotEnoughMemories   = 0
generalError           = 0
```

```
Switch:1>show ipsec statistics vlan 1
```

```
=====
                        Isec  Vlan  Stats
=====
Ifindex                = 2049
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories    = 0
InAHESPReplays         = 0
InAHFailures           = 0
InESPFailures          = 0
OutSuccesses           = 0
OutSPViolations        = 0
OutNotEnoughMemories   = 0
generalError           = 0
```

View IPsec statistics for a loopback interface:

```
Switch:1>show ipsec statistics loopback 1
```

```
=====
                        Isec  LoopBack  Stats
=====
Ifindex                = 1344
InSuccesses            = 0
InSPViolations         = 0
InNotEnoughMemories    = 0
InAHESPReplays         = 0
InESPReplays           = 0
InAHFailures           = 0
InESPFailures          = 0
OutSuccesses           = 0
OutSPViolations        = 0
OutNotEnoughMemories   = 0
generalError           = 0
```

Variable Definitions

Use the data in the following table to use the **show ipsec statistics** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>loopback <1-256></code>	Identifies the loopback interface.
<code>system</code>	Shows statistics for the entire system.
<code>vlan <1-4059></code>	Specifies the VLAN.

IPsec configuration using EDM

The following section provides procedures to configure Internet Protocol security (IPsec).

Create an IPsec Policy

Use the following procedure to configure an IPsec policy for an IPv4 or an IPv6 interface. An IPsec policy defines the level of security for different types of traffic.



Note

- You can only configure the IPsec policies for IPv4 addresses for UDP, TCP, and ICMPv4 protocols. You can continue to configure IPsec policies for IPv6 addresses for ICMPv6, OSPFv3, TCP, and UDP.
- If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.

About This Task

You cannot delete or modify a policy if the policy links to a security association, or if the policy links to a port or VLAN interface. If you need to modify a policy you must first unlink the policy from the security association, and the port or VLAN interface.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IPSec**.
3. Select the **Policy** tab.
4. Select **Insert**.
5. In **Name**, type a policy name.
6. Complete the remaining optional configuration to customize the policy.

7. Select **Insert**.**Policy** Field Descriptions

Use the data in the following table to use the **Policy** tab.

Name	Description
Name	Specifies the IPsec policy name.
DstAddress	Specifies the remote address. This field accepts IPv4 and IPv6 address, depending on the selected source address type.
SrcAddress	Specifies the local address. The local address is optional that you can configure to have multiple local addresses for each remote (destination) address. This field accepts IPv4 and IPv6 address, depending on the selected source address type.
SrcPort	Specifies the source port for TCP and UDP. Leave this field empty to configure any port as the source port. The default is value is 1.
DstPort	Specifies the destination port for TCP and UDP. Leave this field empty to configure any port as the destination port. The default value is 1.
AdminFlag	Enables or disables the policy. The default is disabled.
L4Protocol	Specifies the protocol, as one of the following: <ul style="list-style-type: none"> • tcp • udp • icmp • icmpv6 • ospfv3 IPv4 interfaces only support TCP, UDP, and ICMP. The default is TCP.
Action	Specifies the action the policy takes. The default is to permit the packet.

Create an IPsec Security Association

Use the following procedure to create an IPsec security association. A security association (SA) is a group of algorithms and parameters used to encrypt and authenticate the flow of IP traffic in a particular direction. An SA contains the information IPsec needs to process an IP packet.

About This Task

You cannot delete or modify a security association if the security association links to a policy. To modify a parameter in the security association or to delete the security association, you must first unlink the security association from a policy.

You can only unlink a security association from a policy if the policy does not link to an interface. If a policy links to an interface, you must first unlink the policy from the interface, and then unlink the policy from the security association.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IPSec**.
3. Select the **Security Association** tab.
4. Select **Insert**.
5. In **Name**, type a name to identify the SA.
6. In **SPI**, type the security parameters index.



Note

For IPsec to function, each peer must have the same SPI value configured for a particular policy.

7. Complete the remaining optional configuration.
8. Select **Insert**.

Security Association Field Descriptions

Use the data in the following table to use the **Security Association** tab.

Name	Description
Name	Specifies the name of the security association.
Spi	Specifies the security parameters index (SPI) value, which is a unique value. SPI is a tag IPsec adds to the IP header. The tag enables the system that receives the IP packet to determine under which security association to process the received packet. For IPsec to function, each peer must have the same SPI value configured for a particular policy. The default value is 0.

Name	Description
HashAlgorithm	<p>Specifies the authorization algorithm, which includes one of the following values:</p> <ul style="list-style-type: none"> • AESXCBC • MD5 • SHA1 • SHA2 <p>The default authentication algorithm name is MD5.</p>
EncryptAlgorithm	<p>Specifies the encryption algorithm value as one of the following:</p> <ul style="list-style-type: none"> • DES3CBC • AES128CBC • AESCTR • NULL—Only use the NULL parameter to debug. Do not use this parameter in any other circumstance. <p>The default encryption algorithm is AES128CBC. You can only access the encryption algorithm parameters if you configure the encapsulation protocol to ESP.</p>
AuthMethod	<p>Specifies the encapsulation protocol:</p> <ul style="list-style-type: none"> • ah—Specifies authentication header. • es—Specifies encapsulation security payload. <p>If you configure the encapsulation protocol as ah, you cannot configure the encryption algorithms and other encryption related attributes. You can only access the encryption algorithm parameters if you configure the encapsulation protocol to es.</p> <p>The default value is es.</p>
Mode	<p>Specifies the mode as transport, which encapsulates the IP payload and provides a secure connection between two end points.</p> <p>Note: The IPsec implementation on the switch only supports transport mode.</p>
KeyMode	<p>Specifies the key-mode as one of the following:</p> <ul style="list-style-type: none"> • manual • auto <p>The default is manual.</p>

Name	Description
EncryptKeyName	Specifies the encryption key.
EncryptKeyLength	Specifies the numbers of bits used in the encryption key. The key length values are as follows: <ul style="list-style-type: none">• DES3CBC is 48• AES128CBC is 32, 48, 64• AESCTR is 32
HashKeyName	Specifies the authentication key.
HashKeyLength	Specifies the numbers of bits used in the hash key. The key length values are as follows: <ul style="list-style-type: none">• AESXCBC is 32• MD5 is 32• SHA1 is 40
LifetimeSeconds	Specifies the lifetime value in seconds. The lifetime determines the traffic that can pass between IPsec peers using a security association before that security association expires. The default lifetime value in seconds is 28800.
LifetimeKbytes	Specifies the lifetime value in kilobytes. The lifetime determines the traffic that can pass between IPsec peers using a security association before that security association expires. The default lifetime value in bytes is 4294967295.

Link the IPsec Security Association to an IPsec Policy

Use the following procedure to link the security association to an IPsec policy.

Before You Begin

- The IPsec security association and IPsec policy must exist.

About This Task

You cannot delete or modify a security association if the security association links to a policy. To modify a parameter in the security association, or to delete the security association, you must first unlink the security association from the policy. You can only unlink a security association from a policy if the policy does not link to an interface. If a policy links to an interface, you must first unlink the policy from the interface, and then unlink the policy from the security association.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.

2. Select **IPSec**.
3. Select the **Policy SA Link** tab.
4. Select **Insert**.
5. In **PolicyName**, type the IPsec policy name.
6. In **SAName**, type the security association name.
7. Select **Insert**.

Policy SA Link Field Descriptions

Use the data in the following table to use the **Policy SA Link** tab.

Name	Description
PolicyName	Specifies the name of the IPsec policy.
SAName	Specifies the name of the security association.

Enable IPsec on an IPv6 Interface

Use the following procedure to enable IPsec on an IPv6 interface.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IPSec**.
3. Select the **IPv6 Interfaces** tab.
4. In the IpsecEnable column, double-click in **IpsecEnable**, and select **enable**.
5. Select **Apply**.

IPv6 Interfaces Field Descriptions

Use the data in the following table to use the **IPv6 Interfaces** tab.

Name	Description
Interface	Specifies the interface.
IpsecEnable	Enables IPsec on the interface. The default is disabled.

Enable IPsec on an IPv4 Interface

Use the following procedure to enable IPsec on an IPv4 interface.



Note

If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.

2. Select **IPSec**.
3. Select the **IPv4 Interfaces** tab.
4. In the IpsecEnable column, double-click in **IpsecEnable**, and select **enable**.
5. Select **Apply**.

IPv4 Interfaces Field Descriptions

Use the data in the following table to use the **IPv4 Interfaces** tab.

Name	Description
Interface	Specifies the interface.
IpsecEnable	Enables IPsec on the interface. The default is disabled.

Link an IPsec Policy to an Interface

Use the following procedure to link an IPsec policy to an interface, and configure a policy direction. By default, the direction is both.

Before You Begin

- You must enable IPsec on the interface first, and then you link the IPsec policy to the interface.

About This Task

You cannot delete or modify an IPsec policy if the policy links to a port or VLAN interface. If you need to modify the policy, first unlink the policy from the port or VLAN interface.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IPSec**.
3. Select the **Interface Policy** tab.
4. Select **Insert**.
5. In **Name**, type the name of the IPsec policy.
6. In **IfIndex**, select either **Port** or **Vlan**, and then select an interface.
7. Select **Okay**.
8. Complete the remaining optional configuration.
9. Select **Insert**.

Interface Policy Field Descriptions

Use the data in the following table to use the **Interface Policy** tab.

Name	Description
Name	Specifies the IPsec policy name.
IfIndex	Links a policy to an interface.

Name	Description
IfEnabled	Shows if the IPsec is enabled on the interface and if the administrative state of the policy is enabled.
IfDirection	Specifies the direction you want to protect with IPsec: <ul style="list-style-type: none"> inbound—Specifies ingress traffic. outbound—Specifies egress traffic. bothDirections—Specifies both ingress and egress traffic. The default is bothDirections.

Display IPsec Interface Statistics

Use this procedure to view IPsec statistics and counter values for each IPsec-enabled interface.

About This Task

If you select an interface on the **Stats** tab, you can select **Graph** to graph particular statistics for that interface.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IPSec**.
3. Select the **Interface Stats** tab.

Interface Stats Field Descriptions

Use the data in the following table to use the **Interface Stats** tab.

Name	Description
IfIndex	Shows the interface index for which the statistic is captured.
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the AH replay check fails.

Name	Description
InESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the ESP replay check fails.
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.
InAhSuccesses	Specifies the number of ingress packets IPsec carries because the AH authentication succeeds.
OutAHSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
InESPSuccesses	Specifies the number of ingress packets IPsec carries since boot time because the ESP authentication succeeds.
OutESPSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutKBytes	Specifies the total number of kilobytes on egress.
OutBytes	Specifies the total number of bytes on egress.
InKBytes	Specifies the total number of bytes on ingress.
InBytes	Specifies the total number of bytes on ingress.
TotalPacketsProcessed	Specifies the total number of packets processed.
TotalPacketsByPassed	Specifies the total number of packets bypassed.
OutAHFailures	Specifies the number of egress packets IPsec discards since boot time because the AH authentication check fails.

Name	Description
OutESPFailures	Specifies the number of egress packets IPsec discards since boot time because the ESP authentication check fails.
InMD5Hmacs	Specifies the number of inbound HMAC MD5 occurrences since boot time.
InSHA1Hmacs	Specifies the number of inbound HMAC SHA1 occurrences since boot time.
InAESXCBCs	Specifies the number of inbound AES XCBC MAC occurrences since boot time.
InAnyNullAuth	Specifies the number of inbound null authentication occurrences since boot time.
In3DESCBCs	Specifies the number of inbound 3DES CBC occurrences since boot time.
InAESCBCs	Specifies the number of inbound AES CBC occurrences since boot time.
InAESCTRs	Specifies the number of inbound AES CTR occurrences since boot time.
InAnyNullEncrypt	Specifies the number of inbound null occurrences since boot time. Used for debugging purposes.
OutMD5Hmacs	Specifies the number of outbound HMAC MD5 occurrences since boot time.
OutSHA1Hmacs	Specifies the number of outbound HMAC SHA1 occurrences since boot time.
OutAESXCBCs	Specifies the number of outbound AES XCBC MAC occurrences since boot time.
OutInAnyNullAuth	Specifies the number of outbound null authentication occurrences since boot time.
Out3DESCBCs	Specifies the number of outbound 3DES CBC occurrences since boot time.
OutAESCBCs	Specifies the number of outbound AES CBC occurrences since boot time.
OutAESCTRs	Specifies the number of outbound AES CTR occurrences since boot time.
OutInAnyNullEncrypt	Specifies the number of outbound null occurrences since boot time. Used for debugging purposes.

Graph IPsec Interface Statistics

Use this procedure to graphically view IPsec statistics and counter values for each IPsec-enabled interface.

About This Task

If you select an interface on the **Stats** tab, you can select **Graph** to graph particular statistics for that interface.

Procedure

1. In the navigation pane, expand **Security > Control Path**.
2. Select **IPSec**.
3. Select the **Interface Stats** tab.
4. Select a row, and Select **Graph**.
5. Select one of the parameters, and Select the appropriate icon in the upper-left corner of the menu bar to draw a line chart, area chart, bar chart, or a pie chart.
6. To clear existing counters, and fix a reference point in time to restart the counters, Select **Clear Contents**.
7. To export the statistical data to a file, select **Export**.
8. To configure a poll interval, select an appropriate value from the **Poll Interval** drop-down list.

Interface Stats - Graph Field Descriptions

Use the data in the following table to use the **Interface Stats - Graph** tab.

Name	Description
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the AH replay check fails.
InESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the ESP replay check fails.
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.

Name	Description
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.
InAhSuccesses	Specifies the number of ingress packets IPsec carries because the AH authentication succeeds.
OutAHSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
InESPSuccesses	Specifies the number of ingress packets IPsec carries since boot time because the ESP authentication succeeds.
OutESPSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutKBytes	Specifies the total number of kilobytes on egress.
OutBytes	Specifies the total number of bytes on egress.
InKBytes	Specifies the total number of bytes on ingress.
InBytes	Specifies the total number of bytes on ingress.
TotalPacketsProcessed	Specifies the total number of packets processed.
TotalPacketsByPassed	Specifies the total number of packets bypassed.
OutAHFailures	Specifies the number of egress packets IPsec discards since boot time because the AH authentication check fails.
OutESPFailures	Specifies the number of egress packets IPsec discards since boot time because the ESP authentication check fails.
InMD5Hmacs	Specifies the number of inbound HMAC MD5 occurrences since boot time.
InSHA1Hmacs	Specifies the number of inbound HMAC SHA1 occurrences since boot time.
InAESXCBCs	Specifies the number of inbound AES XCBC MAC occurrences since boot time.
InAnyNullAuth	Specifies the number of inbound null authentication occurrences since boot time.
In3DESCBCs	Specifies the number of inbound 3DES CBC occurrences since boot time.

Name	Description
InAESCBCs	Specifies the number of inbound AES CBC occurrences since boot time.
InAESCTRs	Specifies the number of inbound AES CTR occurrences since boot time.
InAnyNullEncrypt	Specifies the number of inbound null occurrences since boot time. Used for debugging purposes.
OutMD5Hmacs	Specifies the number of outbound HMAC MD5 occurrences since boot time.
OutSHA1Hmacs	Specifies the number of outbound HMAC SHA1 occurrences since boot time.
OutAESXCBCs	Specifies the number of outbound AES XCBC MAC occurrences since boot time.
OutInAnyNullAuth	Specifies the number of outbound null authentication occurrences since boot time.
Out3DESCBCs	Specifies the number of outbound 3DES CBC occurrences since boot time.
OutAESCBCs	Specifies the number of outbound AES CBC occurrences since boot time.
OutAESCTRs	Specifies the number of outbound AES CTR occurrences since boot time.
OutInAnyNullEncrypt	Specifies the number of outbound null occurrences since boot time. Used for debugging purposes.

Display Switch Level Statistics for IPsec-Enabled Interfaces

Use this procedure to view IPsec statistics and counter values at the switch level for all IPsec-enabled interfaces.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **IPSec**.
3. Select the **Global Stats** tab.

Global Stats Field Descriptions

Use the data in the following table to use the **Global Stats** tab.

Name	Description
InSuccesses	Specifies the number of ingress packets IPsec successfully carries.
InSPViolations	Specifies the number of ingress packets IPsec discards since boot time because of a security policy violation.
InNotEnoughMemories	Specifies the number of ingress packets IPsec discards since boot time because not enough memory is available.
InAHESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the AH replay check fails.
InESPReplays	Specifies the number of ingress packets IPsec discards since boot time because the ESP replay check fails.
InAHFailures	Specifies the number of ingress packets IPsec discards since boot time because the AH authentication check fails.
InESPFailures	Specifies the number of ingress packets IPsec discards since boot time because the ESP authentication check fails.
OutSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutSPViolations	Specifies the number of egress packets IPsec discards since boot time because a security policy violation occurs.
OutNotEnoughMemories	Specifies the number of egress packets IPsec discards since boot time because not enough memory is available since boot time.
generalError	Specifies a general error.
InAHSuccesses	Specifies the number of ingress packets IPsec carries because the AH authentication succeeds.
OutAHSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
InESPSuccesses	Specifies the number of ingress packets IPsec carries since boot time because the ESP authentication succeeds.
OutESPSuccesses	Specifies the number of egress packets IPsec successfully carries since boot time.
OutKBytes	Specifies the total number of kilobytes on egress.
OutBytes	Specifies the total number of bytes on egress.
InKBytes	Specifies the total number of bytes on ingress.

Name	Description
InBytes	Specifies the total number of bytes on ingress.
TotalPacketsProcessed	Specifies the total number of packets processed.
TotalPacketsByPassed	Specifies the total number of packets bypassed.
OutAHFailures	Specifies the number of egress packets IPsec discards since boot time because the AH authentication check fails.
OutESPFailures	Specifies the number of egress packets IPsec discards since boot time because the ESP authentication check fails.
InMD5Hmac	Specifies the number of inbound HMAC MD5 occurrences since boot time.
InSHA1Hmac	Specifies the number of inbound HMAC SHA1 occurrences since boot time.
InAESXCBCs	Specifies the number of inbound AES XCBC MAC occurrences since boot time.
InAnyNullAuth	Specifies the number of inbound null authentication occurrences since boot time.
In3DESCBCs	Specifies the number of inbound 3DES CBC occurrences since boot time.
InAESCBCs	Specifies the number of inbound AES CBC occurrences since boot time.
InAESCTRs	Specifies the number of inbound AES CTR occurrences since boot time.
InAnyNullEncrypt	Specifies the number of inbound null occurrences since boot time. Used for debugging purposes.
OutMD5Hmac	Specifies the number of outbound HMAC MD5 occurrences since boot time.
OutSHA1Hmac	Specifies the number of outbound HMAC SHA1 occurrences since boot time.
OutAESXCBCs	Specifies the number of outbound AES XCBC MAC occurrences since boot time.
OutInAnyNullAuth	Specifies the number of outbound null authentication occurrences since boot time.
Out3DESCBCs	Specifies the number of outbound 3DES CBC occurrences since boot time.
OutAESCBCs	Specifies the number of outbound AES CBC occurrences since boot time.

Name	Description
OutAESCTRs	Specifies the number of outbound AES CTR occurrences since boot time.
OutInAnyNullEncrypt	Specifies the number of outbound null occurrences since boot time. Used for debugging purposes.

Configure IPsec for the OSPF Virtual Link

Use the following procedure to configure and enable IPsec for the OSPF virtual link.

IPsec is disabled by default.

Before You Begin

- Configure the OSPF virtual link.
- Create the IPsec security association.

About This Task

Until you enable IPsec on both sides of the virtual links, the links cannot exchange OSPFv3 control messages, and the system drops OSPFv3 exchange packets.

You must disable IPsec before you can perform virtual link policy configuration changes.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Panel**.
2. Select **IPSec**.
3. Select the **OSPF Virtual Link** tab.
4. Select **Insert**.
5. Specify the area ID.
6. Specify the neighbor address.
7. Complete the remaining optional configuration.
8. Select **Insert**.

OSPF Virtual Link Field Descriptions

Use the data in the following table to use the **OSPF Virtual Link** tab.

Name	Description
Areaid	Identifies the OSPF virtual link area.
Neighbor	Identifies the OSPF virtual link neighbor.
SAName	Links the security association to the OSPF virtual link.
AdminStatus	Enables the policy. The default is disabled.

Name	Description
Action	Configures the action of the IPsec policy under the OSPF virtual tunnel to one of the following: <ul style="list-style-type: none">• permit—Permits the IP packets.• drop—Drops the IP packets. The default is permit.
Direction	Specifies the direction you want to protect with IPsec: <ul style="list-style-type: none">• inBound—Specifies ingress traffic.• outBound—Specifies egress traffic.• bothDirections—Specifies both ingress and egress traffic. The default is bothDirections.
SrcAddress	Shows the address of the source interface to which the policy applies.
DstAddress	Shows the address of the destination interface to which the policy applies.
LinkID	Shows a unique ID for the OSPF virtual link. The default is 0.
IfIndex	Shows the interface index to which OSPF virtual link the policy applies.
OperStatus	Shows the operational status of the link, either up or down. The default is down.

IPsec configuration examples

The following section provides examples to configure Internet Protocol Security (IPsec).



Note

If you downgrade your software, the current IPsec configurations are no longer supported. You must boot with the factory default settings for IPsec, and then reconfigure the IPsec features.

IPsec configuration example

Review the following information to understand IPsec configuration.

Use the following steps to configure IPsec.

1. Create and configure an IPsec policy.
2. Enable the policy.
3. Create an IPsec security association to correspond with the IPsec policy.
4. Configure the key mode format.
5. Configure the security association.
6. Link the IPsec security association to the IPsec policy.

7. Enable the IPsec policy on the interface.
8. Link the IPsec policy with the interface.
9. Enable the IPsec on the interface that links to the IPsec policy.

For an example configuration and for more information on IPsec OSPFv3 and OSPFv3 virtual link, see [OSPF](#) on page 2748.

Create a policy named `newpolicy` with a security association named `new_sa` on VLAN 100.

The following displays the IPsec policy configuration:

```
ipsec policy newpolicy raddr 2001:db8:0:0:0:0:0:1
ipsec policy newpolicy laddr 2001:db8:0:0:0:0:0:15
ipsec policy newpolicy protocol tcp sport 4 dport 5
ipsec policy newpolicy action permit
```

The following example displays the IPsec security association:

```
ipsec security-association new_sa
ipsec security-association new_sa key-mode manual
ipsec security-association new_sa mode transport
ipsec security-association new_sa encap-proto ESP
ipsec security-association new_sa Encrpt-algo 3DES-CBC encrypt-key
11111111111111111111111111111111 KeyLength 24
ipsec security-association new_sa auth-algo SHA1 auth-key 11111111111111111111111111111111 KeyLength
20
ipsec security-association new_sa spi 1
ipsec security-association new_sa lifetime seconds 1000
```

IPsec with ICMPv6 configuration example

The following displays configuration of IPsec with ICMPv6.




Figure 149: IPsec configuration with ICMPv6

Switch 10 security association configuration

The following example displays the configuration of the security association on Switch 10.

```
ipsec security-association icmp
ipsec security-association icmp encap-proto ESP
ipsec security-association icmp mode transport
```

```
ipsec security-association icmp spi 1
ipsec security-association icmp auth-algo SHA1 auth-key
1234567890123456789012345678901234567890 keyLength 40
ipsec security-association icmp Encrpt-algo AES-CBC EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association icmp key-mode manual
ipsec security-association icmp lifetime seconds 1
ipsec security-association icmp lifetime bytes 1
```

Switch 10 policy configuration

The following example displays the configuration of the security policy on Switch 10.

```
ipsec policy ICMP_Policy
ipsec policy ICMP_Policy admin enable
ipsec policy ICMP_Policy raddr 2001::2
ipsec policy ICMP_Policy laddr 2001::1
ipsec policy ICMP_Policy protocol icmpv6
ipsec policy ICMP_Policy action permit
ipsec policy ICMP_Policy security-association icmp
```

Switch 10 interface configuration

The following example displays the configuration of IPsec on slot/port 1/10.

```
interface gigabitEthernet 1/10
no shut
interface vlan 3
interface address 2000::1
interface enable
ipv6 ipsec policy ICMP_Policy dir both
ipv6 ipsec enable
```

Switch 10 VLAN configuration

The following example displays the creation and configuration of VLAN 3 with IPsec.

```
interface gigabitEthernet 1/10
no shut
exit
vlan create 3 type port-mstprstp 3
vlan members add 3 1/10 portmember
interface vlan 3
interface enable
interface address 2000::1
ipv6 ipsec policy ICMP_Policy dir both
ipv6 ipsec enable
```

Switch 30 security association configuration

The following example displays the configuration of the security association on Switch 30.

```
ipsec security-association icmp
ipsec security-association icmp encap-proto ESP
ipsec security-association icmp mode transport
ipsec security-association icmp spi 1
ipsec security-association icmp auth-algo SHA1 auth-key
1234567890123456789012345678901234567890 keyLength 40
ipsec security-association icmp Encrpt-algo AES-CBC EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association icmp key-mode manual
```

```
ipsec security-association icmp lifetime seconds 1
ipsec security-association icmp lifetime bytes 1
```

Switch 30 policy configuration

The following example displays the configuration of the security policy on Switch 30.

```
ipsec policy ICMP_Policy
ipsec policy ICMP_Policy admin enable
ipsec policy ICMP_Policy raddr 2001::1
ipsec policy ICMP_Policy laddr 2001::2
ipsec policy ICMP_Policy action permit
ipsec policy ICMP_Policy protocol icmpv6
ipsec policy ICMP_Policy security-association icmp
```

Switch 30 interface configuration

The following example displays the configuration of IPsec on slot/port 1/10.

```
interface gigabitEthernet 1/10
no shut
ipv6 interface enable
ipv6 interface vlan 3
ipv6 interface address 2001::2
ipv6 ipsec policy ICMP_Policy dir both
ipv6 ipsec enable
```


Switch 30 VLAN configuration

The following example displays the creation and configuration of VLAN 3 with IPsec.

```
interface gigabitEthernet 1/10
no shut
exit
vlan create 3 type port-mstprstp 0
vlan members add 3 1/20
interface vlan 3
ipv6 interface enable
ipv6 interface address 2001::2
ipv6 ipsec policy ICMP_Policy dir both
ipv6 ipsec enable
```

OSPFv3 IPsec configuration example

The following example displays a network using IPsec used with OSPFv3.



The following example displays the configuration of IPsec with OSPFv3. For OSPFv3 conceptual and procedural information, see [OSPF](#) on page 2748.

Switch 10 security associations

The following example displays the configuration of security associations for OSPFv3 for Switch 10.

```
ipsec security-association ospf1
ipsec security-association ospf1 encap-proto ESP
ipsec security-association ospf1 mode transport
ipsec security-association ospf1 spi 1
ipsec security-association ospf1 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf1 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf1 key-mode manual
ipsec security-association ospf1 lifetime seconds 1
ipsec security-association ospf1 lifetime bytes 1

ipsec security-association ospf2
ipsec security-association ospf2 encap-proto ESP
ipsec security-association ospf2 mode transport
ipsec security-association ospf2 spi 2
ipsec security-association ospf2 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf2 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf2 key-mode manual
ipsec security-association ospf2 lifetime seconds 1
ipsec security-association ospf2 lifetime bytes 1

ipsec security-association ospf3
ipsec security-association ospf3 encap-proto ESP
ipsec security-association ospf3 mode transport
ipsec security-association ospf3 spi 3
ipsec security-association ospf3 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf3 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf3 key-mode manual
ipsec security-association ospf3 lifetime seconds 1
ipsec security-association ospf3 lifetime bytes 1

ipsec security-association ospf4
ipsec security-association ospf4 encap-proto ESP
ipsec security-association ospf4 mode transport
ipsec security-association ospf4 spi 4
ipsec security-association ospf4 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf4 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf4 key-mode manual
ipsec security-association ospf4 lifetime seconds 1
ipsec security-association ospf4 lifetime bytes 1

ipsec security-association ospf5
ipsec security-association ospf5 encap-proto ESP
ipsec security-association ospf5 mode transport
ipsec security-association ospf5 spi 5
ipsec security-association ospf5 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf5 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
```

```

ipsec security-association ospf5 key-mode manual
ipsec security-association ospf5 lifetime seconds 1
ipsec security-association ospf5 lifetime bytes 1

ipsec security-association ospf6
ipsec security-association ospf6 encap-proto ESP
ipsec security-association ospf6 mode transport
ipsec security-association ospf6 spi 6
ipsec security-association ospf6 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf6 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf6 key-mode manual
ipsec security-association ospf6 lifetime seconds 1
ipsec security-association ospf6 lifetime bytes 1

```

Switch 10 policy configuration

The following example displays the configuration of policies on Switch 10. The link local address is fe80:0:0:0:b2ad:aaff:fe43:100 and the remote link local address is fe80:0:0:0:b2ad:aaff:fe43:4d00. The following displays the policy with the laddr configured to the link local address and raddr configured to the remote link local address, with the direction configured as outbound.

```

ipsec policy ospf1
ipsec policy ospf1 admin enable
ipsec policy ospf1 raddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf1 laddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf1 protocol ospfv3
ipsec policy ospf1 action permit

```

The following example displays the configuration of policies on Switch 10. The link local address is fe80:0:0:0:b2ad:aaff:fe43:100 and the remote link local address is fe80:0:0:0:b2ad:aaff:fe43:4d00. The following displays the policy with the laddr configured to the link local address and raddr configured to the remote link local address, with the direction configured as inbound.

For a policy direction of inbound, laddr and raddr are reversed before storing to the stack. Because of this, even though the policy requires you to configure the laddr as the remote link local address, you need to configure laddr as the link local address in the configuration.

```

ipsec policy ospf2
ipsec policy ospf2 admin enable
ipsec policy ospf2 raddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf2 laddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf2 protocol ospfv3
ipsec policy ospf2 action permit

```

Laddr is configured to the link local and raddr is configured to ff02::05 with the direction configured as outbound.

```

ipsec policy ospf3
ipsec policy ospf3 admin enable
ipsec policy ospf3 raddr ff02::05
ipsec policy ospf3 laddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf3 protocol ospfv3
ipsec policy ospf3 action permit

```

Laddr is configured to the remote link local and raddr is configured to ff02::05 with the direction configured as inbound.

```
ipsec policy ospf4
ipsec policy ospf4 admin enable
ipsec policy ospf4 raddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf4 laddr ff02::05
ipsec policy ospf4 protocol ospfv3
ipsec policy ospf4 action permit
```

Laddr is configured to the link local and raddr is configured to ff02::06 with the direction as outbound.

```
ipsec policy ospf5
ipsec policy ospf5 admin enable
ipsec policy ospf5 raddr ff02::06
ipsec policy ospf5 fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf5 protocol ospfv3
ipsec policy ospf5 action permit
```

Laddr is configured to the remote link local and raddr is configured to ff02::06 with the direction configured as inbound.

```
ipsec policy ospf6
ipsec policy ospf6 admin enable
ipsec policy ospf6 raddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf6 laddr ff02::06
ipsec policy ospf6 protocol ospfv3
ipsec policy ospf6 action permit
```

Switch 10 link table configuration

The following example displays the linking of the policy with the security association on Switch 10.

```
ipsec policy ospf1 security-association ospf1
ipsec policy ospf2 security-association ospf2
ipsec policy ospf3 security-association ospf3
ipsec policy ospf4 security-association ospf4
ipsec policy ospf5 security-association ospf5
ipsec policy ospf6 security-association ospf6
```

Switch 10 OSPFv3 configuration

The following example displays the OSPFv3 configuration on Switch 10.

```
router ospf ipv6-enable
router ospf
ipv6 router-id 1.1.1.1
ipv6 area 0.0.0.1
```

Switch 10 interface configuration

The following example displays the interface configuration on slot/port 1/10.

```
interface gigabitEthernet 1/10
no shut
ipv6 interface vlan 3
ipv6 interface address 2000::1/64
ipv6 interface enable
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
ipv6 ipsec policy ospf1 dir out
```

```

ipv6 ipsec policy ospf2 dir in
ipv6 ipsec policy ospf3 dir out
ipv6 ipsec policy ospf4 dir in
ipv6 ipsec policy ospf5 dir out
ipv6 ipsec policy ospf6 dir in
ipv6 ipsec enable

```

Switch 10 VLAN configuration

The following example displays the creation of VLAN 3 and the configuration of IPsec on VLAN 3.

```

interface gigabitEthernet 1/10
no shut
exit
vlan create 3 type port-mstprstp 3
vlan members add 3 1/10 portmember
interface vlan 3
ipv6 interface enable
ipv6 interface address 2000::1/64
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
ipv6 ipsec policy ospf1 dir out
ipv6 ipsec policy ospf2 dir in
ipv6 ipsec policy ospf3 dir out
ipv6 ipsec policy ospf4 dir in
ipv6 ipsec policy ospf5 dir out
ipv6 ipsec policy ospf6 dir in
ipv6 ipsec enable

```

Switch 30 security associations

The following example displays the configuration of security associations for OSPFv3 for Switch 30.

```

ipsec security-association ospf1
ipsec security-association ospf1 encap-proto ESP
ipsec security-association ospf1 mode transport
ipsec security-association ospf1 spi 2
ipsec security-association ospf1 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf1 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf1 key-mode manual
ipsec security-association ospf1 lifetime seconds 1
ipsec security-association ospf1 lifetime bytes 1
ipsec security-association ospf2
ipsec security-association ospf2 encap-proto ESP
ipsec security-association ospf2 mode transport
ipsec security-association ospf2 spi 1
ipsec security-association ospf2 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf2 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf2 key-mode manual
ipsec security-association ospf2 lifetime seconds 1
ipsec security-association ospf2 lifetime bytes 1
ipsec security-association ospf3
ipsec security-association ospf3 encap-proto ESP
ipsec security-association ospf3 mode transport
ipsec security-association ospf3 spi 4
ipsec security-association ospf3 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf3 Encrpt-algo AES-CTR EncrptKey

```

```

12345678901234567890123456789012 keyLength 32
ipsec security-association ospf3 key-mode manual
ipsec security-association ospf3 lifetime seconds 1
ipsec security-association ospf3 lifetime bytes 1
ipsec security-association ospf4
ipsec security-association ospf4 encap-proto ESP
ipsec security-association ospf4 mode transport
ipsec security-association ospf4 spi 3
ipsec security-association ospf4 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf4 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf4 key-mode manual
ipsec security-association ospf4 lifetime seconds 1
ipsec security-association ospf4 lifetime bytes 1
ipsec security-association ospf5
ipsec security-association ospf5 encap-proto ESP
ipsec security-association ospf5 mode transport
ipsec security-association ospf5 spi 6
ipsec security-association ospf5 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf5 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf5 key-mode manual
ipsec security-association ospf5 lifetime seconds 1
ipsec security-association ospf5 lifetime bytes 1
ipsec security-association ospf6
ipsec security-association ospf6 encap-proto ESP
ipsec security-association ospf6 mode transport
ipsec security-association ospf6 spi 5
ipsec security-association ospf6 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf6 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf6 key-mode manual
ipsec security-association ospf6 lifetime seconds 1
ipsec security-association ospf6 lifetime bytes 1

```

Switch 30 policy configuration

In the example, the local address is fe80:0:0:0:b2ad:aaff:fe43:4d00, and the remote address is fe80:0:0:0:b2ad:aaff:fe43:100. The policy has the laddr configured to the link local address and the raddr is configured to the remote link local address with the direction configured to outbound.

```

ipsec policy ospf1
ipsec policy ospf1 admin enable
ipsec policy ospf1 raddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf1 laddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf1 protocol ospv3
ipsec policy ospf1 action permit

```

Laddr is configured to the remote link local address and raddr is configured to the local link local address with the direction configured to inbound.

```

ipsec policy ospf2
ipsec policy ospf2 admin enable
ipsec policy ospf2 raddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf2 laddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf2 protocol ospfv3
ipsec policy ospf2 action permit

```


Laddr is configured to the link local address and raddr is configured to ff02::05 with the direction configured to outbound.

```
ipsec policy ospf3
ipsec policy ospf3 admin enable
ipsec policy ospf3 raddr ff02::05
ipsec policy ospf3 laddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf3 protocol ospfv3
ipsec policy ospf3 action permit
```

Laddr is configured to the remote link local address and the raddr is configured to ff02::05 with the direction configured to inbound.

```
ipsec policy ospf4
ipsec policy ospf4 admin enable
ipsec policy ospf4 raddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf4 laddr ff02::05
ipsec policy ospf4 protocol ospfv3
ipsec policy ospf4 action permit
```

Laddr is configured to the link local address and raddr is configured to ff02::06 with the direction configured to outbound.

```
ipsec policy ospf5
ipsec policy ospf5 admin enable
ipsec policy ospf5 raddr ff02::06
ipsec policy ospf5 laddr fe80:0:0:0:b2ad:aaff:fe43:4d00
ipsec policy ospf5 protocol ospfv3
ipsec policy ospf5 action permit
```

Laddr is configured to the remote link local address and raddr is configured to ff02::06 with the direction configured to inbound.

```
ipsec policy ospf6
ipsec policy ospf6 admin enable
ipsec policy ospf6 raddr fe80:0:0:0:b2ad:aaff:fe43:100
ipsec policy ospf6 laddr ff02::06
ipsec policy ospf6 protocol ospfv3
ipsec policy ospf6 action permit
```

Switch 30 link table configuration

The following example displays the linking of the policy with the security association on Switch 30.

```
ipsec policy ospf1 security-association ospf1
ipsec policy ospf2 security-association ospf2
ipsec policy ospf3 security-association ospf4
ipsec policy ospf4 security-association ospf3
ipsec policy ospf5 security-association ospf5
ipsec policy ospf6 security-association ospf6
```

Switch 30 OSPFv3 configuration

The following example displays the OSPFv3 configuration on Switch 30.

```
router ospf ipv6-enable
router ospf
ipv6 router-id 2.2.2.2
ipv6 area 0.0.0.1
```

Switch 30 interface configuration

The following example displays the interface configuration on slot/port 1/10.

```
interface gigabitEthernet 1/10
no shut
ipv6 interface vlan 3
ipv6 interface address 2001::2/64
ipv6 interface enable
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
ipv6 ipsec policy ospf1 dir out
ipv6 ipsec policy ospf2 dir in
ipv6 ipsec policy ospf3 dir out
ipv6 ipsec policy ospf4 dir in
ipv6 ipsec policy ospf5 dir out
ipv6 ipsec policy ospf6 dir in
ipv6 ipsec enable
```

Switch 30 VLAN configuration

The following example displays the creation of VLAN 3 and the configuration of IPsec on VLAN 3.

```
interface gigabitEthernet 1/10
no shut
exit
minvlan create 3 type port-mstprstp 0
vlan members add 3 1/10 portmember
interface vlan 3
ipv6 interface enable
ipv6 interface address 2001::2/64
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
ipv6 ipsec policy ospf1 dir out
ipv6 ipsec policy ospf2 dir in
ipv6 ipsec policy ospf3 dir out
ipv6 ipsec policy ospf4 dir in
ipv6 ipsec policy ospf5 dir out
ipv6 ipsec policy ospf6 dir in
ipv6 ipsec enable
```

OSPFv3 virtual link IPsec configuration example

The following example displays a network using IPsec with OSPFv3 virtual link.




Figure 150: OSPFv3 virtual link with IPsec configuration

The following example displays the configuration of IPsec with OSPFv3 virtual link. For OSPFv3 conceptual and procedural information, see [OSPF](#) on page 2748.

Switch 10 security association configuration

The following example displays the configuration of security associations for OSPFv3 for Switch 10.

```
ipsec security-association ospf1
ipsec security-association ospf1 encap-proto ESP
ipsec security-association ospf1 mode transport
ipsec security-association ospf1 spi 1
ipsec security-association ospf1 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf1 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf1 key-mode manual
ipsec security-association ospf1 lifetime seconds 1
ipsec security-association ospf1 lifetime bytes 1
```

Switch 10 OSPFv3 configuration

The following example displays the OSPFv3 configuration on Switch 10.

```
router ospf ipv6-enable
ipv6 forwarding
router ospf
ipv6 router-id 1.1.1.1
ipv6 area 0.0.0.1
ipv6 as-boundary-router
ipv6 area 0.0.0.0
```

Switch 10 virtual link and policy configuration

The following example displays the configuration of a OSPFv3 virtual link.

```
ipv6 area virtual-link 0.0.0.1 3.3.3.3
ipv6 area virtual-link 0.0.0.1 3.3.3.3 ipsec
ipv6 area virtual-link 0.0.0.1 3.3.3.3 ipsec security-association ospf1
ipv6 area virtual-link 0.0.0.1 3.3.3.3 ipsec action permit
ipv6 area virtual-link 0.0.0.1 3.3.3.3 ipsec direction both
ipv6 area virtual-link 0.0.0.1 3.3.3.3 ipsec enable
```

Switch 10 interface configuration

The following example displays the interface configuration on slot/port 1/10.

```
interface gigabitEthernet 1/10
no shut
ipv6 interface vlan 3
ipv6 interface address 2000::1/64
ipv6 interface enable
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
```

Switch 10 VLAN configuration

The following example displays the creation of VLAN 3 and the configuration of IPsec on VLAN 3.

```
interface gigabitEthernet 1/10
no shut
exit
vlan create 3 type port-mstprstp 3
vlan members add 3 1/10 port-member
interface vlan 3
ipv6 interface enable
ipv6 interface address 2000::1/64
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
```

Switch 20 OSPFv3 configuration

The following example displays the OSPFv3 configuration on Switch 20.

```
router ospf ipv6-enable
ipv6 forwarding
router ospf
ipv6 router-id 2.2.2.2
ipv6 area 0.0.0.1
```

Switch 20 interface configuration

The following example displays the interface configuration on slot/port 1/10 and 1/20.

```
interface gigabitEthernet 1/10
no shut
ipv6 interface vlan 3
ipv6 interface address 2000::2/64
ipv6 interface enable
ipv6 ospf area 0.0.0.1
ipv6 ospf enable

interface gigabitEthernet 1/20
no shut
ipv6 interface vlan 4
ipv6 interface address 2001::1/64
ipv6 interface enable
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
```

Switch 20 VLAN configuration

The following example displays the creation of VLAN 3 and the configuration of IPsec on VLAN 3 and VLAN 4.

```
interface gigabitEthernet 1/10
no shut
exit
vlan create 3 type port-mstprstp 0
vlan members add 3 1/10 portmember
interface vlan 3
ipv6 interface enable
ipv6 interface address 2000::2/64
ipv6 ospf area 0.0.0.1
ipv6 ospf enable

interface gigabitEthernet 1/20
no shut
exit
vlan create 4 type port-mstprstp 0
vlan members add 4 1/20 portmember
interface vlan 4
ipv6 interface enable
ipv6 interface address 2001::1/64
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
```

Switch 40 security association configuration

The following example displays the configuration of security associations for OSPFv3 for Switch 40.

```
ipsec security-association ospf1
ipsec security-association ospf1 encap-proto ESP
ipsec security-association ospf1 mode transport
ipsec security-association ospf1 spi 1
ipsec security-association ospf1 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
ipsec security-association ospf1 Enchrpt-algo AES-CTR EnchrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association ospf1 key-mode manual
ipsec security-association ospf1 lifetime seconds 1
ipsec security-association ospf1 lifetime bytes 1
```

Switch 40 OSPFv3 configuration

The following example displays the OSPFv3 configuration on Switch 40.

```
router ospf ipv6-enable
ipv6 forwarding
router ospf
ipv6 router-id 3.3.3.3
ipv6 area 0.0.0.1
ipv6 area 0.0.0.2
ipv6 as-boundary-router
```

Switch 40 OSPFv3 virtual link and policy configuration

The following example displays the configuration of a OSPFv3 virtual link.

```
ipv6 area virtual-link 0.0.0.1 1.1.1.1
ipv6 area virtual-link 0.0.0.1 1.1.1.1 ipsec
ipv6 area virtual-link 0.0.0.1 1.1.1.1 ipsec security-association ospf1
ipv6 area virtual-link 0.0.0.1 1.1.1.1 ipsec action permit
```

```
ipv6 area virtual-link 0.0.0.1 1.1.1.1 ipsec direction both
ipv6 area virtual-link 0.0.0.1 1.1.1.1 ipsec enable
```

Switch 40 interface configuration

The following example displays the interface configuration on slot/port 1/20.

```
interface gigabitEthernet 1/20
no shut
ipv6 interface vlan 4
ipv6 interface address 2001::2/64
ipv6 interface enable
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
```


Switch 40 VLAN interface configuration

The following example displays the creation of VLAN 4 and the configuration of IPsec on VLAN 4.

```
interface gigabitEthernet 1/20
no shut
exit
vlan create 4 type port-mstprstp 0
vlan members add 4 1/20
interface vlan 4
ipv6 interface enable
ipv6 interface address 2001::2/64
ipv6 ospf area 0.0.0.1
ipv6 ospf enable
```

IPsec configuration of TCP

The following example displays the configuration of IPsec for TCP.



Switch 10 IPsec security association configuration

The following example displays the configuration of the IPsec security association for TCP for Switch 10.

```
ipsec security-association tcp1
ipsec security-association tcp1 encap-proto ESP
ipsec security-association tcp1 mode transport
ipsec security-association tcp1 spi 100
ipsec security-association tcp1 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
```

```
ipsec security-association tcp1 Encrpt-algo AES-CTR EncrptKey
12345678901234567890123456789012 keyLength 32
ipsec security-association tcp1 key-mode manual
ipsec security-association tcp1 lifetime seconds 1
ipsec security-association tcp1 lifetime bytes 1
```

Switch 10 IPsec policy configuration

The following example displays the configuration of the IPsec policy for TCP for Switch 10.

```
ipsec policy tcp1
ipsec policy tcp1 admin enable
ipsec policy tcp1 raddr 2000::2
ipsec policy tcp1 raddr 2000::2 laddr 2000::1
ipsec policy tcp1 raddr 2000::2 protocol tcp sport 23 dport 23
ipsec policy tcp1 raddr 2000::2 action permit
```

Switch 10 linking the IPsec policy with the IPsec security association

The following example displays the linking of the IPsec policy with the IPsec security association

```
ipsec policy tcp1 security-association tcp1
```

Switch 10 interface configuration

The following examples displays the configuration of IPsec for slot/port 1/10.

```
interface gigabitEthernet 1/10
no shut
ipv6 interface vlan 3
ipv6 interface address 2000::1/64
ipv6 interface enable
ipv6 ipsec policy tcp1 dir both
ipv6 ipsec enable
```

Switch 10 VLAN configuration

The following example displays the creation and configuration of VLAN 3.

```
interface gigabitEthernet 1/10
no shut
exit
vlan create 3 type port-mstprstp 3
vlan members add 3 1/10 portmember
interface vlan 3
ipv6 interface enable
ipv6 interface address 2000::1/64
ipv6 ipsec policy tcp1 dir both
ipv6 ipsec enable
```

Switch 30 IPsec security association configuration

The following example displays the configuration of the IPsec security association for TCP for Switch 10.

```
ipsec security-association tcp1
ipsec security-association tcp1 encap-proto ESP
ipsec security-association tcp1 mode transport
ipsec security-association tcp1 spi 100
ipsec security-association tcp1 auth-algo MD5 auth-key 12345678901234567890123456789012
keyLength 32
```

```
ipsec security-association tcp1 Encrpt-algo AES-CTR EncrptKey  
12345678901234567890123456789012 keyLength 32  
ipsec security-association tcp1 key-mode manual  
ipsec security-association tcp1 lifetime seconds 1  
ipsec security-association tcp1 lifetime bytes 1
```

Switch 30 IPsec policy configuration

The following example displays the configuration of the IPsec policy for TCP for Switch 10.

```
ipsec policy tcp1  
ipsec policy tcp1 admin enable  
ipsec policy tcp1 raddr 2000::1  
ipsec policy tcp1 raddr 2000::1 laddr 2000::2  
ipsec policy tcp1 raddr 2000::1 protocol tcp sport 23 dport 23  
ipsec policy tcp1 raddr 2000::1 action permit
```

Switch 30 linking the IPsec policy with the IPsec security association

The following example displays the linking of the IPsec policy with the IPsec security association

```
ipsec policy tcp1 security-association tcp1
```

Switch 30 interface configuration

The following examples displays the configuration of IPsec for slot/port 1/10.

```
interface gigabitEthernet 1/10  
no shut  
ipv6 interface vlan 3  
ipv6 interface address 2000::2/64  
ipv6 interface enable  
ipv6 ipsec policy tcp1 dir both  
ipv6 ipsec enable
```

Switch 30 VLAN configuration

The following example displays the creation and configuration of VLAN 3.

```
interface gigabitEthernet 1/10  
no shut  
exit  
vlan create 3 type port-mstprstp 3  
vlan members add 3 1/10 portmember  
interface vlan 3  
ipv6 interface enable  
ipv6 interface address 2000::2/64  
ipv6 ipsec policy tcp1 dir both  
ipv6 ipsec enable
```




IPv4 Routing Basics

[IPv4 Routing Operations Fundamentals](#) on page 2099

[IP routing configuration using the CLI](#) on page 2115

[IP routing configuration using Enterprise Device Manager](#) on page 2142

This section provides conceptual information and procedures to configure IP Routing using CLI and EDM.

The 4220 Series switches support only a single active VRF with IP configuration.



Important

For the 5320 Series, only the following models support more than one VRF with IP configuration:

- 5320-16P-2MXT-2X
- 5320-24T-4X-XT
- 5320-24T-24S-4XE-XT
- 5320-48P-8XE
- 5320-48T-8XE

IPv4 Routing Operations Fundamentals

Use the information in this section to understand IPv4 routing.

For more information about specific routing protocols, see the following sections:

- [BGP](#) on page 536
- [OSPF](#) on page 2748
- [RIP](#) on page 3149

IP Addressing

An IP version 4 address consists of 32 bits expressed in dotted-decimal format (x.x.x.x). The IP version 4 address space is divided into classes, with classes A, B, and C reserved for unicast addresses and accounting for 87.5 percent of the 32-bit IP address space.

Class D is reserved for multicast addressing. The following table lists the breakdown of IP address space by address range and mask.

Class	Address range	Mask	Number of addresses
A	1.0.0.0 to 126.0.0.0	255.0.0.0	126
B	128.0.0.0 to 191.0.0.0	255.255.0.0	$127 * 255$
C	192.0.0.0 to 223.0.0.0	255.255.255.0	$31 * 255 * 255$
D	224.0.0.0 to 239.0.0.0		

To express an IP address in dotted-decimal notation, you convert each octet of the IP address to a decimal number and separate the numbers by decimal points. For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary, has a different boundary point between the network and host portions of the address as illustrated in the following figure. The network portion is a network number field from 8 through 24 bits. The remaining 8 through 24 bits identify a specific host on the network.

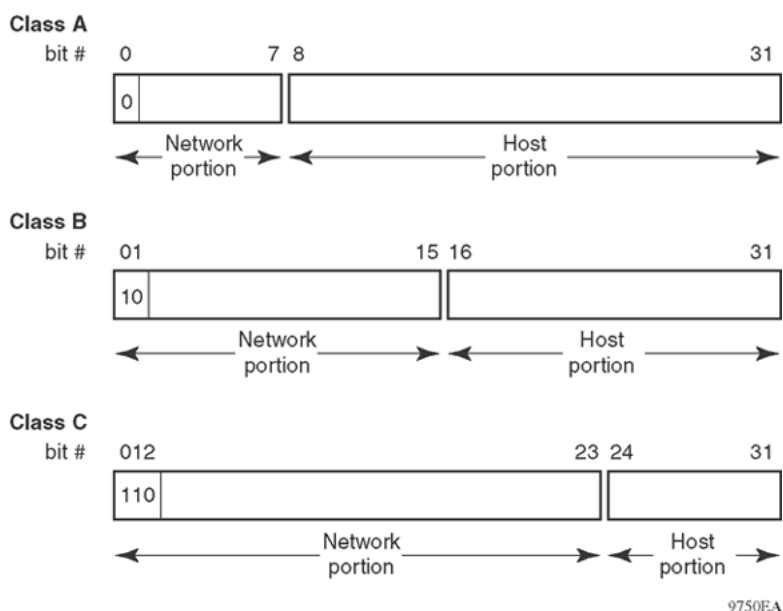


Figure 151: Network and host boundaries in IP address classes

Subnet Addressing

Subnetworks (or subnets) extend the IP addressing scheme an organization uses to one with an IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

You create a subnet address by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is found in the first octet

of the host portion (10). A subnet mask is applied to the IP address and identifies the network and host portions of the address.

The following table illustrates how subnet masks used with class B and class C addresses can create differing numbers of subnets and hosts. This example includes the zero subnet, which is permitted on the switch.

Table 156: Subnet masks for class B and class C IP addresses

Number of bits	Subnet mask	Number of subnets	Number of hosts for each subnet
Class B			
2	255.255.192.0	2	16 382
3	255.255.224.0	6	8 190
4	255.255.240.0	14	4 094
5	255.255.248.0	30	2 046
6	255.255.252.0	62	1 022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1 022	62
11	255.255.255.224	2 046	30
12	255.255.255.240	4 094	14
13	255.255.255.248	8 190	6
14	255.255.255.252	16 382	2
Class C			
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2


You use variable-length subnet masking (VLSM) to divide your intranet into pieces that match your requirements. Routing is based on the longest subnet mask or network that matches. Routing Information Protocol version 2 and Open Shortest Path First are routing protocols that support VLSM.

Supernet Addressing and CIDR

A supernet, or classless interdomain routing (CIDR) address, is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed. You can use

supernetting to address an entire block of class C addresses and avoid using large routing tables to track the addresses.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the class C addresses shown in the following figure. By adding the mask 255.255.128.0 to IP address 192.32.128.0, you aggregate the addresses 192.32.128.0 through 192.32.255.255 and 128 class C addresses use a single routing advertisement. In the bottom half of the following figure, you use 192.32.0.0/17 to aggregate the 128 addresses (192.32.0.0/24 to 192.32.127.0/24).



9577EA

Figure 152: Class C address supernet

Another example is the block of addresses 192.32.0.0 to 192.32.7.0. The supernet address for this block is 11000000 00100000 00000, with the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an address and mask pair:

- The address is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- The mask is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/21.

Although classes prohibit using an address mask with the IP address, you can use CIDR to create networks of various sizes using the address mask. With CIDR, the routers outside the network use the addresses.

IPv4 Interface State

Table 157: IPv4 Interface state product support

Feature	Product	Release introduced
Ability to Enable or Disable an IP Interface	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.10
	5420 Series	Fabric Engine 8.10
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.10
	VSP 7400 Series	VOSS 8.10

To create an IP interface (to enable IP-Layer functionality), you must assign a primary IP address to the VLAN or brouter port. By default, the switch enables the IP interface after you configure the primary IP address.

You can administratively enable or disable an IPv4 interface after you create it or you can instruct the switch to create the IP interface in a disabled state when you configure the primary IP address. The following examples provide different use cases to change the administrative state of an IP interface.

To change the feature configuration on an existing IP interface, you must perform the following actions:

1. Disable the IP interface.
2. Make the configuration change.
3. Reenable the IP interface.

To preconfigure a new interface, including routing protocols on the interface, and enable it later during a maintenance window, you can configure the primary IP address for the interface and configure the disabled state simultaneously.

For more information, see the following CLI tasks:

- [Assign an IP address to a Port](#) on page 2122
- [Modify an Existing IP Interface](#) on page 2124
- [Assign an IP Address to a VLAN](#) on page 4013

Operational Notes

The following list identifies feature restrictions:

- You can change the IP interface administrative state of IPv4 interfaces only.
- You can change the IP interface (IP-layer) administrative state only for VLAN and brouter ports.
- In the case of coexistence for both a NOS routing VLAN and VLAN Management Instance, the IP interface state change only affects the local VLAN.
- Layer 3 information passes to DvR leaf nodes only if you enable the IP interface on the DvR controller.
- You cannot change the administrative state of a VLAN that does not support an IP address configuration, for example, a VLAN used for IP Multicast config-lite for Fabric Connect.
- You cannot change the administrative state of an IP interface on a Shortest Path Bridging MAC (SPBM) VLAN.
- To change the administrative state of a VLAN used with virtual interswitch trunk (vIST), you must first disable vIST.
- The ability to create a new IP interface in a disabled state in one step is available only using CLI and not EDM.

Loopback

Circuitless IP (CLIP) is a virtual (or loopback) interface that is not associated with a physical port. You can use the CLIP interface to provide uninterrupted connectivity to your device as long as a path exists to reach the device.

For example, as shown in the following figure, a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Use an interior Border Gateway Protocol (iBGP) session between two additional addresses, 195.39.128.1/30 (CLIP 1) and 195.39.281.2/30 (CLIP 2).

CLIP 1 and CLIP 2 represent the virtual CLIP addresses that you configure between R1 and R2. These virtual interfaces are not associated with the physical link or hardware interface, which permits the BGP session to continue as long as a path exists between R1 and R2. An IGP (such as OSPF) routes addresses that correspond to the CLIP addresses. After the routers learn all the CLIP addresses in the AS, the system establishes iBGP and exchanges routes.

The system advertises loopback routes to other routers in the domain either as external routes using the route-redistribution process, or after you enable OSPF in passive mode to advertise an OSPF internal route.

You can also use CLIP for PIM-SM, typically, as a Rendezvous Point (RP), or as a source IP address for sending SNMP traps and Syslog messages.




Figure 153: Routers with iBGP connections

The system treats the CLIP interface as an IP interface. The network associated with the CLIP is treated as a local network attached to the device. This route always exists and the circuit is always up because there is no physical attachment.

Static Routes

Table 158: Static routing product support

Feature	Product	Release introduced
Static routing	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Static routes for the management OOB, if supported, and management VLAN (non-routing mode) must use the Segmented Management Instance. For more information, see [Segmented Management](#) on page 106. The management CLIP can use the Segmented Management Instance or routes in the associated VRF routing table manager (RTM). The management VLAN (routing mode) uses the Fabric Engine routing table or routes in the associated VRF.

A static route is a route to a destination IP address that you manually create.

Create static routes to enhance network stability. Use the local next hop option to configure a static route with or without local next hop.

You can configure static routes with a next hop that is not directly connected, but that hop must be reachable. Otherwise, the static route is not enabled.

Layer 3 redundancy supports only address resolution protocol (ARP) and static route. Static ARP must configure the nonlocal next-hop of static routes. No other dynamic routing protocols provide nonlocal next-hop.

**Note**

Static ARP entries are not supported for NLB Unicast or NLB Multicast operations.

You can use a default static route to specify a route to all networks for which no explicit routes exist in the forwarding information base or the routing table. This route has a prefix length of zero (RFC1812). You can configure the switch with a route through the IP static routing table.

To create a default static route, you must configure the destination address and subnet mask to 0.0.0.0.

**Note**

Do not configure static routes on a DvR Leaf node unless the configuration is for reachability to a management network using a brouter port. Also, configuring the preference of static routes is not supported on a Leaf node.

Static Route Tables

A router uses the system routing table to make forwarding decisions. In the static route table, you can change static routes directly. Although the two tables are separate, the static route table manager entries are automatically reflected in the system routing table if the next-hop address in the static route is reachable, and if the static route is enabled.

The system routing table displays only active static routes with a best preference. A static route is active only if the route is enabled and the next-hop address is reachable (for example, if a valid ARP entry exists for the next hop).

You can enter multiple routes (for example, multiple default routes) that have different costs, and the routing table uses the lowest cost route that is available. However, if you

enter multiple next hops for the same route with the same cost, the software does not replace the existing route. If you enter the same route with the same cost and a different next-hop, the first route is used. If the first route becomes unreachable, the second route (with a different next-hop) is activated with no connectivity loss.

IPv4 Discard Static Routes

A discard static route is a route with an invalid next hop that results in traffic matching the route (and not matching the more specific routes) being discarded.

While the router aggregates or injects routes to other routers, the router does not have a path to the aggregated destination. In such cases, the result is a discard static route and a routing loop. To avoid routing loops, configure a discard static route to the destination the router is advertising.

You can configure a preference value for a discard route. However, you must configure that preference value appropriately so that when you want to use the discard route, it is elected as the best route.

Before you add a discard static route, perform a check to ensure that no other static route to that identical destination is enabled. If such a route exists, you cannot add the discard route and the system displays an error message.

If you enable a discard route, you cannot add another static route to that destination. You must first delete or disable the discard route before you add a regular static route to that destination.

Only discard routes that belong to the static type protocol are supported. An inter-VRF discard route is not installed in the routing table of the destination VRF on the same switch.

Routing Protocols

Routers and routing switches use routing protocols to exchange reachability information. Routers use a routing protocol to advertise available paths on which the router can forward data. The routers use the protocol to determine the most efficient path to use. Routers use dynamic routing protocols to avoid sending data to inoperable links, and to send data to links that generally result in the fastest transmission times.

The switch routes frames using one of the following dynamic unicast IP routing protocols for path selection:

- Routing Information Protocol version 1 (RIPv1) (RFC 1058)
- RIPv2 (RFC 2453)
- Open Shortest Path First version 2 (OSPFv2) (RFC 2328)
- OSPFv3 (RFC 2740)
- Border Gateway Protocol version 4 (BGPv4) (RFC 1771)

Unlike static IP routing, where you must create a manual entry in the routing table to specify a routing path, dynamic IP routing uses a learning approach to determine the paths and routes to other routers. Dynamic routing uses two basic types of routing:

distance vector and link-state. Routing Information Protocol (RIP) is a distance vector protocol and Open Shortest Path First (OSPF) is a link-state protocol.

The switch uses routing protocols like OSPF and RIP to populate routing tables. Routers use a routing protocol to exchange network topology information. A router uses the IP address of an incoming data packet to send the packet according to the routing tables.

The most commonly used unicast routing protocols include OSPF, RIP, and BGP. For more information about BGP, see [BGP](#) on page 536. For information about multicast routing protocols, see [IP Multicast](#) on page 1696. For information about OSPFv3 routing protocols, see [OSPF](#) on page 2748.

VLANs and Routing

When traffic is routed on a virtual local area network (VLAN), an IP address is assigned to the VLAN and is not associated with a particular physical port. Brouter ports are VLANs that route IP packets and bridge nonroutable traffic in a single-port VLAN.

Virtual Routing Between VLANs

The switch supports wire-speed IP routing between VLANs. As shown in the following figure, VLAN 1 and VLAN 2 are on the same device, yet for traffic to flow from VLAN 1 to VLAN 2, the traffic must be routed.

When you configure routing on a VLAN, you assign an IP address to the VLAN, which acts as a virtual router interface address for the VLAN (a virtual router interface is not associated with a particular port). You can reach the VLAN IP address through the VLAN ports, and frames are routed from the VLAN through the gateway IP address. Routed traffic is forwarded to another VLAN within the device.




Figure 154: IP routing between VLANs

When Spanning Tree Protocol is enabled in a VLAN, the spanning tree convergence must be stable before the routing protocol begins. This requirement can lead to an additional delay in the IP traffic forwarding.

Because a port can belong to multiple VLANs (some of which are configured for routing on the device and some of which are not), a one-to-one correspondence no longer exists between the physical port and the router interface.

As with an IP address, virtual router interface addresses using Virtual Router Redundancy Protocol (VRRP) are also used for device management. For Simple Network Management Protocol (SNMP) or Telnet management, you can use virtual router interface address to access the device as long as routing is enabled on the VLAN.

Router Ports

The switch also supports router ports. A router port is a single-port VLAN that routes IP packets and bridges all nonroutable traffic. The difference between a router port and a standard IP protocol-based VLAN configured to route traffic is that the routing interface of the router port is not subject to the spanning tree state of the port. A router port can be in the blocking state for nonroutable traffic and still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

Because a router port is a single-port VLAN, each router port decreases the number of available VLANs by one and uses one VLAN ID.

The switch allows IP routing to be enabled on VLANs and router ports. For the maximum number of interfaces, see the Software scaling capabilities section of the *Fabric Engine Release Notes*.

Equal Cost Multipath

Table 159: Equal Cost Multiple Path for IPv4 product support

Feature	Product	Release introduced
Equal Cost Multiple Path (ECMP) for IPv4	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

With Equal Cost Multipath (ECMP), the switch can determine up to eight equal-cost paths to the same destination prefix. You can use multiple paths for load sharing of traffic. These multiple paths provide faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use

links between routers more efficiently when sending IP traffic. Equal Cost Multipath is formed using routes from the same source or protocol.

All IP ECMP routes that share the same combination of ECMP next hops consume the same ECMP group resource. The following list illustrates how shared next hops affect resource consumption:

- Prefix 3.1.1.0/16 is learned as an ECMP route with next hops A and B, and consumes one entry in the ECMP GROUP table.
- Prefix 4.1.1.0/16 is learned as an ECMP route with the same next hops, A and B. No additional resource is taken in the ECMP GROUP table.
- Prefix 5.1.1.0/16 is learned as an ECMP route with next hops B and C, and consumes one additional entry in the ECMP GROUP table.

ECMP is supported on both the Global Routing Table (GRT) and Virtual Routing and Forwarding (VRF).

The ECMP feature supports and complements the following protocols and route types:

- Border Gateway Protocol (BGP)
- Default route
- Intermediate System-to-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Static route
- VRF

ECMP Pathlist

Use the ECMP Pathlist feature to control how many equal-cost paths to add to the routing manager for the same destination.

IP Source Routing

Table 160: IP Source Routing product support

Feature	Product	Release introduced
IP Source Routing enable or disable	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

IP Source Routing allows the sender of a packet to specify the route that the packet must travel through the network. When the Source Route option is not enabled, the router chooses the primary routing path to send the packets. If IP Source Routing flag is on, the source host dictates the datapath for the packet to reach the destination using the information contained in the IP header.

IP Source Routing is considered a security risk because it allows the users to specify their own path through the network outside of the primary forwarding path. This can cause packets to bypass the security devices. Therefore, IP Source Routing is disabled by default.

Multihoming

The switch uses the multihoming feature to support clients or servers that have multiple IP addresses associated with a single MAC address. Multihomed hosts can be connected to port-based and policy-based VLANs.

The IP addresses associated with a single MAC address on a host must be in the same IP subnet.

IPv4 ICMP Broadcast

Table 161: Internet Control Message Protocol product support

Feature	Product	Release introduced
Internet Control Message Protocol (ICMP)	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 161: Internet Control Message Protocol product support (continued)

Feature	Product	Release introduced
ICMP broadcast and multicast enable or disable	4220 Series	Fabric Engine 9.2 IPv4 only The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 161: Internet Control Message Protocol product support (continued)

Feature	Product	Release introduced
Fragmented ICMP Packet Filtering for IPv4	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5
Fragmented ICMP Packet Filtering for IPv6	4220 Series	Not Supported
	5320 Series	For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

On IPv4 networks, a packet can be directed to an individual machine or broadcast to an entire network. When a packet is sent to an IP broadcast address from a machine on the local network, that packet is delivered to all machines on that network.

If a packet that is broadcast is an ICMP echo request packet, the machines on the network receive this ICMP echo request packet and send an ICMP echo reply packet back. When all the machines on a network respond to this ICMP echo request, the result can be severe network congestion or outages.

The switch always responds to IPv4 ICMP packets sent to a broadcast address. You can disable the processing of these IPv4 ICMP packets sent to the broadcast address. On disabling the ICMP broadcast processing, all the packets containing ICMP sent to the broadcast addresses will be dropped when the packets reach the control plane.

You can disable or enable the IPv4 ICMP broadcast processing at the VRF level.

Fragmented ICMP Packet Filtering

ICMP fragmentation distributed denial-of-service (DDoS) attacks flood the destination resources with fragmented packets and overwhelm the network because of massive volumes of traffic. With Fragmented ICMP packet filtering, the system inspects each incoming IPv4 ICMP packet to determine if it should drop the packet or forward it.

You can configure ICMP drop packet filtering globally, on a specific VRF, and on the following management interfaces:

- Out-of-Band (OOB) management
- Circuitless IP (CLIP) management
- VLAN management

IP routing configuration using the CLI

Configure the IP router interface so that you can configure and use routing protocols and features on the interface. This section contains instructions for both the Global Router and Virtual Router Forwarding (VRF) instances.

Enable Routing Globally or on a VRF Instance

Use IP forwarding (routing) on a global level so that the device supports routing. You can use the IP address of an interface for IP-based network management.

About This Task

You cannot configure this option on a DvR Leaf node.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Activate IP forwarding:

```
ip routing
```

3. View the forwarding configuration:

```
show ip routing [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Activate IP forwarding and view the forwarding configuration:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#ip routing
Switch:1(router-vrf)#show ip routing

=====
                        IP - GlobalRouter
=====

IP Forwarding is enabled
IP ECMP feature is disabled
Maximum ECMP paths number is 1
ECMP 1 pathlist :
ECMP 2 pathlist :
ECMP 3 pathlist :
ECMP 4 pathlist :
ECMP 5 pathlist :
ECMP 6 pathlist :
ECMP 7 pathlist :
ECMP 8 pathlist :
Gratuitous-Arp : enable
IP Alternative Route feature is enabled
IP More Specific Non Local Route feature is disabled
IP ICMP Unreachable Message is disabled
IP Supernetting is disabled
IP Icmp-echo-broadcast-request is enabled

IP Default TTL is 255 seconds
IP ARP life time is 360 minutes
IP Source Route Option is disabled
```

Variable Definitions

Use the data in the following table to use the **show ip routing** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF instance by VRF name.
<code>vrfids WORD<0-512></code>	Specifies a VRF instance by VRF number.

Configure IP Routing Globally or for a VRF

Configure the IP routing protocol stack to specify which routing features the device can use. You can configure global parameters before or after you configure the routing protocols.

About This Task

To configure IP routing globally for a VRF instance, use VRF Router Configuration mode rather than Global Configuration mode.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the default TTL for all routing protocols to use:

```
ip ttl <1-255>
```

This value is placed into routed packets that have no TTL specified.

3. Activate the alternative route feature globally:

```
ip alternative-route
```

4. Configure the remaining global parameters as required.

Variable Definitions

The following table defines parameters for the **ip** command.

Variable	Value
<i>alternative-route</i>	Enables or disables the alternative route feature. The default value is enabled. If the alternative-route parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are readed.
<i>max-routes-trap enable</i>	Enables the device to send a trap after the maximum number of routes is exceeded.
<i>more-specific-non-local-route</i>	Enables the more-specific-non-local-route feature. If enabled, the device can enter a more-specific nonlocal route into the routing table. The default is disabled.
<i>routing</i>	Enables routing. You cannot enable IP routing on a DvR Leaf node.
<i>supernet</i>	Enables or disables supernetting. If you globally enable supernetting, the device can learn routes with a route mask of less than eight bits. Routes with a mask length less than eight bits cannot have ECMP paths, even if the ECMP feature is globally enabled. The default is disabled.
<i>ttl <1-255></i>	Configures the default time-to-live (TTL) value for a routed packet. The TTL is the maximum number of seconds before a packet is discarded. The default value of 255 is used whenever a time is not supplied in the datagram header.

The following table defines parameters for the **ip icmp** commands.

Variable	Value
<i>unreachable</i>	Enables the device to send ICMP unreachable messages. When enabled, this variable generates Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this router. These messages help determine if the device is reachable over the network. The default is disabled.

Delete a Dynamically Learned Route

About This Task

Delete a dynamically learned route from the routing table if you do not want the switch to use the route. Exercise caution when you delete entries from the routing table.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. View IP route information:

```
show ip route [<A.B.C.D>] [-s default|-s <A.B.C.D/X>] [alternative]
[count-summary] [spbm-nh-as-mac][preference] [vrf WORD<1-16>] [vrfids
WORD<0-512>] [static]
```

3. Delete the dynamically learned route:

```
no ip route <A.B.C.D A.B.C.D | A.B.C.D/X> <A.B.C.D> dynamic
```

Example

Delete a dynamically learned route:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#no ip route 192.0.2.32 255.255.255.0 198.51.100.31 dynamic
```

Variable Definitions

Use the data in the following table to use the **show ip route** commands.

Variable	Value
<i><A.B.C.D></i>	Specifies the IP address of the route to the network.
<i>alternative</i>	Displays the alternative routes.
<i>count-summary</i>	Displays a summary of the number of routes learned from each routing protocol for each VRF.
<i>preference</i>	Displays the route preference.
<i>-s <A.B.C.D/X></i>	Indicates the IP address and subnet mask for which to display routes.
<i>-s default</i>	Indicates the default subnet.
<i>static</i>	Displays the static route information.
<i>vrf WORD<1-16></i>	Displays the route for a particular VRF.
<i>vrfids WORD<0-512></i>	Displays the route for a particular VRF number.
<i>spbm-nh-as-mac</i>	Displays the spbm route next hop as mac.

Use the data in the following table to use the **no ip route** command.

Variable	Value
<code><A.B.C.D A.B.C.D A.B.C.D/X><A.B.C.D></code>	Specifies the IP address, the subnet mask, and the next-hop IP address, respectively. Specify the IP address and mask for the route destination in one of the following formats: <ul style="list-style-type: none"> A.B.C.D A.B.C.D A.B.C.D/X
<code>dynamic</code>	Specifies that a dynamic route is to be deleted.
<code>enable</code>	Disables the route.
<code>local-next-hop enable</code>	Disables the local-next-hop option.
<code>preference</code>	Deletes the value of the route preference.
<code>next-hop-vrf WORD<1-16></code>	Specifies the name of the next-hop VRF router.

Configure IP Route Preferences

Before You Begin

- Disable ECMP before you configure route preferences.



Important

Changing route preferences can affect system performance and network accessibility while you perform the procedure. You must therefore change a prefix list or a routing protocol before you activate the protocols.

About This Task

Configure IP route preferences to give preference to routes learned for a specific protocol. You must disable ECMP before you configure route preferences.

To configure route preferences for a VRF, access VRF Router Configuration mode, rather than Global Configuration mode.

Procedure

- Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
Optional: router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the route preference:

```
ip route preference protocol <static|ospf-intra|ospf-inter|ebgp|ibgp|
rip|ospf-extern1|ospf-extern2|isis-external|isis-internal> <0-255>
```

3. Confirm that the configuration is correct:

```
show ip route preference [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#show ip route preference
```

```
=====
```

IP Route Preference - GlobalRouter		
PROTOCOL	DEFAULT	CONFIG

LOCAL	0	0
STATIC	5	5
ISIS-INTERNAL	7	140
OSPF_INTRA	20	20
OSPF_INTER	25	25
EBGP	45	45
RIP	100	100
OSPF_E1	120	120
OSPF_E2	125	125
ISIS-EXTERNAL	130	9
IBGP	175	175

View the route preference configuration for a specific VRF, for example 444.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf test
Switch:1(router-vrf)#show ip route preference vrf 444
```

```
=====
```

IP Route Preference - VRF 444		
PROTOCOL	DEFAULT	CONFIG

LOCAL	0	0
STATIC	5	5
ISIS-INTERNAL	7	135
OSPF_INTRA	20	20
OSPF_INTER	25	25
EBGP	45	45
RIP	100	100
OSPF_E1	120	120
OSPF_E2	125	125
ISIS-EXTERNAL	130	11
IBGP	175	175

Variable Definitions

Use the data in the following table to use the **ip route preference protocol** command.

Variable	Value
<i>ebgp</i>	Protocol type for ebgp.
<i>ibgp</i>	Protocol type for ibgp.
<i>isis-external</i>	Protocol type for IS-IS external routes.
<i>isis-internal</i>	Protocol type for IS-IS internal routes.
<i>ospf-extern1</i>	Protocol type for ospf-extern1 routes.
<i>ospf-extern2</i>	Protocol type for ospf-extern2 routes.
<i>ospf-intra</i>	Protocol type for ospf-intra routes.
<i>ospf-inter</i>	Protocol type for ospf-inter routes.
<i>rip</i>	Protocol type for rip routes.
<i>static</i>	Protocol type for static IP routes.
<0-255>	Specifies the preference value for the specified protocol.

Assign an IP address to a Port

Assign an IP address to a port so that it supports routing operations.

Before You Begin

- Ensure routing (forwarding) is globally enabled.
- Ensure the VLAN is configured.
- If required, ensure the VRF instance exists.

About This Task

Use a brouter port to route IP packets and to bridge all nonroutable traffic. The routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic and still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

If you configure an IP interface without specifying the VRF instance, it maps to VRF 0 by default.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Assign an IP address to the port:

```
brouter port {slot/port[/sub-port]} vlan <2-4059> subnet <A.B.C.D/X>  
[mac-offset <0-511> | name WORD <0-64> | state-disabled]
```

3. (Optional) If required, associate the port with a VRF:

```
vrf WORD<1-16>
```

4. Confirm that the configuration is correct:

```
show brouter [<1-4059>]
```

Example

Assign an IP address to a port:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitethernet 1/11  
Switch:1(config-if)#brouter port 1/11 vlan 2202 subnet 47.17.10.31/255.255.255.0
```

Variable Definitions

Use the data in the following table to use the **brouter port** command.

Variable	Value
<i>mac-offset</i> <0-511>	Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated.
<i>name</i> WORD <0-64>	Specifies the IP address name.
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Variable	Value
<i>state-disabled</i>	Creates the IP interface in an administratively disabled state. This parameter is useful if you want to preconfigure an IP interface, including the routing protocol configuration for that interface, and enable it later during a maintenance window. By default, the switch enables an IP interface after you configure the primary IP address.
<i>subnet</i> <A.B.C.D/X>	Specifies the IP address and subnet mask.
<2-4059>	Specifies the VLAN ID that is used if the port is tagged (802.1q encapsulation). The VLAN ID is unique to the switch and is not used if the port is untagged.

Use the data in the following table to use the **show router** command.

Variable	Value
<1-4059>	Specifies the VLAN ID that is used if the port is tagged (802.1q encapsulation). The VLAN ID is unique to the switch and is not used if the port is untagged.

Modify an Existing IP Interface

By default, the switch enables an IP interface (a VLAN or router port) after you configure the primary IP address. You can administratively disable, and then reenabling, an IPv4 interface after you create it. For example, if you want to change the routing protocol configuration on the interface.

About This Task

You can also instruct the switch to create the IP interface in a disabled state when you initially configure the primary IP address. For more information, see [Assign an IP address to a Port](#) on page 2122 or [Configure a Primary IP Address for a VLAN](#) on page 4050.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable the IP interface:

```
no ip interface enable
```

3. Make feature configuration changes for the IP interface.
4. Reenable the IP interface:

```
ip interface enable
```

View Primary IP Addresses for all Router Interfaces

Perform the following procedure to display information about all primary IP interfaces configured on the device.

About This Task

This command displays only primary IP interfaces. To view Secondary IP Interfaces, see [Display All Configured IP Addresses](#) on page 4015.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. Show the IP interfaces and addresses on the device:

```
show ip interface
```

Example

Show the primary IP interfaces and addresses on the device:

```
Switch:1>show ip interface
=====
                        IP Interface - GlobalRouter
=====
```

INTERFACE	IP ADDRESS	NET MASK	ADMIN STATUS	OPER STATE	VLAN ID	BROUTER PORT	IPSEC STATE	ACCESS	IP NAME
Port1/2	192.0.2.1	255.255.255.0	enable	up	20	true	disable	GlobalRouter	
Vlan10	198.51.100.0	255.255.255.0	disable	down	10	false	disable	GlobalRouter	

```

All 2 out of 2 Total Num of IP interfaces displayed

```

Variable Definitions

Use the data in the following table to **show ip interface** command.

Variable	Value
<i>gigabitethernet</i>	Displays IP interface information for Gigabit Ethernet ports.
<i>vrf</i>	Displays interface information for a particular VRF.
<i>vrfids</i>	Displays interface information for particular VRF IDs.

View the IPv4 and IPv6 Interface Configuration on a Port

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Verify the interface configuration on the port:

For IPv4, enter: `show ip interface gigabitethernet`

For IPv6, enter: `show ipv6 interface gigabitethernet`

Example

Verify the IPv4 configuration on the port:

```
Switch:1(config-if)# show ip interface gigabitethernet
```

Brouter Port Ip									
PORT NUM	VRF NAME	IP_ADDRESS	NET_MASK	ADMIN STATUS	OPER STATUS	ADVERTISE WHEN_DOWN	DIRECT BCAST	RPC	RPCMODE
1/1	Glob~	192.0.2.3	255.255.255.0	enable	up	disable	disable	disable	exist-only
1/10	spbo~	198.51.100.4	255.255.255.0	enable	up	disable	disable	disable	exist-only

PORT NUM	VRF NAME	BROADCAST	REASM MAXSIZE
1/1	GlobalRouter	ones	1500
1/10	spboip	ones	1500

Verify the IPv6 configuration on the port:

```
Switch:1(config-if)#show ipv6 interface gigabitethernet
```

Port Ipv6 Interface													
IFINDX INDX	BROUTER	PHYSICAL ADDRESS	ADMIN STATE	OPER STATE	TYPE	MTU	HOP LMT	REACHABLE TIME	RETRANSMIT TIME	MCAST STATUS	IPSEC	RPC	RPCMODE
192	1/16	e4:5d:52:3c:65:02	enable	down	ETHER	1500	2	30000	1000	disable	disable	disable	existonly

Port Ipv6 Address					
IPV6 ADDRESS	BROUTER	TYPE	ORIGIN	STATUS	
2001:db8:0:0:0:0:ffff/64	1/16	UNICAST	MANUAL	INACCESSIBLE	INF
2001:db8:0:0:e65d:52ff:fe3c:6502/64	1/16	UNICAST	LINKLAYER	INACCESSIBLE	INF

1 out of 5 Total Num of Interface Entries displayed.
2 out of 10 Total Num of Address Entries displayed.

Configure ECMP

Enable Equal Cost MultiPath (ECMP) to permit routers to determine up to eight equal-cost paths to the same destination prefix. You can use the multiple paths for load-sharing of traffic, which provides fast convergence to alternative paths. By maximizing load sharing among equal-cost paths, you can maximize the efficiency of links between routers.

About This Task

To configure ECMP for a VRF instance, after you enable ECMP globally, use VRF Router Configuration mode rather than Global Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable ECMP globally:

```
ip ecmp
```
3. (Optional) Configure the maximum number of ECMP paths:

```
ip ecmp max-path <1-8>
```
4. (Optional) Configure a prefix-list for the target destination:

```
ip prefix-list WORD<1-64> <A.B.C.D/X> [ge <0-32>] [le <0-32>]
```
5. (Optional) Configure an ECMP pathlist to specify routes with the required number of paths:

```
ip ecmp pathlist-<1-8> WORD<1-64>
```
6. (Optional) Return to Privileged EXEC mode:

```
end
```
7. (Optional) Apply changes to all ECMP pathlist configurations:

```
ip ecmp pathlist-apply [vrf WORD<1-16>]
```

Example

Define which IP prefixes use ECMP and which do not.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ecmp
Switch:1(config)#ip prefix-list ecmpAllowed 192.0.2.0/24 ge 24 le 24
Switch:1(config)#ip prefix-list ecmpDenied 0.0.0.0/0 ge 0 le 32
Switch:1(config)#ip ecmp pathlist-2 ecmpAllowed
Switch:1(config)#ip ecmp pathlist-1 ecmpDenied
Switch:1(config)#end
Switch:1#ip ecmp pathlist-apply
```

Variable Definitions

The following table defines parameters for the **ip ecmp** command.

Variable	Value
<code>max-path <1-8></code>	Specifies the maximum number of ECMP paths. The default is 1.
<code>pathlist-1 WORD<0-64></code>	Specifies one equal-cost path to the same destination prefix. To remove the policy, enter a blank string.

Variable	Value
<i>pathlist-2 WORD<0-64></i>	Specifies up to two equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-3 WORD<0-64></i>	Specifies up to three equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-4 WORD<0-64></i>	Specifies up to four equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-5 WORD<0-64></i>	Specifies up to five equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-6 WORD<0-64></i>	Specifies up to six equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-7 WORD<0-64></i>	Specifies up to seven equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-8 WORD<0-64></i>	Specifies up to eight equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.
<i>pathlist-apply [vrf WORD<1-16>]</i>	Applies the pathlist configuration changes. You can optionally specify a VRF name.

The following table defines parameters for the **ip prefix-list** command.

Variable	Value
<i>WORD<0-64></i>	Specifies the prefix list name.
<i><A.B.C.D/X></i>	Specifies the IP address and network mask in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d/x • a.b.c.d/x.x.x.x • default
<i>ge <0-32></i>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
<i>le <0-32></i>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.

Configure Static Routes

Perform the steps in this task to complete one of the following configurations:

- Create static routes for data traffic in either the GRT or a specific VRF context for any platform.
- Create static routes for a VRF associated with a Segmented Management Instance CLIP or VLAN (routing mode) interface. Specify the name of the VRF context in [Step 1](#).
- Configure a discard static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.
- Configure a default route to specify a route to all networks for which there are no explicit routes in the forwarding information base or the routing table. This route has a prefix length of zero (RFC 1812). You can configure the switch with a static default route, or they can learn it through a dynamic routing protocol. To create a default static route, you configure the destination address and subnet mask to 0.0.0.0.

Static routes for the management OOB, if supported, and management VLAN (non-routing mode) must use the Segmented Management Instance. For more information, see [Segmented Management](#) on page 106. The management CLIP can use the Segmented Management Instance or routes in the associated VRF routing table manager (RTM). The management VLAN (routing mode) uses the Fabric Engine routing table or routes in the associated VRF.

Before You Begin

- If a discard static route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

About This Task



Note

Only discard routes that belong to the static type protocol are supported. An inter-VRF discard route is not installed in the routing table of the destination VRF on the same switch.

When you configure a static route with a next-hop-vrf context, you can specify a next-hop IP address that is a locally owned VRRP IP address of the system itself. However, this is not a supported configuration. The best practice is to implement an alternative method of inter-vrf route sharing, such as route redistribution or ISIS accept polices.

For route scaling information and for information on the maximum number of static routes supported on the hardware platform, see *Fabric Engine Release Notes*.



Note

As a best practice, do not configure static routes on a DVR Leaf node unless the configuration is for reachability to a management network using a Brouter port.

You cannot configure the preference of static routes on a Leaf node.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create one of the following types of static routes:

- A regular static route: `ip route {A.B.C.D | A.B.C.D/X} {A.B.C.D} weight <1-65535>`
- A discard static route by keyword: `ip route {A.B.C.D | A.B.C.D/X} discard weight <1-65535> [name WORD<0-64>]`
- A discard route by next-hop address: `ip route {A.B.C.D | A.B.C.D/X} {A.B.C.D} weight <1-65535> [name WORD<0-64>]`
- A default static route: `ip route 0.0.0.0 0.0.0.0 <A.B.C.D> weight <1-65535>]`

3. Enable one of the following types of static route:

- A regular static route: `ip route <A.B.C.D A.B.C.D | A.B.C.D/X> <A.B.C.D> enable`
- A discard static route: `ip route <A.B.C.D A.B.C.D | A.B.C.D/X> {discard | {A.B.C.D}} enable`
- A default static route: `ip route 0.0.0.0 0.0.0.0 <A.B.C.D>`

4. Use the following variable definitions table to configure other static route parameters as required.

When you specify a route preference, appropriately configure the preference so that when a discard route is used, it is elected as the best route.

5. View existing IP static routes for the device:

```
show ip route static
```

Examples

Create and enable an IP static route, and then view the existing IP static routes for the device. If the static route is inactive, the Access column displays no value.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip route 192.0.2.2 255.255.0.0 198.51.100.24 weight 20 preference 1 name ExtSer
Switch:1(config)#ip route 192.0.2.2 255.255.0.0 198.51.100.24 enable
Switch:1(config)#show ip route static
```

```
=====
                        IP Static Route - GlobalRouter
=====
```

DEST	MASK	NEXT	NH-VRF	COST	PREF	LCLNHOP	STATUS	ENABLE	ORIGIN	ACCESS	NAME
192.0.2.2	255.255.0.0	198.51.100.24			20						ExtSer
192.0.2.2	255.255.0.0	198.51.100.24									


```
-----
192.0.2.2 255.255.255.0 198.51.100.24 GlobalRouter 20 1 TRUE ACTIVE TRUE CONFIG ExtSer
```

Create and enable a discard static route.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip route 192.0.2.0 255.255.0.0 255.255.255.255 weight 200
Switch:1(config)#ip route 192.0.2.0 255.255.0.0 255.255.255.255 enable
```

Create and enable a default static route.

```
Switch:1>enable
Switch:1>configure terminal
Switch:1(config)#ip route 0.0.0.0 0.0.0.0 192.0.2.128 weight 100
Switch:1(config)#ip route 0.0.0.0 0.0.0.0 192.0.2.128 enable
```

Variable Definitions

The following table defines parameters for the **ip route** command.

Variable	Value
<A.B.C.D A.B.C.D A.B.C.D/X> <A.B.C.D> <A.B.C.D> <A.B.C.D> <A.B.C.D>	Specifies the IP address and mask for the route destination in one of the following formats: <ul style="list-style-type: none"> A.B.C.D A.B.C.D A.B.C.D/X The third <A.B.C.D> specifies the IP address of the next-hop router (the next router at which packets must arrive on this route). When you create a discard static route, configure this parameter to 255.255.255.255. This indicates that the traffic to this route is discarded.
<i>disable</i>	Disables a route to the router or VRF.
<i>discard</i>	Specifies the static route as a discard static route.
<i>enable</i>	Adds a static route to the router or VRF.
<i>local-next-hop enable</i>	Enables the local next hop for this static route.
<i>name</i> <1-64>	Specifies the name of the static route. You can name the route before or after it is created. Only 32 characters display. The tilde (~) symbol indicates that the name is truncated.
<i>next-hop-vrf</i> <WORD 1-16>	Specifies the next-hop VRF instance by name. After you configure the <i>next-hop-vrf</i> parameter, the static route is created in the local VRF, and the next-hop route is resolved in the next-hop VRF instance (next-hop-vrf).
<i>preference</i> <1-255>	Specifies the route preference.
<i>weight</i> <1-65535>	Specifies the static route cost. Removed note for VOSS-32396

The following table defines optional parameters for the **show ip route static** command.

Variable	Value
<A.B.C.D>	Specifies the route by IP address.
-s { <A.B.C.D> <A.B.C.D> default }	Specifies the route by IP address and subnet mask.
name <1-64>	Specifies the name of the static route. You can name the route before or after it is created. Only 32 characters display. The tilde (~) symbol indicates that the name is truncated.
trusted-segment {vrf WORD<1-16> vrfids WORD<0-512>}	Displays only routes from the trusted segment.
unrestricted-segment {vrf WORD<1-16> vrfids WORD<0-512>}	Displays only routes from the unrestricted segment, if the associated VRF is a Segmented VRF.
untrusted-segment {vrf WORD<1-16> vrfids WORD<0-512>}	Displays only routes from the untrusted segment, if the associated VRF is a Segmented VRF.
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

Configure IPv4 Fragmented ICMP Packet Filtering

About This Task

Use this task to enable IPv4 Fragmented ICMP packet filtering globally or on a specific VRF. This feature is disabled by default.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable Fragmented ICMP packet filtering:

```
ip icmp drop-fragments
```

Enable ICMP Router Discovery Globally

About This Task

Enable Router Discovery globally so that the device supports Router Discovery. Use ICMP Router Discovery to enable hosts attached to the broadcast network to discover the IP addresses of their neighboring routers.

If you enable ICMP Router Discovery globally, you automatically enable it for all VLANs. If you do not require ICMP Router Discovery on a specific VLAN, you must manually disable the feature.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable ICMP Router Discovery on the device:

```
ip irdp
```

3. Confirm that Router Discovery is enabled:

```
show ip irdp [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Enable ICMP router discovery on the device and confirm that router discovery is enabled.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip irdp
Switch:1(config)#show ip irdp
      VRF "GlobalRouter" (Global Routing Table) : Router Discovery Enabled
```

Variable Definitions

Use the data in the following table to **show ip irdp** command.

Variable	Value
<i>interface</i>	Displays route discovery interface information.
<i>vrf WORD<1-16></i>	Displays route discovery for particular VRF.
<i>vrfids WORD<0-512></i>	Displays route discovery for particular VRF IDs.

Enable or Disable IPv4 ICMP Broadcast Globally

On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when they reach the control plane.

About This Task

Use these commands to enable or disable the IPv4 ICMP broadcast feature on the global router.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable IPv4 ICMP broadcast feature, enter:

```
ip icmp echo-broadcast-request
```

3. Disable IPv4 ICMP broadcast feature, enter:

```
no ip icmp echo-broadcast-request
```

4. Set the IPv4 ICMP broadcast feature to default state, enter:

```
default ip icmp echo-broadcast-request
```



Note

By default, the IPv4 ICMP broadcast feature is enabled.

5. View the IPv4 ICMP broadcast feature state:

```
show ip routing
```

Enable or Disable IPv4 ICMP Broadcast per VRF

On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when they reach the control plane.

About This Task

Use these commands to enable or disable the IPv4 ICMP broadcast feature on the VRF router.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable IPv4 ICMP broadcast feature, enter:

```
ip icmp echo-broadcast-request
```

3. Disable IPv4 ICMP broadcast feature, enter:

```
no ip icmp echo-broadcast-request
```

4. Set the IPv4 ICMP broadcast feature to default state, enter:

```
default ip icmp echo-broadcast-request
```



Note

By default, the IPv4 ICMP broadcast feature is enabled.

5. View the IPv4 ICMP broadcast feature state:

```
show ip routing
```

View IPv4 ICMP Statistics

View the collective IPv4 ICMP statistics for all VRF instances.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View IPv4 ICMP statistics:

```
show ip icmp statistics
```

Example

View IPv4 ICMP statistics:

```
Switch:1#show ip icmp statistics

Icmp Statistics:
-----
ICMP IN STATISTICS (includes GRT and all VRF instances) :
IcmpInMsgs           = 0
IcmpInErrors         = 0
IcmpInDestUnreachs   = 0
IcmpInTimeExcds     = 0
IcmpInParmProbs      = 0
IcmpInSrcQuenchs     = 0
IcmpInRedirects      = 0
IcmpInEchos          = 0
IcmpInEchoReps       = 0
```

```

IcmpInTimestamps      = 0
IcmpInTimestampReps   = 0
IcmpInAddrMasks       = 0
IcmpInAddrMaskReps    = 0
ICMP OUT STATISTICS (includes GRT and all VRF instances) :
IcmpOutMsgs           = 0
IcmpOutErrors         = 0
IcmpOutDestUnreachs   = 0
IcmpOutTimeExcds      = 0
IcmpOutParmProbs      = 0

IcmpOutSrcQuenchs     = 0

IcmpOutRedirects      = 0

IcmpOutEchos          = 0
IcmpOutEchoReps       = 0

IcmpOutTimestamps     = 0
IcmpOutTimestampReps  = 0

IcmpOutAddrMasks      = 0
IcmpOutAddrMaskReps   = 0

0 messages dropped due to rate limiting

```

Configure Router Discovery on a Port or VLAN

Enable Router Discovery so that the device forwards Router Discovery Advertisement packets to the VLAN or port.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Specify the address placed in advertisement packets:

```
ip irdp address <A.B.C.D>
```

3. Enable the interface to send the advertisement packets:

```
ip irdp multicast
```

4. Configure other Router Discovery parameters for the interface as required.

Example

Log on to the GigabitEthernet Interface mode:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/16
```

Specify the address placed in advertisement packets to the all-systems multicast address:

```
Switch:1(config-if)# ip irdp address 244.0.0.1
```

Enable the interface to send the advertisement packets:

```
Switch:1(config-if)# ip irdp multicast
```

Configure the lifetime for advertisements:

```
Switch:1(config-if)# ip irdp holdtime 180
```

Variable Definitions

Use the data in the following table to use the **ip irdp** command.

Variable	Value
<i>address</i> <A.B.C.D>	Specifies the IP destination address use for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default address is 255.255.255.255.
<i>holdtime</i> <4-9000>	Configures the lifetime for advertisements. The default form of this command is <code>default ip irdp holdtime</code> .
<i>maxadvertinterval</i> <4-1800>	Specifies the maximum time (in seconds) that elapses between unsolicited router advertisement transmissions from the router interface. The default is 600 seconds.
<i>minadvertinterval</i> <3-1800>	Specifies the minimum time (in seconds) that elapses between unsolicited router advertisement transmissions from the interface. The range is 3 seconds to <i>maxadvertinterval</i> . The default is 450 seconds.
<i>multicast</i>	Specifies if multicast advertisements are sent. The no form of this command is <code>no ip irdp multicast</code> .
<i>preference</i> <-2147483648-2147483647>	Specifies the preference (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet. The default is 0.

Configure a CLIP Interface

About This Task

Configure a circuitless IP (CLIP) interface to provide a virtual interface that is not associated with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to your device.

For scaling information and for information on the maximum number of CLIP interfaces you can configure on your device, see *Fabric Engine Release Notes*.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create or access a CLIP interface:

```
interface loopback <1-256>
```

<1-256> indicates the identification number for the CLIP.

The command prompt changes to indicate you now access the Loopback Interface Configuration mode.

3. Configure an IP address and name for the interface:

```
ip address [<1-256>] <A.B.C.D/X> [vrf WORD<1-16> | vrf WORD<1-16>
<trusted | unrestricted>] [name WORD <0-64>]
```

4. View the IP address on the CLIP interface:

```
show ip interface
```

Example

Create a CLIP interface.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 3
Switch:1(config-if)#ip address 23.23.23.23 255.255.255.0 name ExtNet
Switch:1>show ip interface
```

IP Interface - GlobalRouter									
INTERFACE	IP ADDRESS	NET MASK	ADMIN STATUS	OPER STATUS	VLAN ID	BROUTER PORT	IPSEC STATE	ACCESS	IP NAME
Port1/2	10.3.4.2	255.255.255.0	enable	up	3	true	disable	GlobalRouter	address1
Clip1	1.2.3.4	255.255.0.0	enable	up	--	false	disable	GlobalRouter	
Clip2	2.3.4.5	255.255.255.0	enable	up	--	false	disable	GlobalRouter	
Clip23	23.23.23.23	255.255.255.0	enable	up	--	false	disable	GlobalRouter	ExtNet
Vlan2	192.0.2.5	255.255.255.0	enable	up	2	false	disable	GlobalRouter	
Vlan55	55.55.55.55	255.255.255.0	enable	up	55	false	disable	GlobalRouter	Boston

All 6 out of 6 Total Num of IP interfaces displayed

Variable Definitions

Use the data in the following table to use the **ip** commands.

Variable	Value
<code>address [<1-256>] <A.B.C.D/X> [vrf WORD<1-16>] [name WORD<0-64]</code>	Specifies the IP address for the CLIP interface. <1-256> specifies the interface. <A.B.C.D/X> specifies the IP address and mask. vrf WORD<1-16> specifies an associated VRF by name. name WORD<1-16> specifies a name for the IP address.
<code>area [<1-256>] {A.B.C.D} [vrf WORD<1-16>]</code>	Designates an area for the CLIP interface. vrf WORD<1-16> specifies an associated VRF by name
<code>ospf [<1-256>] [vrf WORD<1-16>]</code>	Enables OSPF on the CLIP interface.

Display BFD Configurations on the Loopback Interface

About This Task

Use the following procedure to display all BFD configurations on the loopback interface.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the BFD configurations:

```
show ip bfd interfaces loopback
```

Example

The following example displays BFD configurations on the Loopback Interface:

```
Switch:1>enable
Switch:1#show ip bfd interfaces loopback
=====
Circuitless IP Interface Bfd
=====
```

INTF ID	STATUS	MIN_RX	INTERVAL	MULTIPLIER	VRF-ID
1	enable	200	200	3	0
2	enable	200	200	3	2

Display the IPv4 Global Configuration

About This Task

Use this task to view Fragmented ICMP packet filtering on a IPv4 network.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display the ICMP IPv4 information:

```
show ip global [vrf WORD<1-16> | vrfids WORD<0-512>]
```

View TCP and UDP Information

Use this procedure to view TCP and UDP configuration information for IPv4.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the IPv4 TCP connection information:

```
show ip tcp connections
```

3. View the IPv4 TCP connection information for a specific vrf or vrfids:

```
show ip tcp connections vrf WORD<0-16>
```

```
show ip tcp connections vrfids WORD<0-512>
```

4. View IPv4 TCP properties:

```
show ip tcp properties
```

5. View IPv4 TCP statistics

```
show ip tcp statistics
```

6. View IPv4 udp endpoints

```
show ip udp endpoints
```

7. View IPv4 udp statistics

```
show ip udp statistics
```

Example

```
Switch:1#show ip tcp connections
```

```
=====
                        TCP connection table info
=====
```

LOCALPORT	LOCALADDR	REMOTEPORT	REMOTEADDR	STATE	VRF ID
21	0.0.0.0	0	0.0.0.0	listen	0
22	0.0.0.0	0	0.0.0.0	listen	0
23	0.0.0.0	0	0.0.0.0	listen	0
80	0.0.0.0	0	0.0.0.0	listen	0
443	0.0.0.0	0	0.0.0.0	listen	0
23	192.0.2.146	52583	198.51.100.30	established	0

```
Switch:1#show ip tcp properties
```

```
show ip tcp global properties command:
```

```
-----
RtoAlgorithm      constant
RtoMin            5002 milliseconds
RtoMax            60128 milliseconds
MaxConn           127
```

```
Switch:1>show ip tcp statistics
```

```
show ip tcp global statistics:
```

```
-----
```

```

ActiveOpens:      0
PassiveOpens:    37
AttemptFails:    0
EstabResets:     34
CurrEstab:       1
InSegs:          6726
OutSegs:         7267
RetransSegs:     10
InErrs:          0
OutRsts:         10

Switch:1#show ip udp endpoints

```

```

=====
                        UDP endpoint table info
=====
-----
LOCALPORT  LOCALADDRESS  REMOTEPORT  REMOTEADDRESS  PROCESS  INSTANCE
-----
26          0.0.0.0 0      0.0.0.0 1          0
67          0.0.0.0 0      0.0.0.0 1          0
69          0.0.0.0 0      0.0.0.0 1          0
161         0.0.0.0 0      0.0.0.0 1          0

```

```
Switch:1#show ip udp statistics
```

```
show ip udp info:
```

```

-----
InDatagrams:      887
NoPorts:          0
InErrors:          0
OutDatagrams:     887
HCInDatagrams:    887
HCOutDatagrams:   887

```

Variable Definitions

Use the data in the following table to use the **show ip tcp** command.

Variable	Value
connections	<p>Specifies the TCP connection information. Use the following parameters:</p> <ul style="list-style-type: none"> vrf WORD<0–16> Specifies a virtual routing and forwarding (VRF) by name. vrfids WORD<0–512> Specifies the IDs of a VRF path as an integer from 1 to 512. <p>Example: show ip tcp connections vrf 0</p>
properties	Specifies the TCP global properties information.
statistics	Specifies the TCP global statistics.

Use the data in the following table to use the **show ip udp** command.

Variable	Value
endpoints	Specifies the IP UDP endpoint information.
statistics	Specifies IP UDP statistics information.

IP routing configuration using Enterprise Device Manager

Configure the IP router interface so that you can use routing protocols and features on the interface. This section contains instructions for both the Global Router and Virtual Router Forwarding (VRF) instances.

Enable Routing for a Router or a VRF Instance

Enable IP forwarding (routing) on a router or a Virtual Router Forwarding (VRF) instance so that they support routing. You can use the IP address of any physical or virtual router interface for an IP-based network management.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

You cannot configure this option on a DvR Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.
4. To enable routing, select **forwarding**.
5. Select **Apply**.

Delete a Dynamically-Learned Route

Use the **Routes** tab to view and manage the contents of the system routing table. You can also delete a dynamically learned route using this table. Exercise caution if you delete entries from the route table.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

To delete a static route, use the **StaticRoute** tab.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Routes** tab.
4. To delete a route, select the route, and then select **Delete**.

Configure IP Route Preferences

Change IP route preferences to force the routing protocols to prefer a route over another. Configure IP route preferences to override default route preferences and give preference to routes learned for a specific protocol.

Before You Begin

- Disable ECMP before you configure route preferences.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task



Important

Changing route preferences is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. Therefore, if you want to change default preferences for routing protocols, do so before you enable the protocols.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **RoutePref** tab.
4. For **ConfiguredValue**, change the preference for the given protocol.
5. Select **Apply**.

RoutePref Field Descriptions

Use the data in the following table to use the **RoutePref** tab.

Name	Description
DefaultValue	Specifies the default preference value for the specified protocol.
Protocol	Specifies the protocol name.
ConfiguredValue	Specifies the preference value for the specified protocol.

Assign an IP Address to a Port

Assign an IP address to a port so that it acts as a routable VLAN (a brouter port) and supports IP routing.

Before You Begin

- Ensure routing (forwarding) is globally enabled.
- Ensure the VLAN is configured.
- If required, ensure the VRF instance exists.

About This Task

After you configure the IP address, you cannot edit the IP address, and you can assign only one IP address to a brouter port.

You cannot assign an IP address to a brouter port that is a member of a routed VLAN. To assign an IP address to the brouter port, you must first remove the port from the routed VLAN.

If you want to assign a new IP address to a VLAN or brouter port that already has an IP address, first delete the existing IP address and then insert the new IP address.

The ability to create an IP interface in a disabled state in one step is available only using CLI.

Procedure

1. In Device Physical View, select the port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **IP Address** tab.
5. Select **Insert**.
6. In the **Insert IP Address** dialog box, type the IP address, network mask, and VLAN ID.
7. Select **Insert**.

IP Address *Field Descriptions*

Use the data in the following table to help use the **IP Address** tab.

Name	Description
Interface	Specifies the router interface. <ul style="list-style-type: none"> The name of the VLAN followed by the VLAN designation identifies virtual router interfaces. The slot and port number of the brouter port identifies brouter interfaces.
Ip Address	Specifies the IP address of the brouter interface on this port. You can define only one IP address on a given port interface.
Net Mask	Specifies the subnet mask of the brouter interface on this port. The mask is an IP address with all the network bits set to 1 and all the hosts bits set to 0.
Name	Specifies the name of the brouter interface. The value ranges from 0 to 64 characters.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface.
ReasmMaxSize	Specifies the size of the largest IP packet which the interface can reassemble from fragmented incoming IP packets.
VlanId	Specifies the ID of the VLAN associated with the brouter port. This parameter is used to tag ports.
BrouterPort	Indicates whether this is a brouter port.
MacOffset	Specifies a number by which to offset the MAC address of the VLAN from the chassis MAC address. This ensures that each IP address has a different MAC address.
VrfId	Specifies the associated VRF interface. The VrfId associates VLANs or brouter ports to a VRF after the creation of VLANs or brouter ports. VRF ID 0 is reserved for the Global Router.

View Primary IP Addresses

Use this procedure to view primary IP addresses, and their associated router interfaces, from one central location.

About This Task

You can also use this tab to configure a name or MAC offset for the interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Primary Addresses** tab.

Primary Addresses *Field Descriptions*

Use the data in the following table to use the **Primary Addresses** tab.

Name	Description
Interface	Displays the router interface.
Ip Address	Displays the IP address of the router interface.
Net Mask	Displays the subnet mask of the router interface.
Name	Specifies the name assigned to the interface.
BcastAddrFormat	Displays the IP broadcast address format used on this interface; that is, whether 0 (zero) or one is used for the broadcast address. The switch uses 1.
ReasmMaxSize	Displays the size of the largest IP packet that this interface can reassemble from incoming fragmented IP packets.
VlanId	Displays the VLAN associated with this entry. This value corresponds to the lower 12 bits in the IEEE 802.1Q VLAN tag.
BrouterPort	Displays if the interface is a brouter port, as opposed to a routable VLAN.
MacOffset	Specifies a number by which to offset the MAC address of the VLAN from the chassis MAC address. This offset ensures that each IP address has a different MAC address.
IfType	Specifies the type of interface associated with the address.
VrfId	Displays the VRF ID, if one is associated with the address.

View all IP Addresses

Use this procedure to view primary or secondary IP addresses on all interface types from one EDM tab.

About This Task

You can create a Secondary IP Interface only on a VLAN interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **All Addresses** tab.

All Addresses *Field Descriptions*

Use the data in the following table to use the **All Addresses** tab.

Name	Description
Interface	Identifies the IP interface.
Type	Displays the address type, either primary or secondary.
Ip Address	Displays the IP address.
Net Mask	Displays the subnet mask for the address.
VlanId	Displays the VLAN ID, if applicable.
VrfId	Displays the associated VRF ID, if applicable.

Modify the Administrative State of an IP Interface on a Brouter Port

By default, the switch enables an IP interface after you configure the primary IP address. You can administratively disable, and then reenabling, an IPv4 interface after you create it. For example, if you want to change the routing protocol configuration on the interface.

About This Task

This procedure applies to the IP Layer and functionality for the brouter port.

The ability to create an IP interface in a disabled state in one step is available only using CLI.

Procedure

1. On the Device Physical View tab, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **IP Interface** tab.
5. Change the **AdminStatus** value to change the administrative state.
6. Select **Apply**.

IP Interface *Field Descriptions*

Use the data in the following table to use the **IP Interface** tab.

Name	Description
AdminStatus	Modifies the administrative state of the IP Interface. By default, the switch enables an IP interface after you configure the primary IP address.
OperState	Displays the operational state of the IP Interface.

Configure IP Routing Features Globally

Configure the IP routing protocol stack to determine which routing features the switch can use.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.
4. To globally enable routing, select **forwarding**.
5. To globally configure the default TTL parameter type a value in the **DefaultTTL** field.
This value is placed into routed packets that have no TTL specified.
6. To globally enable IPv4 ICMP broadcast, select **IcmpEchoBroadcastRequestEnable**.
7. To globally enable IPv4 Fragmented ICMP packet filtering, select **IcmpDropFragmentsEnable**.
8. To globally enable the Alternative Route feature, select **AlternativeEnable**.
9. To globally enable ICMP Router Discovery, select **RouteDiscoveryEnable**.
10. To globally enable IP Source Routing, select **SourceRouteEnable**.
11. To globally enable ECMP, select **EcmpEnable**.
12. Configure the remaining parameters as required.
13. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Forwarding	Configures the system for forwarding (routing) or for dropping. The default value is forwarding. You cannot configure this option on a DvR Leaf node.
DefaultTTL	Configures the default time-to-live (TTL) value for a routed packet. TTL indicates the maximum number of seconds elapsed before a packet is discarded. Enter an integer from 1 to 255. The default value of 255 is used if a value is not supplied in the datagram header.

Name	Description
ReasmTimeout	Specifies the maximum number of seconds that received fragments are held while they wait for reassembly. The default value is 30 seconds.
ICMPUnreachableMsgEnable	<p>Enables the generation of Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this system. These messages help determine if the system is reachable over the network. The default is disabled.</p> <p>Important: As a best practice, enable icmp-unreach-msg only if it is absolutely required. If icmp-unreach-msg is enabled and a packet is received for which there is no route in the routing table, CPU utilization can dramatically increase.</p>
ICMPRedirectMsgEnable	Enables or disables the system sending ICMP destination redirect messages.
IcmpEchoBroadcastRequestEnable	Enables or disables IP ICMP echo broadcast request feature. The default is enabled.
IcmpDropFragmentsEnable	Enables or disables IPv4 Fragmented ICMP packet filtering globally. The default is disabled.
AlternativeEnable	<p>Globally enables or disables the Alternative Route feature.</p> <p>If the alternative-route parameter is disabled, all existing alternative routes are removed. After the parameter is enabled, all alternative routes are re-added. The default is enabled.</p>
RouteDiscoveryEnable	Enables the ICMP Router Discovery feature. The default is disabled (not selected). Use ICMP Router Discovery to enable hosts attached to multicast or broadcast networks to discover the IP addresses of neighboring routers.
AllowMoreSpecificNonLocalRouteEnable	Enables or disables a more-specific nonlocal route. If enabled, the system can enter a more-specific nonlocal route into the routing table. The default is disabled.

Name	Description
SuperNetEnable	Enables or disables supernetting. If supernetting is globally enabled, the system can learn routes with a route mask less than 8 bits. Routes with a mask length less than 8 bits cannot have ECMP paths, even if you globally enable the ECMP feature. The default is disabled.
UdpChecksumEnable	Enables or disables the UDP checksum calculation. The default is enable.
SourceRouteEnable	Enables or disables IP Source Routing globally. It is disabled by default.
SpbMulticastPolicyEnable	Enables the IP SPB Multicast Policy .
SpbMulticastPolicyRmap	Configures the name of the route-map for IP SPB Multicast Policy .
SpbMulticastPolicyApply	Applies the IP SPB Multicast Policy . The default value is noAction.
ARPLifeTime	Specifies the lifetime of an ARP entry within the system, global to the switch. The default value is 360 minutes.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled. After ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.
EcmpMaxPath	Globally configures the maximum number of ECMP paths. You cannot configure this feature unless ECMP is enabled globally.
Ecmp1PathList	Selects a preconfigured ECMP path.
Ecmp2PathList	Selects a preconfigured ECMP path.
Ecmp3PathList	Selects a preconfigured ECMP path.
Ecmp4PathList	Selects a preconfigured ECMP path.
Ecmp5PathList	Selects a preconfigured ECMP path.
Ecmp6PathList	Selects a preconfigured ECMP path.
Ecmp7PathList	Selects a preconfigured ECMP path.
Ecmp8PathList	Selects a preconfigured ECMP path.
EcmpPathListApply	Applies changes in the ECMP pathlist configuration, or in the prefix lists configured as the pathlists.

Name	Description
TcpAdjustMssEnable	Adjusts the TCP maximum segment size (MSS) to improve the throughput for the TCP session over a Fabric Extend (FE) adjacency. The default is disabled.
TcpAdjustMssStatus	Displays the activation status of the MSS adjustment functionality.
TcpAdjustMssType	Displays if the MSS adjustment value is manually configured or auto-derived. The software does not support auto-derived values for this feature.
TcpAdjustMssValue	Configures the MSS adjustment value. • The default value is 1300.

Configure ECMP

Enable Equal Cost MultiPath (ECMP) to permit routers to determine up to eight equal-cost paths to the same destination prefix. You can use the multiple paths for load-sharing of traffic, which allows fast convergence to alternative paths. By maximizing load sharing among equal-cost paths, you can maximize the efficiency of links between routers.

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- To configure an ECMP pathlist, you must first configure a prefix list that you reference in the pathlist configuration.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.
4. Select **EcmpEnable**.
5. (Optional) In **EcmpMaxPath**, type the preferred maximum number of equal-cost paths.
6. (Optional) Configure an ECMP pathlist to specify routes with the required number of paths.
7. (Optional) If you modified the ECMP pathlist configuration, select **EcmpPathListApply**.

8. Select **Apply**.

Configure Static Routes

Perform the steps in this task to complete one of the following configurations:

- Create static routes for data traffic in either the GRT or a specific VRF context for any platform.
- Create static routes for a VRF associated with a Segmented Management Instance CLIP or VLAN (routing-mode) interface.
- Configure a discard static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.
- Configure a default route to specify a route to all networks for which there are no explicit routes in the forwarding information base or the routing table. This route has a prefix length of zero (RFC 1812). You can configure the switch with a static default route, or they can learn it through a dynamic routing protocol. To create a default static route, you configure the destination address and subnet mask to 0.0.0.0.

Static routes for the management OOB, if supported, and management VLAN (non-routing mode) must use the Segmented Management Instance. The management CLIP can use the Segmented Management Instance or routes in the associated VRF routing table manager (RTM). The management VLAN (routing mode) uses the Fabric Engine routing table or routes in the associated VRF.

Before You Begin

- If a discard static route is enabled, you must first delete or disable it before you can add a regular static route to that destination.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task



Note

Only discard routes that belong to the static type protocol are supported. An inter-VRF discard route is not installed in the routing table of the destination VRF on the same switch.

When you configure a static route with a next-hop-vrf context, you can specify a next-hop IP address that is a locally owned VRRP IP address of the system itself. However, this is not a supported configuration. The best practice is to implement an alternative method of inter-vrf route sharing, such as route redistribution or IS-IS accept polices.

For route scaling information, see *Fabric Engine Release Notes*.



Note

Do not configure static routes on a DVR Leaf node unless the configuration is for reachability to a management network using a brouter port.

Also, configuring the preference of static routes is not supported on a Leaf node.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Static Routes** tab.
4. Select **Insert**.
5. (Optional) If required, in the **OwnerVrflid** check box, select the appropriate VRF ID. By default, the VRF is the GlobalRouter VRF 0.
6. In the **Dest** field, type the IP address.
For a default static route, type 0.0.0.0.
7. In the **Mask** field, type the subnet mask.
For a default static route, type 0.0.0.0.
8. In the **NextHop** field, type the IP address of the router through which the specified route is accessible.
9. (Optional) In the **NextHopVrflid** field, select the appropriate value.
10. Select **Enable** to enable the static route.
11. (Optional) In the **Metric** field, type the metric value.
12. (Optional) In the **Preference** field, type the route preference.
When you specify a route preference, be sure to appropriately configure the preference so that when a discard static route is used, it is elected as the best route.
13. (Optional) If required, select **LocalNextHop**.
Use this option to create Layer 3 static routes.
14. Select **Insert**.

Static Routes Field Descriptions

Use the data in the following table to use the **Static Routes** tab.

Name	Description
OwnerVrflid	Specifies the VRF ID for the static route.
Dest	Specifies the destination IP address of this route. A value of 0.0.0.0 is a default route. The system can display multiple routes to a single destination in the table, but access to such multiple entries is dependent on the table-access mechanisms defined by the network management protocol in use.

Name	Description
Mask	Specifies the mask that the system operates a logically AND function on, with the destination address, to compare the result to the Route Destination. For systems that do not support arbitrary subnet masks, an agent constructs the Route Mask by determining whether it belongs to a class A, B, or C network, and then uses one of: 255.0.0.0—Class A 255.255.0.0—Class B 255.255.255.0—Class C If the Route Destination is 0.0.0.0 (a default route) then the mask value is also 0.0.0.0.
NextHop	Specifies the IP address of the next hop of this route. In the case of a route bound to an interface which is realized through a broadcast media, the Next Hop is the IP address of the agent on that interface. If the IPv4 next hop is intended for a discard static route, the field autocompletes with the 255.255.255.255. This value cannot be modified. This address indicates the traffic to this route is discarded.
Discard	Enables or disables IPv4 discard static routes. The default is disabled.
NextHopVrflid	Specifies the next-hop VRF ID in interVRF static route configurations. Identifies the VRF in which the ARP entry resides.
Name	Specifies the name for the static route.
Enable	Determines whether the static route is available on the port. The default is enable. If a static route is disabled, it must be enabled before it can be added to the system routing table.
Status	Specifies the status of the route. The default is enabled.
Metric	Displays the primary routing metric for this route. The semantics of this metric are determined by the routing protocol specified in the route RouteProto value. Note: To change the metric for an existing static route, remove the route and recreate it with the proper metric.
IfIndex	Specifies the route index of the Next Hop. The interface index identifies the local interface through which the next hop of this route is reached.
Preference	Specifies the routing preference of the destination IP address. If more than one route can be used to forward IP traffic, the route that has the highest preference is used. The higher the number, the higher the preference.

Name	Description
LocalNextHop	Enables and disables LocalNextHop. If enabled, the static route becomes active only if the system has a local route to the network. If disabled, the static route becomes active if the system has a local route or a dynamic route.
Origin	Specifies the origin of the static route configuration, either manually configured through CLI or EDM, or through Zero Touch Fabric Configuration.

Delete a Static Route

Delete static routes that are no longer needed to prevent routing errors.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Static Routes** tab.
4. Select the static route you want to delete.
5. Select **Delete**.

View IP Routes

View IP routes learned on the device.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Routes** tab to view IP routes learned on the device.
4. If you want to limit the routes displayed, select **Filter** to show a smaller subset of the learned routes.
5. In the Filter dialog box, select an option, or options, and enter information to limit the routes to display in the Routes table.
6. Select **Filter** to display only the routes that match the options and information that you enter.

Routes Field Descriptions

Use the data in the following table to use the **Routes** tab.

Name	Description
Dest	Specifies the destination IP network of this route. An entry with a value of 0.0.0.0 is a default route. The system can display multiple routes to a single destination in the table, but access to multiple entries depends on the table access mechanisms defined by the network management protocol in use.
Mask	Indicates the network mask to logically add with the destination address before comparison to the destination IP network.
NextHop	Specifies the IP address of the next hop of this route.
AltSequence	Indicates the alternative route sequence. The value of 0 denotes the best route.
NextHopId	Displays the MAC address or hostname of the next hop.
HopOrMetric	Displays the primary routing metric for this route. The semantics of this metric are specific to different routing protocols.
Interface	Specifies the router interface for this route. <ul style="list-style-type: none"> Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation. Brouter interfaces are identified by the slot and port number of the brouter port.
Proto	Specifies the routing mechanism through which this route was learned: <ul style="list-style-type: none"> other—none of the following local—nonprotocol information, for example, manually configured entries static ICMP EGP GGP Hello RIP IS-IS ES-IS Cisco IGRP bbnSpfIgp OSPF BGP Inter-VRF Redistributed Route
Age	Displays the number of seconds since this route was last updated or otherwise determined to be correct.

Name	Description
PathType	<p>Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths.</p> <ul style="list-style-type: none"> • iA indicates Indirect Alternative route without an ECMP path • iAE indicates Indirect Alternative ECMP path • iB indicates Indirect Best route without ECMP path • iBE indicates Indirect Best ECMP path • dB indicates Direct Best route • iAN indicates Indirect Alternative route not in hardware • iAEN indicates Indirect Alternative ECMP route not in hardware • iBN indicates Indirect Best route not in hardware • iBEN indicates Indirect Best ECMP route not in hardware • dBN indicates Direct Best route not in hardware • iAU indicates Indirect Alternative Route Unresolved • iAEU indicates Indirect Alternative ECMP Unresolved • iBU indicates Indirect Best Route Unresolved • iBEU indicates Indirect Best ECMP Unresolved • dBU indicates Direct Best Route Unresolved • iBF indicates Indirect Best route replaced by FTN • iBEF indicates Indirect Best ECMP route replaced by FTN • iBV indicates Indirect best IPVPN route • iBEV indicates Indirect best ECMP IP VPN route • iBVN indicates Indirect best IP VPN route not in hardware • iBEVN indicates Indirect best ECMP IP VPN route not in hardware
Pref	Displays the preference.
NextHopVrflid	Specifies the VRF ID of the next-hop address.
AccessType	<p>Displays the VRF segment name for entries within a Segmented VRF, based on the access level of the entry.</p> <p>EDM only displays this field if you launch the VRF context view.</p>

Configure ICMP Router Discovery Globally

About This Task

Configure ICMP Router Discovery so that it can operate on the system.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Globals** tab.

4. Select **RouteDiscoveryEnable**.
5. To select a preconfigured ECMP path, select the **EcmpPathList** ellipsis button.
6. Select **OK**.
7. Select **Apply**.

Configure the ICMP Router Discovery Table

Before You Begin

- ICMP Router Discovery must be globally enabled.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100.

About This Task

Configure the ICMP Router Discovery table to ensure correct ICMP operation for all interfaces that use Router Discovery.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Router Discovery** tab.
4. Configure the Router Discovery parameters to suit the network.
5. Select **Apply**.

Router Discovery Field Descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
Interface	Indicates the VLAN ID or the port.
AdvAddress	Specifies the IP destination address used for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address 224.0.0.1, or the limited-broadcast address 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface. The default value is true (advertise address).
AdvLifetime	Specifies the time to-live-value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds. The default value is 1800 seconds.

Name	Description
MaxAdvInterval	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 4 to 1800 seconds. The default value is 600 seconds.
MinAdvInterval	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet. The range is -2147483648 to 2147483647. The default value is 0.

Configure ICMP Router Discovery for a Port

Before You Begin

- You must globally enable ICMP Router Discovery.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100.

About This Task

When you enable ICMP Router Discovery, the port sends Router Discovery advertisement packets.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **Router Discovery** tab.
5. To enable Router Discovery, select **AdvFlag**.
6. Configure other parameters as required for proper operation.
7. Select **Apply**.

Router Discovery *Field Descriptions*

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
AdvAddress	Specifies the destination IP address used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address 224.0.0.1, or the limited-broadcast address 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface. The default value is True (advertise address).
AdvLifetime	Specifies the time to live value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 4 seconds to 1800 seconds. The default value is 600 seconds.
MinAdvInterval	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet. The accepted values are -2147483648 to 2147483647. The default value is 0.

Configure ICMP Router Discovery on a VLAN

Before You Begin

- You must globally enable ICMP Router Discovery.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100.

About This Task

When you enable ICMP Router Discovery, the system sends Router Discovery advertisement packets to the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.

2. Select **VLANs**.
3. Select the VLAN ID that you want to configure to participate in Router Discovery.
4. Select **IP**.
5. Select the **Router Discovery** tab.
6. To enable Router Discovery for the VLAN, select **AdvFlag**.
7. Configure other parameters as required for proper operation.
8. Select **Apply**.

Router Discovery *Field Descriptions*

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
AdvAddress	Specifies the destination IP address used for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface. The default value is true (advertise address).
AdvLifetime	Specifies the time to-live-value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	Specifies the maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The range is 4 seconds to 1800 seconds. The default value is 600 seconds.
MinAdvInterval	The minimum time (in seconds) allowed between unsolicited broadcast or multicast router advertisements sent from the interface. The range is 3 seconds to MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The range is -2147483648 to 2147483647. The default value is 0.

Configure a Circuitless IPv4 Interface

About This Task

You can use a circuitless IPv4 (CLIPv4) interface to provide uninterrupted connectivity to your system.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select **Insert**.
5. In **Interface**, assign a CLIP interface number.
6. Type the IP address.
7. Type the network mask.
8. Select **Insert**.
9. Select **Apply**.

Circuitless IP Field Descriptions

Use the data in the following table to use the **Circuitless IP** tab.

Name	Description
Interface	Specifies the number assigned to the interface. You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.
Ip Address	Specifies the IP address of the CLIP. You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.
Net Mask	Specifies the network mask. You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.
Name	Specifies a name assigned to the IPv4 CLIP address. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 5720 Series, 7520 Series, and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link</code> .

Name	Description
IfType	<p>Specifies the interface type.</p> <p>Note: Exception: circuitlessIPMAVirtualLinkHome and circuitlessIPMAVirtualLinkRemote apply to 5720 Series, 7520 Series, and 7720 Series only.</p> <p>You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.</p>
AccessType Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	<p>Specifies to which segment of a Segmented VRF the interface belongs. If you use the <i>trusted</i> parameter, the interface is part of the trusted segment and is only accessible from hosts in the trusted and unrestricted segments. Use the <i>unrestricted</i> parameter to make the interface accessible from hosts in all three segments. The default is trusted.</p> <p>EDM only displays this field if you launch the VRF context view. This field does not apply to the GlobalRouter.</p> <p>You can only configure this parameter through the Insert Circuitless IP dialog. The value on the tab is read-only.</p>

View TCP Global Information

About This Task

Perform this procedure to view information about the handshake (SYN) configuration and the maximum number of TCP connections you can create on your system. When you initiate a TCP connection, both end points send handshake information to create the channel. The retransmission algorithm and fields display the configured timeout value and minimum and maximum retransmission times that your system uses to terminate a connection attempt that falls outside your specified parameters.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **TCP/UDP**.
3. Select the **TCP Globals** tab.

TCP Global *Field Descriptions*

Use the data in the following table to use the **TCP Global** tab.

Name	Description
RtoAlgorithm	Determines the timeout value used for retransmitting unacknowledged octets.
RtoMin	Displays the minimum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
RtoMax	Displays the maximum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
MaxConn	Displays the maximum connections for the device.

View TCP Connections Information

About This Task

Perform this procedure to view the status of connections that traverse your switch.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **TCP/UDP**.
3. Select the **TCP Connections** tab.

TCP Connections *Field Descriptions*

Use the data in the following table to use the **TCP Connections** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
RemAddressType	Displays the type (IPv6 or IPv4) for the remote address of the TCP connection.
RemAddress	Displays the IPv6 address for the remote TCP connection.
RemPort	Displays the remote port number for the TCP connection.

Name	Description
State	Displays an integer that represents the state for the connection: <ul style="list-style-type: none">• closed• listen• synSent• synReceived• established• finWait1• finWait2• closeWait• lastAck(9)• closing• timeWait• deleteTCB
Process	Displays the process ID for the system process associated with the TCP connection.

View TCP Listeners Information

About This Task

Perform this procedure to view a detailed list of systems that are in the listening state.

When a connection is in the listen state an end point system is waiting for a returning handshake (SYN). The normal listening state should be very transient, changing all of the time. Two or more systems going to a common system in an extended listening state indicates the need for further investigation. End systems in an extended listening state can indicate a broken TCP connection or a DOS attack on a resource. This type of DOS attack, known as a SYN attack, results from the transmission of SYNs with no response to return replies. While many systems can detect a SYN attack, the TCP listener statistics can provide additional forensic information.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **TCP/UDP**.
3. Select the **TCP Listeners** tab.

TCP Listeners *Field Descriptions*

Use the data in the following table to use the **TCP Listeners** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
Process	Displays the process ID for the system process associated with the TCP connection.



IPv6 Routing Basics

[Origins of IPv6](#) on page 2168
[Advantages of IPv6](#) on page 2168
[Comparison of IPv4 and IPv6](#) on page 2169
[IPv6 packet](#) on page 2169
[IPv6 header](#) on page 2170
[IPv6 extension headers](#) on page 2170
[IPv6 address component summary](#) on page 2172
[IPv6 address formats](#) on page 2173
[Address types](#) on page 2173
[IP address prefix](#) on page 2178
[Interface ID](#) on page 2178
[How to write an IPv6 address](#) on page 2179
[ICMPv6](#) on page 2180
[Path MTU discovery](#) on page 2181
[Routing](#) on page 2181
[IPv6 Discard Static Routes](#) on page 2187
[Route scaling](#) on page 2187
[IPv6 Circuitless IP](#) on page 2187
[Equal Cost Multipath](#) on page 2189
[ECMP with static routes](#) on page 2189
[Disable IPv6 ICMP multicast](#) on page 2190
[View IPv6 Connections](#) on page 2190
[IPv6 Basic Configuration using CLI](#) on page 2191
[Viewing IPv6 Connections Using CLI](#) on page 2221
[IPv6 Basic Configuration using EDM](#) on page 2223
[Viewing IPv6 Connections using EDM](#) on page 2255

The following sections provide concepts and procedures to complete basic IPv6 configuration, for example, IPv6 forwarding and static routes.

4220 Series supports IPv6 for management only.

For the 5320 Series, only the following models support more than one VRF with IP configuration:

- 5320-16P-2MXT-2X
- 5320-24T-4X-XT

- 5320-24T-24S-4XE-XT
- 5320-48P-8XE
- 5320-48T-8XE

Origins of IPv6

The growth of IP address use is exponential.

Predictions indicated that the IPv4 address pool could be exhausted as early as 1994.

So, in July 1991, the Internet Engineering Task Force (IETF) began researching a replacement for IPv4.

That replacement is IPv6.

The Internet Assigned Numbers Authority (IANA) free pool of IPv4 addresses reached 0% in February 2011, according to the American Registry for Internet Numbers (ARIN).

While IPv4 addresses may remain available for some time within reserved pools, no further IPv4 addresses are available for reservation.

Although IPv6 is designed to replace IPv4, IPv6 is not backward-compatible and IPv4 and IPv6 need to coexist within your network during and after the transition to IPv6.

Advantages of IPv6

IPv6 can provide more addresses and support more networks than IPv4. For example, IPv6 offers enough addresses for every person on Earth to have 1 million addresses.

Because IPv6 offers a larger address space it offers improved scalability.

Following are additional advantages of IPv6 over IPv4:

- With 128 bit addresses, the larger IPv6 address space offers global access and scalability and solves the pending exhaustion of IP addresses.
- Network Address Translation (NAT) is no longer required.

Flat address space and transparency are restored by IPv6 because NAT is eliminated.

- Routing efficiency is improved due to the hierarchical network architecture.

IPv6 allows for hierarchical routing and effective route summarization.

- IPv6 supports Auto-configuration.
- IPv6 supports plug-and-play.
- Enhanced support is included for mobile IP and mobile computing devices.

Addresses can be permanently assigned to end devices such as DSL, PDAs, mobile terminals and PCs.

- Neighbor discovery (ND) replaces ARP in IPv6.

ND combines the IPv4 services for IPv4 Address Resolution Protocol (ARP) and router discovery.

Comparison of IPv4 and IPv6

The following table compares the key differences between IPv4 and IPv6.



Note

This information may not reflect IPv6 support in the current release.

Table 162: IPv4 and IPv6 key differences compared

Feature	IPv4	IPv6
Address length	32 bits	128 bits
IPsec support	Optional	Required
QoS support	Limited	Improved
Fragmentation	Hosts and routers	Hosts only
Minimum MTU packet size	576 bytes	1280 bytes
Checksum in header	Yes	No
Options in header	Yes	Yes
Link-layer address resolution	ARP (broadcast)	Multicast neighbor discovery messages
Multicast membership	IGMP	Multicast Listener Discovery
Router discovery	Optional	Optional
Uses broadcasts	Yes	No
Address configuration	Manual, DHCP	Automatic, DHCP

IPv6 packet

Each IPv6 packet can include mandatory and non-mandatory components.

An IPv6 packet includes:

- The basic header, which has a fixed length and is mandatory
- Extension header(s) , which has a variable length and is not mandatory
- Payload, which has a variable length and is not mandatory

The following figure illustrates the components of an IPv6 packet.




Figure 155: IPv6 packet components



Note

Nodes must be able to handle packets up to 1,280 octets in length.

IPv6 header

The IPv6 header basic length is fixed at 40 octets (bytes) and it contains the following fields:

Table 163: Fields in the IPv6 header

Field	Size in bits
Ver—Internet Protocol version number, with a value of 6	4
DS byte—Traffic class field, similar to Type of Service in IPv4	8
Flow label—identifies traffic flow for additional Quality of Service (QoS)	20
Payload Length—Unsigned integer, the length of the IPv6 payload	16
Next header selector—identifies the next header	8
Hop limit unsigned integer—decrements by 1 each time a node forwards the packet (nodes discard packets with hop limit values of 0)	8
Source address	128
Destination address	128

The following figure illustrates the basic IPv6 header, without extension headers.




Figure 156: IPv6 header

IPv6 extension headers

IPv6 extension headers describe processing options.

Each extension header contains a separate category of options and is identified by a number, similar to protocol identification numbers.

An IPv6 packet can include extension headers, but they are not mandatory.

The following figure illustrates the IPv6 header with extension headers.




Figure 157: IPv6 header with extension headers

IPv6 examines the destination address in the main header of each packet it receives.

This examination determines whether the router is

- the packet destination - if the router is the packet destination, IPv6 examines the header extensions that contain options for destination processing.
- an intermediate node in the packet data path - if the router is an intermediate node, IPv6 examines the header extensions that contain forwarding options.

By examining only the extension headers that apply to the operations it performs, IPv6 reduces the amount of time and resources required to process a packet.

IPv6 defines the following extension headers as described in the following table:

Table 164: IPv6 extension headers

Extension header name	Description
hop-by-hop	Contains optional information, and sub-options for Router Alert and Jumbo Payload, that all intermediate IPv6 routers examine between the source and the destination.
destinations-options	Contains optional information for the destination node. The system can display this option twice, once for way points and once for final destination.
source-routing	Contains a list of one or more intermediate nodes that define a path for the packet to follow through the network to the destination. The packet source creates this list. The source-routing function is similar to the IPv6 source routing options.
fragmentation	Uses an IPv6 source to send packets larger than the size specified for the path maximum transmission unit (MTU).
authentication	provides security for IPv6 datagrams
encapsulated security payload (ESP)	provides security for IPv6 datagrams
The authentication extension header and the encapsulated security payload extension header can be used together to provide security services for IPv6 datagrams.	

The following extension header order is:

- Hop-by-hop
- Destination option 1
- Routing
- Fragmentation
- Authentication/ESP
- Destination Option 2

The presence of particular extension headers within a packet can cause slower packet processing if the IPv6 implementation handles only certain headers and diverts others to a slow path. For example, many IPv6 implementations usually process Hop-by-Hop extension headers on the control plane.

IPv6 address component summary

The IPv6 Internet is divided into addressing zones and IPv6 addresses can be categorized by type and scope.

IPv6 addressing is represented in RFC 4291.

Address types

IPv6 addresses are divided into the following types:

- Unicast

Unicast addresses provide one-to-one communication.

- Multicast

Multicast addresses are similar in operation to IPv4 and provide one-to-many communication.

- Anycast

An Anycast address is a Unicast address used for several devices to allow them to communicate with the device closest to the source; one-to-nearest communication.

- Broadcast

In IPv6, broadcast addresses have been superseded by multicast addresses per RFC 4291.

For more information about address types and scopes, see [Address types](#) on page 2173.

Address scopes

Following are IPv6 address scopes:

- node-local
- link-local
- global

The switch does not support site-local addresses and, according to RFC 4193, site-local addresses will be replaced by unique-local addresses.

For more information about address types and scopes, see [IPv6 address formats](#) on page 2173.

Address zones

The IPv6 Internet is divided into zones.

For example:

- Each node is a separate zone of the node-local scope.
- Each link is a separate zone of the link-local scope.
- The entire Internet is a single zone of global scope.

Zones of the same scope do not overlap.

IPv6 address formats

IPv6 addresses are 128 bits long. In comparison, IPv4 addresses are 32 bits in length.

The IPv6 address contains an

- address type
- address prefix
- interface ID

The following figure illustrates the IPv6 address format.

Type	Address prefix	Interface ID
------	----------------	--------------

Figure 158: IPv6 address format

Address types

IPv6 uses three main address types to help route packets.

Address types are:

- Unicast: global, link—local, special unspecified, special loopback
- Multicast
- Anycast

Difference between multicast and anycast

Anycast address delivery is from one to one-of-many, whereas multicast address delivery is from one to many.

Unicast addresses

Unicast addresses provide one-to-one communication.

Global

A Unicast global address identifies a single interface and is similar to an IPv4 public address.

Unicast global addresses are globally routable in the same manner as IPv4 addresses.

The following figure illustrates the Unicast global address parts.




Figure 159: Unicast global address parts

An IPv6 Unicast global address is composed of the following 3 levels:

- public topology (48 bit Global Routing Prefix)
 - 001, specifies an IPv6 Unicast global address
 - Top Level Aggregation Identifier (TLA ID), the highest level in routing hierarchy
 - Res, reserved for future use
 - Next Level Aggregation Identifier (NLA ID), specifies a customer site
- site topology (16 bit Subnet ID)
 - Site Level Aggregation Identifier (SLA ID); assigned within the site, an ISP cannot affect the SLA ID, enables up to 65,536 subnets within a site
- interface ID (64 bits)
 - specifies the interface for a node on a subnet

The system uses the 48 bit global routing prefix for the route prefix exchange.

The IPv6 Prefix for Unicast global is 2000::/3 (RFC3513).

Link-local

Hosts on the same link/subnet use automatically configured IPv6 Unicast link-local addresses to communicate with each other.

Link-local addresses are automatically configured on all interfaces.

Routers do not forward packets containing a destination or source address with a link-local address.

IPv6 uses neighbor discovery (ND) for address resolution.

The IPv6 prefix for link-local Unicast addresses is FE80::/10 (RFC3513).

The following figure illustrates the parts of a Unicast Link-local address.




Figure 160: Unicast Link-local address

Special addresses

The Unicast/special/unspecified address indicates the absence of an address and is the only valid SRC address for IPv6 Duplicate Address Detection (DAD).

Equivalent to the IPv4 unspecified address 0.0.0.0, represented as 0:0:0:0:0:0:0:0 or ::1; an IPv6 host that does not have a valid address uses the unspecified address as its source address when it sends a packet to discover whether an address is used by another node (during the boot process when the host requests address configuration information).



Note

Do not assign an unspecified address, either statically or dynamically, to an interface.

The Unicast/special/loopback address is a special case Unicast address only found inside a single node.

The switch does not support the loopback address.

Equivalent to the IPv4 loopback address 127.0.0.1, represented as 0:0:0:0:0:0:0:1 or ::1; a node uses a loopback address to send a packet to itself.

The loopback address is beneficial in troubleshooting and testing the IP stack because you can use it to send a packet to the protocol stack without sending it onto the subnet.



Note

Do not assign a loopback address, either statically or dynamically, to an interface.

Both Loopback and Unspecified addresses are not valid destination addresses.

An example of a unicast IPv6 address is 1080:0:0:0:8:8000:200C:417A

An example of a link-local Unicast IPv6 address is FE80::4445:4eff:fe54:1212

Multicast addresses

Multicast addresses provide one-to-many communication

Multicast addresses provide one-to-many communication.

An IPv6 multicast address identifies a group of nodes.

The scope is built into the multicast address structure.

The system uses a multicast address to send traffic to multiple destinations. In this situation traffic experiences less delay with a multicast address than it would with Unicast address.

The following figure shows the format of an IPv6 multicast address.

8 bits	4 bits	4 bits	112 bits
11111111	flags	scope	group ID

Figure 161: IPv6 multicast address format

A value of FF (11111111) in the 8 high-order bits of an IPv6 address indicates that the address is an IP multicast address.

The Multicast IPv6 Prefix is FF00::/8 (RFC3513).

Flags

The 4-bit flags field indicates whether the group is permanent or transient. The first 3 bits are reserved and the 4th bit represents the Transient flag. Currently only the Transient (T) flag is defined. A T flag set to 0 specifies a permanently assigned multicast address. A T flag set to 1 specifies a transient address.

Group ID

The 112 bit group ID identifies the multicast group.

An example of a multicast address is FF01:0:0:0:0:0:0:101

Scope field

The 4-bit scope field within the group ID specifies the multicast traffic scope.

Following is a list of the scope options that limit the scope of the multicast address:

- 1 - node-local
- 2 - link-local
- 3 - subnet local
- 4 - admin local
- 5 - site-local – not supported
- 8 - organization-local
- B - community-local
- E - global

Examples of multicast addresses

All-nodes addresses look like this:

FF01::1 (Node Local), FF02::1 (Link Local)

All-routers addresses look like this:

FF01::2 (Node Local), FF02::2 (Link Local)

A solicited node or host address looks like this:

FF02::1:FF1E:8329.

In this case the MAC is 00-02-B3-1E-83-29 and the IPv6 address is fe80::202:B3FF:FE1E:8329.

The following table lists some well-known multicast IPv6 addresses

Table 165: Well-known multicast IPv6 addresses

Name	Address
All Nodes	FF02:0:0:0:0:0:1
All Routers	FF02:0:0:0:0:0:2
OSPFv2	FF02:0:0:0:0:0:5
OSPFv2 Designated Routers	FF02:0:0:0:0:0:6
All PIM Routers	FF02:0:0:0:0:0:D
VRRP	FF02:0:0:0:0:0:12
All MLDv2-capable routers	FF02:0:0:0:0:0:16
All DHCP agents	FF02:0:0:0:0:0:2
Solicited Node address	FF02::1:FF00:0000/104

Anycast Addresses

Anycast addresses provide one-to-nearest (one to one-of-many) communication.

Anycast addresses provide one-to-nearest (one to one-of-many) communication.

An anycast address designates a set of interfaces that share an address.

A packet sent to an anycast address goes only to the nearest member of the group. Considering routing distance, the system delivers packets with anycast addresses only to the nearest member of a group of multiple interfaces.

Restrictions

An anycast address must not be:

- used as the source address in an IPv6 packet
- assigned to an IPv6 host (you can assign an anycast address to an IPv6 router)

Anycast address scopes

Anycast addresses have the following scopes:

- Link-local—the local link; nodes on the same subnet
- Global—IPv6 Internet addresses

Similar to anycast IPv4 addresses, IPv6 anycast addresses are more efficient. They use the unicast address space but identify multiple interfaces.

IPv6 delivers a packet bearing an anycast address to the nearest interface identified by the address.

Currently anycast addresses are assigned to routers and are used as destination addresses. Because packets bearing anycast addresses are delivered to the closest router, you can also access the closest name server or time server with an anycast address.

Visually there is no distinction between an anycast address and a unicast address.



Note

The switch supports only the subnet-router anycast address.

You cannot configure any specific anycast addresses beyond the automatic, generic subnet-router anycast address.

IP address prefix

Address prefixes represent one of the following:

- network identifier
- fixed address part

Examples of IP address prefixes

2001:10F2::/48 represents a summarized route prefix

2001:10F2:0:102F::/64 represents a subnet or link prefix

FF00::/8 represents Multicast IPv6

Interface ID

Interface identifiers identify interfaces on a link.

As long as the interfaces are attached to different subnets, you can use the same identifier on more than one interface on a single node.

The IPv6 interface ID is as unique as the MAC address.

The interface ID is derived by a formula that uses the link layer 48-bit MAC address. In most cases, the interface ID is a 64-bit interface ID that contains the 48-bit MAC

address. If you enter less than 64 bits, the system adds leading zeroes to extend the interface ID length to 64 bits.

You can configure the interface ID in the following ways:

- Manual configuration
- DHCPv6 (can configure the whole address)
- Automatic derivation from EUI-64 (MAC address or other HW serial)—enables serverless or stateless auto-configuration when combined with high order part of address learned from router advertisements
- Pseudo-random generation (client privacy)—enables serverless or stateless auto-configuration when combined with high order part of address learned from router advertisements

The switch supports manual interface ID configuration or automatic derivation from EUI-64.

**Note**

You must manually specify the network prefix, regardless of the interface ID formation method.

For stateless autoconfiguration, the ID is 64 bits in length.

For more information about stateless autoconfiguration, see [Host autoconfiguration](#) on page 2265.

How to write an IPv6 address

The appearance of IPv6 addresses differs from IPv4 addresses and you express them differently.

Hexadecimal IPv6 address representations

The 128 bits in an IPv6 address are divided into 8 blocks of 16 bits each.

Following is the preferred IPv6 address format:

XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

Each 16 bit block in an IPv6 address is converted into a 1 to 4 digit hexadecimal number separated by colons (:).

The format to represent an IPv6 address is n:n:n:n:n:n:n, where n is the hexadecimal representation of 16 bits in the address; for example, 2001:0:0:0:0:0:0:43.

Each nonzero field must contain at least one numeral.

Within a hexadecimal field, you do not need leading zeros.

Certain classes of IPv6 addresses commonly include multiple adjacent fields that contain hexadecimal 0.

The sample address—2001::43—includes six adjacent fields that contain zeroes represented by a double colon (::) .

You can use a double colon to compress the leading zero fields in a hexadecimal address.

The system can display a double colon only once in an address.

Four more ways to write an IPv6 address

```
2001:DB8:0000:0000:25AB:0000:0000:0001
2001:DB8:0:0:25AB:0:0:1
2001:DB8:0:0:25AB::1
2001:DB8::25AB:0:0:1
```

ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) maintains and improves on features from ICMP for IPv4.

ICMPv6 reports the delivery of forwarding errors.

For example:

- Destination unreachable
- Packet too big (path MTU)
- Time exceeded (fragmentation)
- Parameter problem

ICMPv6 also delivers information messages such as ping, otherwise known as

- Echo request
- Echo reply



Important

By providing a framework for informational messages, ICMPv6 plays an important role in IPv6 features such as

- Neighbor discovery (ND)
- Path MTU discovery
- Multicast Listener Discovery (MLD)

You can identify an IPv6 ICMP packet because the Next Header field in the IPv6 packet header is 58.

Internet Protocol Security (IPsec) with ICMPv6

You can configure IPsec with ICMPv6. For a configuration example of IPsec with ICMPv6, see [IPsec with ICMPv6 configuration example](#) on page 2083.

Fragmented ICMP Packet Filtering

ICMP fragmentation distributed denial-of-service (DDoS) attacks flood the destination resources with fragmented packets and overwhelm the network because of massive volumes of traffic. With Fragmented ICMP packet filtering, the system inspects each incoming IPv6 ICMP packet to determine if it should drop the packet or forward it.

You can configure ICMP packet filtering globally, on a specific VRF, and on the following management interfaces:

- Out-of-Band (OOB) management
- Circuitless IP (CLIP) management
- VLAN management

Path MTU discovery

IPv6 routers do not fragment packets.

The source node can send packets less than or equal to the maximum transmission unit (MTU) of the link layer.

As the packet travels through the network to the source it can encounter a link with a smaller MTU. If so, the router sends the source node an ICMP error message that contains the MTU size of the next link. The source IPv6 node then resends a packet equal to the size of the MTU included in the ICMP message.

The default Layer 3 IPv6 MTU value is 1500 where the system MTU default value is 1950.

The default IPv6 MTU value is always less than the default System MTU value.

You can configure the MTU for each IPv6 interface.



Note

The MTU value for an IPv6 interface is not configurable on 5320 Series, 5420 Series, and 5720 Series.



Note

To configure separate Layer 3 MTU values for IPv4 and IPv6 packets on the same VLAN interface, you must disable Unicast Reverse Path Forwarding (uRPF) mode. If you enable the uRPF mode using the command **boot config flags urpf-mode**, the MTU values for both IPv4 and IPv6 packets on the same VLAN are matched. Different Layer 3 MTU sizes on the same VLAN are not allowed in uRPF mode.

Routing

A routing table is present on all nodes.

The routing table stores information about IPv6 network prefixes and how to reach them.

**Note**

The switch requires routing protocols, such as OSPFv3 to exchange IPv6 routing prefixes.

For each incoming packet, the switch checks the destination neighbor cache first. If the destination is not in the destination neighbor cache, the routing table determines:

- the next-hop interface (the interface used for forwarding)
- the next-hop address

**Note**

The system uses the IPv6 Neighbor Cache for on-link, directly-connected destinations only. Off-link destinations go through a next-hop router, as determined by the next-hop address lookup.

IPv6 routes in a routing table can be:

- directly attached network routes using a 64-bit prefix
- remote network routes using a 64-bit or lower prefix
- host routes using a 128-bit prefix length
- the default route using a prefix of ::/0

The switch supports the following IPv6 routing protocols:

- BGP+ (over 6in4 tunnels)
- BGPv6
- IPv6 Shortcuts (over Fabric Connect)
- OSPFv3
- RIPng

You can redistribute IPv6 routes between any of these routing protocols.

To configure IPv6 routing on a VLAN, an IP address is assigned to the VLAN. This IP address is not associated with any particular physical port, but is used on all ports where this VLAN is a member.

On a brouter port, a single port VLAN is used to route the traffic. IPv4 and IPv6 traffic is routed in the single-port brouter VLAN.

Other VLANs (which are multiple port VLANs) can bridge and route the traffic.

Virtual routing between IPv6 subnets

The switch supports IPv6 routing between subnets.

When you add an IP address to the VLAN, the system maps an IP subnet to the VLAN.

As shown in the following figure, although VLAN 1 and VLAN 2 reside on the same switch, for traffic to flow from VLAN 1 to VLAN 2, you must route the traffic.

You must enable IPv6 forwarding to route IPv6 traffic between VLANs. And you must enable IPv6 both globally and on a specific VLAN basis in order for forwarding to function. You can enable or disable IPv6, either globally or on a specific VLAN basis.

IPv6 forwarding is enabled by default.




Figure 162: IPv6 routing between VLANs

When you configure routing on a VLAN, an IPv6 address assigned to the VLAN is the VLAN IP interface.

The VLAN IPv6 address can be reached through any VLAN port, and frames route from the VLAN through the gateway IPv6 address.

You can forward traffic to any IPv6 subnet in the switch. A VLAN can be reached only if it has an IPv6 interface configured on it.

Because a port can belong to multiple VLANs, a one-to-one correspondence no longer exists between the physical port and the router interface when VLAN tagging is enabled.

If you do not enable VLAN tagging a single port can belong only to one port-based VLAN, but that same single port can belong to multiple policy-based VLANs.

As with any IPv6 address, you can use any VLAN IP interface for device management.

For the Simple Network Management Protocol (SNMP) or Telnet management, you can use any VLAN IP interface to access the switch while routing is enabled on the VLAN.

Brouter ports

A brouter port is a single-port VLAN that can route IP packets and bridge all nonroutable traffic.

The difference between a brouter port and a standard protocol-based VLAN configured for routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic while it routes IP traffic.



Note

Because a brouter port is a one-port VLAN, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

Static routes

Static routes provide an alternative method for establishing route reachability.

Static routes, with dynamic routes, provide routing information from the forwarding database.

Only enabled static routes whose nexthop address is reachable are submitted to the Route Table Manager (RTM), which determines the best route based on reachability, route preference, and cost.

The RTM communicates all updates to best routes.

If the nexthop is not reachable you can use the **show ipv6 route static** command to display the status. If the nexthop is not reachable, the status is `TryToResolve` and the system does not display the route in the RTM until the nexthop address is resolved.

For directly-connected IPv6 Subnets you do not need to specify a nexthop address; you can specify outgoing Tunnel-ID, VLAN, or port. If you use outgoing Tunnel-ID, VLAN, or port, the implied nexthop value is `0::0`.

When you configure IPv6 static routes only by interface (VLAN or brouter), it lets the traffic to reach IPv6 prefixes configured on the link that is directly connected to the interface provided in the static route configuration. For example: `ipv6 route 180:0:0:0:0:0:0:0/64 cost 1 vlan 631`.

When you configure static routes with a link-local nexthop, you must also specify the outgoing Tunnel-ID, VLAN, or port because link-local addresses are ambiguous unless the proper interface binding is attached. For example: `ipv6 route 1234::/64 cost 1 next-hop fe80::1 vlan 1900`.

You must provide the following options to configure a static route:

- local or nonlocal hop option

Configure a static route either with a next hop that exists on a locally attached network or a next hop that is reachable through a dynamic route. The static route is available as long as the next hop is reachable.

- route preference

You can specify the route preference for the static routes as follows:

- Global value for all static routes: the preference is either static or dynamic routes.
- Preference for each static route entry: if specified, this value overrides the global value for the entry which provides flexibility to change the general behavior of a specific static route.

- Administrative status

Controls when the static route is considered for forwarding. Administrative status differs from the operational status. An admin-enabled static route can still be unreachable and not used for forwarding. An admin-disabled static route is operationally a nonexistent route.

- Multiple static routes

Specify alternative paths to the same destination. Multiple static routes provide stability and load balancing.

To configure a default static route, supply a value of 0 for the prefix and the prefix length.

The following table describes events that affect static route operation.

Table 166: Events and their affects on static route operation

Action	Result
Change the administrative status of the static route	Makes the static route unavailable for forwarding You can use one CLI command to administratively enable or disable all static routes as follows <code>ipv6 route static enable</code> . You can administratively disable all routes but preserve the static route configuration when you use the CLI command: <code>no ipv6 static route enable</code> .
Delete the IPv6 addresses of a VLAN or brouter port	Permanently deletes the static routes with the corresponding local neighbors from the RTM, the forwarding database, and the configuration database
Delete a VLAN	Removes static routes with a local next-hop option from the configuration database. Static routes with a nonlocal next-hop option become inactive (they are removed from the forwarding database).
Disable forwarding on a VLAN or brouter port	Static routes reachable through the locally attached network become inactive
Disable a VLAN or brouter port	Makes the static route inactive
Disable IPv6 forwarding globally	Stops forwarding all IPv6 traffic
Learn changes about a dynamically learned neighbor	After a neighbor becomes unreachable or is deleted, the static route with the neighbor becomes inactive, and the configuration is not affected. The static route with the neighbor becomes active in the configuration and is added to the RTM and forwarding database when the neighbor becomes reachable.
Enable a static route	Adds the route to the RTM to change certain static routes to active.
Delete a static route	Permanently deletes a static route from the configuration.
Disable a static route	Stops traffic on the static route but does not remove the route from the configuration.
Change a route preference	After the static route preference changes, the best routes for the entries use both static and dynamic paths.
Delete or disable a tunnel	Removes the tunnel entry from the forwarding table
Enable a tunnel	Activates the tunnel static routes and adds an entry to the forwarding table.

The local-nexthop flag is not required for IPv6.

An IPv4 device cannot learn a neighbor ARP entry unless the device uses a local route entry.

In IPv6, a host can learn a neighbor entry if the device is physically connected to the neighbor (one hop).

The static route becomes active when the next hop is reachable by a dynamic route neighbor resolution. The static route takes the forwarding information from the dynamic route. If the next hop is reachable using a local route, the neighbor resolution is required.

Static route table

The static route table is separate from the system routing table that the router uses to make forwarding decisions.

You can use the static route table to directly change static routes.

Although the tables are separate, the system routing table automatically reflects the static routing table manager entries if the next-hop address in the static route is reachable and if the static route is enabled.

The static route table is indexed by four attributes:

- destination network
- destination mask
- next hop
- interface

You can insert static routes by using the static route table, and you can delete static routes by using either the static route table or the system routing table. For information on route scaling, see *Fabric Engine Release Notes*.



Important

The system routing table stores only active static routes with the best route preference. A static route is active only if the route is enabled and if the next-hop address is reachable; for example, if a valid IPv6 neighbor cache entry exists for the next hop.

You can enter multiple routes (for example, multiple default routes) that use different costs and the lowest cost route that is reachable, the system displays the lowest cost route in the routing table.

If you enter multiple next hops for the same route with the same cost, the switch does not replace the existing route.

If you enter the same route with the same cost and a different next hop, the switch uses the first route. If that first route becomes unreachable, the system activates the second route, with a different next-hop, with no connectivity loss.

NEW! IPv6 Discard Static Routes

A discard static route is a route with an invalid next hop that results in traffic matching the route (and not matching the more specific routes) being discarded. While the router aggregates or injects routes to other routers, the router does not have a path to the aggregated destination. In such cases, the result is a discard static route and a routing loop. To avoid routing loops, configure a discard static route to the destination the router is advertising.

You can configure a preference value for a discard static route. However, you must configure that preference value appropriately so that when you want to use the discard static route, it is elected as the best route.

IPv6 discard static route distribution supports the following IPv6 routing protocols:

- BGP+ (over 6in4 tunnels)
- BGPv6
- IPv6 Shortcuts (over Fabric Connect)
- OSPFv3
- RIPng (GRT support only)



Note

BFD and ECMP do not apply to IPv6 discard routes.

Before you add an IPv6 discard static route, perform a check to ensure that no other static route to that identical destination is enabled. If such a route exists, you cannot add the discard route and the system displays an error message. If you enable a discard static route, you cannot add another static route to that destination. You must first delete the discard static route before you add a regular static route to that destination.

Only discard static routes that belong to the static type protocol are supported. An inter-VRF discard static route is not installed in the routing table of the destination VRF on the same switch.

Route scaling

IPv4 and IPv6 route scaling depends on the combination of the `ipv6-mode` and `urpf-mode` boot config flags. For more information, see *Fabric Engine Release Notes*.

IPv6 Circuitless IP

IPv6 Circuitless IP (CLIP) is a virtual interface that is not associated with any physical port. You can use an IPv6 CLIP interface to provide uninterrupted connectivity to your switch as long as an actual path exists to reach the device. The system treats the IPv6 CLIP interface like an IPv6 interface and treats the network associated with the IPv6 CLIP as a local network attached to the device. This route always exists and the circuit is always up because no physical attachment exists.

You can use an IPv6 CLIP address as a logical IPv6 address for network management, as well as for other purposes. The IPv6 CLIP is typically a host address with any prefix length. You can redistribute this address as part of any other routing protocol update, so that the CLIP address is known to neighbors and available for use in routing or other types of connectivity. You can use IPv6 CLIP for many kinds of management connectivity, such as Telnet or SSH. You can also use IPv6 CLIP as a source IP address for sending Syslog messages.

For scaling information on IPv6 CLIP, see *Fabric Engine Release Notes*.

IPv6 CLIP restrictions and limitations

This section describes the restrictions and limitations associated with IPv6 CLIP.

- Stateless address autoconfiguration (SLAAC) is not supported on IPv6 CLIP interfaces.
- IPv6 CLIP does not support link-local address configuration.
- To configure an IPv6 address with a prefix length from 65 to 127 on a CLIP interface, you must enable the IPv6 mode flag.
- Neighbor discovery (ND) does not run on an IPv6 CLIP interface. Therefore, the system does not detect duplicate IPv6 address assignment to this interface.
- Multiple IPv6 address configuration on an IPv6 CLIP interface is not supported.
- IPv6 CLIP interface is enabled by default and it cannot be disabled.
- You cannot configure an IPv6 CLIP interface as the source or destination endpoint of an IPv6-in-IPv4 tunnel.

Equal Cost Multipath

Table 167: Equal Cost Multiple Path for IPv6 product support

Feature	Product	Release introduced
ECMP for IPv6	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

With Equal Cost Multipath (ECMP), the switch can determine equal-cost paths to the same destination prefix. You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP traffic. ECMP is formed using routes from the same source or protocol, and the ECMP routes are evaluated within each routing protocol.

The ECMP feature supports and complements the following protocols and route types:

- OSPFv3
- Static routes
- BGP+
- RIPng
- IPv6 Shortcuts

ECMP with static routes

ECMP supports and complements static routes.

The following points need to be considered while configuring ECMP with static routes:

- When ECMP is globally enabled, the equal cost static routes are added in the Routing Table Manager (RTM).
- Static routes that are configured only using an interface, such as VLAN or port, do not support ECMP as these routes have a preference value of 0 and are treated like local routes.
- If your switch supports a management interface, then static routes configured on the management interface do not support ECMP.
- Static routes configured by next-hop are not considered equal cost with routes that are configured by tunnel even if the routes have the same cost and preference.
- A static route configured by tunnel is the least preferred and is programmed only when a next-hop or an interface does not configure a static route.
- Static routes configured by next-hop that resolves their next-hop using another static route will be in the `notReachable` state even if the next-hop can be pinged.
- If there are two static routes configured by next-hop, and both next-hops are resolved via a dynamic protocol to the same value, then only one route will be in the `reachableInRtm` state. The state of the other route will be `reachableNotInRtm`.

Disable IPv6 ICMP multicast

On IPv6 networks, a packet can be directed to an individual machine or multicasted to a set of hosts. When a packet is sent to an IPv6 multicast address from a machine on the local network, that packet is delivered to a subset or all machines on that network.

If the packet that is sent to the multicast address is an ICMP Echo Request packet, the machines on the network will receive this ICMPv6 echo request packet and send an ICMP echo reply packet back. When all the machines on a network respond to this ICMPv6 echo request, the result can be severe network congestion or outages.

Network devices always respond to the IPv6 ICMP packets sent to a multicast address. However, you can disable the processing of IPv6 ICMP packets sent to a multicast address on the device. On disabling the ICMP multicast processing, all the packets containing ICMP sent to multicast addresses are dropped when the packets reach the control plane.

View IPv6 Connections

You can establish network connectivity with the following protocols:

- Transmission Control Protocol (TCP), for connection-oriented sessions
- User Datagram Protocol (UDP), for connectionless sessions

When you view TCP information you can:

- check the status of the connections, from the switch perspective, as they traverse the network
- detect intermittent connectivity
- detect attacks on resources
- determine which applications are active by checking the port numbers

UDP endpoint information tells you about local and remote UDP activity.

When you view UDP information you can:

- determine which applications are active by checking the local and remote port numbers
- identify processes within a UDP session to enable multiplexing of a port mapping for UDP

IPv6 Basic Configuration using CLI

Use the procedures in this section to configure IPv6 basics using CLI.

Enable the IPv6-Mode Boot Config Flag

Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits.

**Note**

This procedure does not apply to 4220 Series and 7830 Series.

About This Task

This flag is disabled by default. Use this procedure to enable (set to true) the IPv6-mode boot config flag.

When the IPv6-mode boot config flag is enabled, the maximum number of IPv4 routing table entries decreases. For scaling information, see *Fabric Engine Release Notes*.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the IPv6-mode boot config flag:
`boot config flags ipv6-mode`
3. Save the configuration, and then reboot the switch for the change to the IPv6-mode boot config flag to take effect.
4. After you reboot the switch, verify that the IPv6-mode boot config flag is set to true:
`show boot config flags`

Configure an IPv6 Static Neighbor Address

You can use static IPv6 neighbors to manually specify the link-layer address for a given IPv6 endpoint.

About This Task

Under normal operation you do not need to configure static IPv6 neighbors.

However, IPv6 static neighbors can be used to:

- avoid the overhead associated with dynamic neighbor discovery protocol traffic
- help troubleshoot specific network scenarios



Note

- IPv6 static neighbors are not supported on SMLT.
- When you add or remove IPv6 static neighbors that point to a nexthop router, make sure that you have disabled the IPv6 interface.

Not all parameters are available in non-default VRFs.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure an IPv6 neighbor address:

```
ipv6 neighbor WORD<0-128> port {slot/port[sub-port]} mac  
<0x00:0x00:0x00:0x00:0x00:0x00> [vlan <1-4059>
```

3. Configure optional parameters if the default values do not meet your requirements:

- a. Configure the hop limit:

```
ipv6 hop-limit <0-255>
```

The default is 64.

- b. Configure ICMP network address unreachable messages:

```
ipv6 icmp addr-unreach
```

- c. Configure the ICMP error interval:

```
ipv6 icmp error-interval <1-2147483647>
```

The interval is in milliseconds. An interval of 0 results in no error messages. The default is 1000.

- d. Configure the ICMP error quota:

```
ipv6 icmp error-quota <0-2000000>
```

The default is 50.

- e. Configure ICMP port unreachable messages:

```
ipv6 icmp port-unreach
```

- f. Enable response to icmp echo multicast packets:

```
ipv6 icmp echo-multicast-request
```

The default is disabled.

- g. Enable ICMP network unreachable messages:

```
ipv6 icmp unreachable-msg
```

The default is disabled.

Examples

Add an IPv6 neighbor for a brouter port:

```
Switch:1(config)#ipv6 neighbor 3000:0:0:0:0:0:2 port 1/11 mac 00:0c:42:07:35:90
```

Add an IPv6 neighbor for a VLAN:

```
Switch:1(config)#ipv6 neighbor 3000::3 port 1/12 mac 01:02:03:04:05:06 vlan 20
```

Variable Definitions

Use the data in the following table to use the **ipv6** commands in this procedure.

Variable	Value
<i>forwarding</i>	Configures whether this entity is an IPv6 router with respect to the forwarding of datagrams received by, but not addressed to, this entity. Enable forwarding to act as a router. The default is enabled.
<i>hop-limit <0-255></i>	Configures the hop limit. The default is 64.
<i>icmp addr-unreach</i>	Enables ICMP address unreachable messages. The default is enabled.
<i>icmp echo-multicast-request</i>	Enable response to icmp echo multicast packets. The default is enable.
<i>icmp error interval <1-2147483647></i>	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000.
<i>icmp error-quota <0-2000000></i>	Configures the number of ICMP error messages that can be sent during the ICMP error interval. A value of zero instructs the system not to send an ICMP error messages. The default value is 50.
<i>icmp port-unreach</i>	Enables ICMP port unreachable messages. The default is enabled.

Variable	Value
<code>icmp unreachable-msg</code>	Enables ICMP network unreachable messages. The default is disabled.
<code>neighbor WORD<0-128> port {slot/port [/sub-port]} mac <0x00:0x00:0x00:0x00:0x00:0x00> [vlan <1-4059>]</code>	<p>Creates a static IPv6 neighbor with the following variables:</p> <ul style="list-style-type: none"> <code>WORD<0-128></code> specifies the IPv6 address of the neighbor in hexadecimal colon format. <code>{slot/port [/sub-port]}</code> specifies the router port to use for the neighbor. <p>Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <ul style="list-style-type: none"> <code>mac <0x00:0x00:0x00:0x00:0x00:0x00></code> specifies the MAC address of the neighbor. <code>vlan <1-4059></code> specifies the VLAN ID to use for the neighbor. <p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <code>vrf-scaling</code> and <code>spbm-config-mode</code> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>Static IPv6 neighbors do not maintain any state machine and the system assumes that they are always reachable.</p>

Configure an IPv6 Interface

The information in this section can help you configure an IPv6 interface to make IPv6 active on the interface and fine-tune IPv6 neighbor discovery to control the frequency of protocol traffic.

By default, IPv6 forwarding is enabled on an interface.

Compared to IPv4/ARP, the IPv6 neighbor discovery mechanism maintains more protocol state, timers, and protocol traffic overhead.

There are two important tunable parameters for IPv6 ND that can control the frequency of protocol traffic:

- `ipv6 interface reachable-time`
- `ipv6 interface retransmit-timer`

Before You Begin

- Before you can assign an IPv6 address to the interface, you must configure an IPv6 interface for a VLAN or brouter port.

You must configure a VLAN before you can give the VLAN an interface identifier or an IPv6 address.

The switch supports port-based and IPv6 protocol-based VLANs.

For information about how to configure VLANs, see the followings:

- [VLAN Configuration Using CLI](#) on page 4007
- [MultiLink Trunking and Split MultiLink Trunking](#) on page 2657

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create IPv6 interface:

```
ipv6 interface
```

3. Configure optional parameters to meet your requirements:

- a. Enable IPv6 router advertisement on the interface:

```
ipv6 nd send-ra
```

- b. Configure the maximum number of hops before packets drop:

```
ipv6 interface hop-limit <1-255>
```

- c. Configure the link-local address:

```
ipv6 interface link-local WORD<0-19>
```

- d. Configure the mac offset:

```
ipv6 interface mac-offset <0-511>
```

- e. Configure the maximum transmission unit (MTU):



Note

This step does not apply to 5320 Series , 5420 Series, or 5720 Series.

```
ipv6 interface mtu <1280-9500>
```

- f. Configure an interface description:

```
ipv6 interface name WORD<0-255>
```

- g. Configure the time a neighbor is considered reachable after receiving a reachability confirmation:

```
ipv6 interface reachable-time <1-3600000>
```

- h. Configure the time between retransmissions of neighbor solicitation messages to a neighbor when resolving the address or when probing the reachability of a neighbor:

```
ipv6 interface retransmit-timer <1-4294967295>
```

- i. Configure a brouter port as part of an IPv6 VLAN:

```
ipv6 interface vlan <1-4059>
```

- j. Configure the interface to perform IPv6 unicast reverse path forwarding:

- To enable urpf-mode boot flag, enter:

```
boot config flags urpf-mode
```

- To configure unicast reverse path forwarding, enter:

```
ipv6 rvs-path-chk mode {strict|exist-only}
```

Example



Note

In contrast to IPv6 interface creation and address assignment in EDM, you use the **ipv6 interface** CLI command to create an interface and specify a single global address in one step.

Create and administratively enable the interface:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/4
Switch:1(config-if)#ipv6 interface enable
```

Variable Definitions

Use the data in the following table to use the **ipv6 interface** command.

Variable	Value
<i>hop-limit</i> <1-255>	Configures the maximum hops. The default is 64.
<i>link-local WORD</i> <0-19>	Specifies the 64-bit interface ID used to calculate the actual link-local address as a name up to 19 characters long.

Variable	Value
<i>mac-offset</i> <0-511>	<p>Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated.</p> <p>Note: This parameter applies only to VLANs.</p> <p>You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.</p>
<i>mtu</i> <1280-9500>	<p>Configures the maximum transmission unit for the interface: 1280-1500, 1850, or 9500. This value must be the same for all addresses defined on this interface.</p> <p>The default is 1500.</p> <p>Different hardware platforms support different MTU values. To see what values your switch supports, use the CLI command completion help.</p>
<i>name</i> WORD<0-255>	<p>Assigns a descriptive name. The network management system also configures this string.</p>
<i>reachable-time</i> <0-3600000>	<p>Controls how long IPv6 neighbor entries learned on an interface remain in the REACHABLE state (as described in RFC 4861).</p> <p>The system randomizes the value you configure, per RFC specifications, to be 50%-150% of the configured value.</p> <p>By default the reachable-time base value is 30 seconds, with an actual 15-45 second range when you consider the randomization factor.</p> <p>The default is 3000 milliseconds</p>

Variable	Value
<code>retransmit-timer <0-4294967295></code>	<p>Controls the time, in milliseconds, between retransmission of Neighbor Solicitation messages when the system attempts to resolve or reconfirm the reachability of an IPv6 neighbor.</p> <p>By default, the system sends three Neighbor Solicitation messages with a one second interval between each message. If the system does not receive a corresponding Neighbor Advertisement within an interval equal to 3 X retransmit-timer milliseconds, the system declares the IPv6 neighbor unreachable.</p> <p>Tip: You can increase the retransmit-timer to extend the interval that the switch waits until it declares the neighbor unreachable. For example: a retransmit-timer value of 5000 means the switch waits 3 X 5000 milliseconds which equals 15000 milliseconds or 15 seconds.</p> <p>The default is 1000 milliseconds</p>
<code>vlan <1-4059></code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>This value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.</p>

Use the data in the following table to use the **ipv6 rvs-path-chk** command.

Variable	Value
<code>mode {strict exist-only}</code>	<p>Specifies the mode for Unicast Reverse Path Forwarding (uRPF).</p> <p>In strict mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. If the incoming interface is not the best reverse path, the packet check fails and uRPF drops the packet.</p> <p>In exist-only mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. The packet is dropped only if the source address is not reachable via any interface on that router.</p>

Use the data in the following table to use the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Assign IPv6 Addresses to a Brouter Port or VLAN

Assign IPv6 addresses to interfaces to configure IPv6 routing for the interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Assign an IPv6 address:

```
ipv6 interface address WORD<0-255>
```

Example

Assign an IPv6 address specifying the full 128 bits of the address:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/5
Switch:1(config-if)#ipv6 interface address 30:0:0:0:0:0:0:ffff/64
```

Variable Definitions

The following table defines parameters for the **ipv6 interface address** command.

Variable	Value
<i>WORD</i> <0–255>	Specifies the IPv6 address for the port or VLAN.

The following table defines parameters for the **interface** command.

Variable	Value
<1–4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{ <i>slot/port</i> [/ <i>sub-port</i>] [- <i>slot/port</i> [/ <i>sub-port</i>]] [<i>,...</i>] [<i>slot/all</i>] [<i>all</i>] }	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure IPv6 Route Preferences

Configure IPv6 route preferences to give preference to routes learned for a specific protocol.

Before You Begin



Important

Changing route preferences can affect system performance and network accessibility while you perform the procedure. Change a prefix list or a routing protocol before you activate the protocols.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the route preference:

```
ipv6 route preference protocol <static|ebgp|ibgp|ospfv3-intra|ospfv3-  
inter|ospfv3-extern1|ospfv3-extern2|ripng|spbm-level1> <0-255>
```

3. Confirm that the configuration is correct:

```
show ipv6 route preference [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Example

Configure the route preference to RIPng and confirm the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ipv6 route preference protocol ripng 10
Switch:1(config)#show ipv6 route preference
```

IPv6 Route Preference - GlobalRouter		
PROTOCOL	DEFAULT	CONFIG
LOCAL	0	0
STATIC	5	5
SPBM_L1	7	7
OSPFv3_INTRA	20	20
OSPFv3_INTER	25	25
EBGP	45	45
RIPNG	100	100
OSPFv3_E1	120	120
OSPFv3_E2	125	125
IBGP	175	175

Variable definitions

Use the data in the following table to use the **ipv6 route preference** and the **show ipv6 route preference** commands.

Variable	Value
<0-255>	Assigns a route preference value.
<i>ebgp</i>	Configures the preference for protocol type EBGp.
<i>ibgp</i>	Configures the preference for protocol type IBGP.

Variable	Value
<i>ospfv3-extern1</i>	Configures the preference for protocol type OSPFv3 external type 1.
<i>ospfv3-extern2</i>	Configures the preference for protocol type OSPFv3 external type 2.
<i>ospfv3-intra</i>	Configures the preference for protocol type OSPFv3 intra-area.
<i>ospfv3-inter</i>	Configures the preference for protocol type OSPFv3 inter-area.
<i>ripng</i>	Configures the preference for protocol type RIPng.
<i>spbm-level1</i>	Configures the preference for protocol type spbm-level1.
<i>static</i>	Configures the preference for protocol type static.

View Global IPv6 Information

View and manage general IPv6 information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display IPv6 information for an interface:

```
show ipv6 interface [gigabitethernet {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]} [loopback <1-256>][tunnel <1-
2000>][vlan <1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
3. Display IPv6 tunnel information:

```
show ipv6 interface tunnel <1-2000>
```
4. Display IPv6 address information for the specified slot and port:

```
show ipv6 address interface gigabitethernet {slot/port[/sub-port]}[-
slot/port[/sub-port]][,...][slot/all][all]} [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```
5. Display IPv6 address information for the specified IPv6 address:

```
show ipv6 address interface ip WORD<0-46> [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```
6. Display IPv6 address information for the specified tunnel:

```
show ipv6 address interface tunnel <1-2000>
```
7. Display IPv6 address information for the specified VLAN:

```
show ipv6 address interface vlan <1-4059>
```
8. Display the current state of IPv6 forwarding:

```
show ipv6 forwarding [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
9. Display information on the current state of IPv6 functionality:

```
show ipv6 global [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
10. Display IPv6 Gigabit Ethernet (GbE) router advertisement information:

```
show ipv6 nd interface gigabitethernet [{slot/port[-slot/port]}[,...]]
```


11. Display IPv6 Gigabit Ethernet (GbE) router advertisement information:

```
show ipv6 nd interface gigabitethernet [{slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]]
```

12. Display IPv6 VLAN router advertisement information:

```
show ipv6 nd interface vlan [<1-4059>]
```

13. Display detailed information in IPv6 router advertisements:

```
show ipv6 nd-prefix detail [vrf WORD<1-16> | vrfids WORD<0-512>]
```

14. Display GbE interface information in IPv6 router advertisements:

```
show ipv6 nd-prefix interface gigabitethernet [{slot/port[-slot/port]}
[,...]]
```

15. Display GbE interface information in IPv6 router advertisements:

```
show ipv6 nd-prefix interface gigabitethernet [{slot/port[/sub-port]}[-
slot/port[/sub-port]][,...][slot/all][all]]
```

16. Display VLAN interface information in IPv6 router advertisements:

```
show ipv6 nd-prefix interface vlan [<1-4059>]
```

17. Display VLAN information in IPv6 router advertisements:

```
show ipv6 nd-prefix vlan <1-4059>
```

18. Display IPv6 neighbor entries with specific brouter port numbers:

```
show ipv6 neighbor interface gigabitethernet {slot/port[/sub-port]}
```

19. Display IPv6 neighbor information for neighbors of the specified type:

```
show ipv6 neighbor type <1-4>
```

20. Display IPv6 neighbor information:

```
show ipv6 neighbor [WORD<0-46>]
```

Examples

```
Switch:1(config)#show ipv6 interface vlan
```

```
=====
                        Vlan Ipv6 Interface
=====
IFINDX VLAN PHYSICAL      ADMIN  OPER  TYPE  MTU  HOP REACHABLE  RETRANSMIT  MCAST  IPSEC  RPC  RPCMODE
INDX   ADDRESS          STATE   STATE                LMT TIME      TIME      STATUS
-----
2070   22   e4:5d:52:3c:65:02 disable down  ETHER 1500 64  30000      1000      disable disable disable existonly
=====
```

```
=====
                        Vlan Ipv6 Address
=====
IPV6 ADDRESS                VLAN-ID  TYPE  ORIGIN  STATUS
-----
fe80:0:0:0:e65d:52ff:fe3c:6502  V-22    UNICAST LINKLAYER INACCESSIBLE
=====
```

```
1 out of 2 Total Num of Interface Entries displayed.
1 out of 2 Total Num of Address Entries displayed.
```

```
Switch:1#show ipv6 interface tunnel
```

```
=====
                        Tunnel Ipv6 Interface - GlobalRouter
=====
IF  Descr      VLAN PHYSICAL      ADMIN  OPER  TYPE  MTU  HOP REACHABLE  RETRANSMIT  MCAST
INDX          ADDRESS          STATE   STATE                LMT TIME      TIME      STATUS
-----
6656 T-1      0    n/a              enable down  P2P   1280 64  30000      1000      disable
=====
```

```

Tunnel Ipv6 Address
=====
IPV6 ADDRESS                                TUNNEL-ID    TYPE      ORIGIN    STATUS
-----
11:0:0:0:0:11:11/32                        T-1          UNICAST   MANUAL    INACCESSIBLE INF
INF
fe80:0:0:0:0:b0b:b0b/64                    T-1          UNICAST   LINKLAYER INACCESSIBLE INF
INF

1 out of 5 Total Num of Interface Entries displayed.
2 out of 9 Total Num of Address Entries displayed.

Switch:1#show ipv6 address interface tunnel 2

=====
Address Information
=====
IPV6 ADDRESS/PREFIX LENGTH                  VID/BID/TID  TYPE      ORIGIN    STATUS
-----
44:211:0:0:0:0:2/64                        T-2          UNICAST   MANUAL    PREFERRED
fe80:0:0:0:0:d301:3702/64                  T-2          UNICAST   LINKLAYER PREFERRED

2 out of 407 Total Num of Address Entries displayed.
Switch:1#show ipv6 address interface vlan 100

=====
Address Information
=====
IPV6 ADDRESS/PREFIX LENGTH                  VID/BID/TID  TYPE      ORIGIN    STATUS
-----
10:1:50:0:0:0:7/64                         V-100        UNICAST   MANUAL    PREFERRED
fe80:0:0:0:b2ad:aaff:fe46:f19a/64          V-100        UNICAST   LINKLAYER PREFERRED

2 out of 407 Total Num of Address Entries displayed.
Switch:1#show ipv6 forwarding
      Ipv6 forwarding - GlobalRouter : enable
      ecmp                          : disable
      ecmp-max-path                  : 1

Switch:1#show ipv6 global
=====
IPv6 Global Information - GlobalRouter
=====
forwarding                               : disable
default-hop-cnt                         : 64
number-of-interfaces                    : 1
icmp-error-interval                     : 1000
icmp-error-quota                        : 50
icmp-unreach-msg                        : disable
icmp-addr-unreach-msg                   : enable
icmp-port-unreach-msg                   : enable
icmp-echo-multicast-request              : enable
static-route-admin-status                : enable
alternative-route                       : enable
ecmp                                     : disable
ecmp-max-path                           : 1
source-route                            : disable
host-autoconfig                         : disable

Switch:1#show ipv6 nd interface vlan

=====
Vlan Ipv6 Nd - GlobalRouter
=====
IFID VLAN  RTR-ADV MAX-INT MIN-INT LIFETIME MANAG OTHER DAD-NS MTU  HOP

```

REACHABLE RETRANSMIT

TIME TIME FLAG CONF LIMIT

```

2092 V-44 True 600 200 1800 False False 1 0(d) 64(d) 30000(i) 1000(i)
2081 V-33 True 600 200 1800 False False 1 0(d) 64(d) 30000(i) 1000(i)

```

Note: (s) = Set, (d) = Default, (i) = inherit.

2 out of 2 Total Num of Ipv6 ND Entries displayed.

Switch:1#show ipv6 nd-prefix interface gigabitethernet

Port Ipv6 Nd Prefix

INTF INDEX	IPV6 ADDRESS/PREFIX	BTR	VALID LIFE	PREF LIFE	EUI
344	2011:beef:4004:0:0:0:0/64		5/25	2592000	604800 1

1 out of 9 Total Num of Ipv6 ND prefix Entries displayed.

Switch:1#show ipv6 nd-prefix interface vlan

Vlan Ipv6 Nd Prefix

INTF INDEX	IPV6 ADDRESS/PREFIX	VLAN ID	VALID LIFE	PREF LIFE	EUI
2148	2011:beef:100:0:0:0:0/64	100	2592000	604800	1
2158	2011:beef:110:0:0:0:0/64	110	2592000	604800	1
2248	2011:beef:200:0:0:0:0/64	200	2592000	604800	1
2258	2011:beef:210:0:0:0:0/48	210	2592000	604800	1
2548	2011:beef:500:0:0:0:0/64	500	2592000	604800	1
2648	2011:beef:600:0:0:0:0/64	600	2592000	604800	1
2948	2011:beef:900:0:0:0:0/64	900	2592000	604800	1

7 out of 9 Total Num of Ipv6 ND prefix Entries displayed.

Switch:1#show ipv6 nd-prefix vlan 100

Nd-Prefix Address Information

INTF INDEX	IPV6 ADDRESS/PREFIX	VLAN ID	VALID LIFE	PREF LIFE	EUI
2148	2011:beef:100:0:0:0:0/64	100	2592000	604800	1

Legend: EUI: eui-not-used(1), eui-used-with-ul-complement(2) eui-used-without-ul-complement(3)

Switch:1#show ipv6 neighbor type 4

Neighbor Information - GlobalRouter

NET ADDRESS/ PHYSICAL ADDRESS	IPV6 INTF	PHYS INTF	TYPE	STATE	LAST UPD	TUNNEL
123:0:0:0:0:0:123/ 00:00:00:00:00:01	C-1	cpp	LOCAL	REACHABLE	0 days 00h:00m:01s	

1 out of 1 Total Num of Neighbor Entries displayed.

Switch:1(config)#show ipv6 neighbor

```

=====
Neighbor Information - GlobalRouter
=====
NET ADDRESS/          IPV6  PHYS      TYPE  STATE  LAST      TUNNEL
PHYSICAL ADDRESS      INTF  INTF
-----
22:0:0:0:0:0:22/      V-22  1/2       DYNAMIC REACHABLE  0 days 00h:00m:01s
84:83:71:49:38:82
22:0:0:0:0:0:44/      V-22  cpp       LOCAL  REACHABLE  0 days 00h:00m:00s
d4:ea:0e:c8:8c:81
fe80:0:0:0:8683:71ff:fe49:3882/ V-22  1/2       DYNAMIC DELAY      0 days 00h:00m:00s
84:83:71:49:38:82
fe80:0:0:0:d6ea:eff:fec8:8c81/ V-22  cpp       LOCAL  REACHABLE  0 days 00h:00m:00s
d4:ea:0e:c8:8c:81
=====

```

Variable Definitions

Use the data in the following table to use the **show ipv6** commands.

Variable	Value
<code>address interface ip WORD<0-46></code>	Specifies the IPv6 address.
<code>neighbor [WORD<0-46>]</code>	Specifies the IPv6 address of the neighbor.
<code>loopback <1-256></code>	Specifies the loopback interface ID value. If you do not specify a value, the output includes all IPv6 loopback interfaces.
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>type <1-4></code>	<p>Specifies the neighbor type as one of the following:</p> <ul style="list-style-type: none"> • 1 - other • 2 - dynamic • 3 - static • 4 - local
<code>tunnel <1-2000></code>	Specifies the tunnel ID.
<code>vlan <1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrfWORD<1-16></code>	Specifies the VRF name.
<code>vrfidsWORD<1-256></code>	Specifies the VRF ID.

Create IPv6 Static Routes

Perform the steps in this task to:

- Create static routes for data traffic in either the GRT or a specific VRF context for any platform.
- Create static routes for a VRF associated with a Segmented Management Instance CLIP or VLAN (routing mode) interface. Specify the name of the VRF context in [Step 1](#).
- Configure an IPv6 discard static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.

Static routes for the management OOB, if supported, and management VLAN (non-routing mode) must use the Segmented Management Instance. For more information, see [Segmented Management](#) on page 106. The management CLIP can use the Segmented Management Instance or routes in the associated VRF routing table manager (RTM). The management VLAN (routing mode) uses the Fabric Engine routing table or routes in the associated VRF.

Before You Begin

If an IPv6 discard static route already exists for the destination, you must delete it before you add a regular IPv6 static route.

About This Task

Not all parameters are available in non-default VRFs.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable IPv6 static routes globally:

```
ipv6 route static enable
```

If you disable static routes globally, the system removes all enabled static routes from the RTM and does not add new static routes to the RTM.

3. Create one of the following types of static routes:

- A regular static route:
`ipv6 route WORD<0-46> [cost <1-65535>]`
- A discard static route by keyword:
`ipv6 route WORD<0-46> cost <1-65535> discard [name] WORD<0-64>`
- A discard static route by next-hop address:
`ipv6 route WORD<0-46> cost <1-65535> next-hop
ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff: [preference <1-255>]`
- A default static route:
`ipv6 route ::/0 cost <1-65535>`

4. Enable one of the following types of static routes:

- A regular static route: `ipv6 route WORD<0-46> enable`
- A discard static route: `ipv6 route WORD<0-46> enable discard`
- A default static route: `ipv6 route ::/0 enable`

Examples

Enable IPv6 static routes globally:

```
Switch:1(config)#ipv6 route static enable
```

Create and enable a static route through a global nexthop:

```
Switch:1(config)#ipv6 route 4000::/64 cost 1 next-hop 3000::2 enable
```

Create and enable a static route through an outgoing interface (VLAN or brouter port):

```
Switch:1(config)#ipv6 route 4000::/64 cost 1 vlan 1900 enable
```

Create and enable a static route through a link local nexthop and an outgoing interface:

```
Switch:1(config)#ipv6 route 4000::/64 cost 1 next-hop fe80::1 vlan 1900 enable
```

In the preceding example, you must specify the outgoing interface so that the system can apply the correct context to the link-local address.

Create an IPv6 discard route by keyword:

```
ipv6 route 4000::/64 cost 1 discard
```

Create an IPv6 discard route by next-hop address:

```
ipv6 route 4000::/64 cost 1 next-hop ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
```

Variable Definitions

Use the data in the following table to use the **ipv6 route** command.

Variable	Value
<i>WORD</i> <0-46>	Specifies the IPv6 destination address and prefix.
<i>enable</i>	Enables the static route. The default is enabled.

Variable	Value
<i>discard</i>	Specifies the static route as a discard static route.
<i>cost</i> <1-65535>	Specifies the cost or distance ratio to reach the destination for this static route. The default is 1.
<i>next-hop Word</i> <0-46>	Specifies the IPv6 address of the next hop on this route. You do not need to specify the next hop if the devices directly connect to one another. Configure the next hop if the two nodes do not share the same network prefix but reside on the same link. For a discard static route, configure the next hop address as ffff.ffff.ffff.ffff.ffff.ffff.ffff.ffff. This value indicates the traffic to this route is discarded.
<i>preference</i> <1-255>	Specifies the routing preference of the destination IPv6 address. The default is 5.
<i>{slot/port[/sub-port]}</i>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>tunnel</i> <1-2000>	Specifies the tunnel ID.
<i>vlan</i> <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

View Routes Information

View routes information to view the current configuration.

About This Task

Not all parameters are available in non-default VRFs.

IPv6 host routes created for the IPv6 local interfaces do not display in the routing table.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show route information for alternative routes:

```
show ipv6 route alternative [vrf WORD<1-16> | vrfids WORD<1-256>]
```
3. Show the number of OSPF, RIP, static, and local routes:

```
show ipv6 route count-summary [vrf WORD<1-16> | vrfids WORD<1-256>]
```
4. Show route information for a destination:

```
show ipv6 route dest WORD<0-46> [vrf WORD<1-16> | vrfids WORD<1-256>]
```
5. Show route information for a port:

```
show ipv6 route gigabitethernet {slot/port[sub-port]}
```

6. Show route information for a next-hop address:

```
show ipv6 route next-hop WORD<0-46> [vrf WORD<1-16> | vrfids WORD<1-256>]
```

7. Show route information for an SPBM IPv6 route:

```
show ipv6 route spbm-nh-as-mac
```

8. Show route information for a static route:

```
show ipv6 route static [vrf WORD<1-16> | vrfids WORD<1-256>]
```

9. Show route information for a tunnel:

```
show ipv6 route tunnel <1-2000>
```

10. Show route information for a VLAN:

```
show ipv6 route vlan <1-4059>
```

Example

```
Switch:1(config-if)#show ipv6 route
```

```
=====
IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen      NEXT HOP      NH VRF/ISID      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
-----
22:0:0:0:0:0:0:0/64              42-VSP-8404   -               V-22         LOCAL   1     0    B    0
123:0:0:0:0:0:0:0/64             39-4900-24XE  -               C-1          LOCAL   1     0    B    0
=====
```

```
2 out of 2 Total Num of Route Entries displayed.
```

```
TYPE Legend:
```

```
A=Alternative Route, B=Best Route, E=Ecmp Route
```

```
Switch:1#show ipv6 route count-summary
```

```
=====
IPv6 Route Summary - GlobalRouter
=====
VRF NAME      TOTAL    OSPF    RIP    BGP    STATIC    LOCAL    ISIS
-----
GlobalRouter   13       10      0      3      1         2        7
=====
```

```
Switch:1#show ipv6 route static
```

```
=====
Static Route Information - GlobalRouter
=====
DEST-IP      NET IFINDX(VID/BRT/TUN)  ENABLE  STATUS      NAME
NEXT-HOP      PREFERENCE
-----
22:0:0:0:0:0:0:0      64 2059 (V-11 ) enable  NotReachable  ExtSer10
66:0:0:0:0:0:0:66      5
=====
```

```
1 out of 1 Total Num of Static Routes displayed.
```

```
Global IPv6 Static Routes Admin Status: enable
```

```
Switch:1#show ipv6 route spbm-nh-as-mac
```

```
=====
IPv6 Routing Table Information - GlobalRouter
=====
Destination Address/PrefixLen      NEXT HOP      NH VRF/ISID      VID/BID/TID  PROTO  COST  AGE  TYPE  PREF
```



```

-----
2001:cdab:0:0:0:0:0:0/32      42-VSP-8404      -      V-611      LOCAL  1  0  B  0
2002:cdab:0:0:0:0:0:0/32      VSP-8404         -      V-10       ISIS   1  0  B  0
-----

```

Variable Definitions

Use the data in the following table to use the **show ipv6 route** command.

Variable	Value
<i>count-summary</i>	Shows the total number of OSPF, static, and local routes.
<i>dest WORD<0-46></i>	Specifies the IPv6 destination network address. The prefix value must match the prefix length.
<i>next-hop WORD<0-46></i>	Specifies the IPv6 address of the next hop on this route.
<i>spbm-nh-as-mac</i>	Shows the BMAC address as the next hop rather than the host name.
<i>{slot/port[/sub-port]}</i>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>static</i>	Shows static IPv6 routes.
<i>tunnel <1-2000></i>	Shows route entries for a specific tunnel ID.
<i>vlan<1-4059></i>	Shows route entries for a specific VLAN ID.
<i>vrfWORD<1-16></i>	Shows the VRF name.
<i>vrfidsWORD<1-256></i>	Shows the VRF ID.

Create an IPv6 CLIP Interface

About This Task

Perform this procedure to create an IPv6 CLIP interface and associate it with a specific VRF.

Procedure

1. Enter Loopback Interface Configuration mode


```
enable

configure terminal

interface Loopback <1-256>
```
2. Create an IPv6 loopback interface address:


```
ipv6 interface address WORD <0-255> vrf WORD <1-16>
```

3. (Optional) Enter a name for the IPv6 address:

`ipv6 interface name WORD <0-64>`

4. Ensure the configuration is correct:

`show ipv6 interface loopback <1-256>`

Example

```
Switch:1#show ipv6 interface loopback
=====
                        Loopback Ipv6 Interfaces
=====
IF   VRF      Descr   VLAN PHYSICAL  ADMIN  OPER  TYPE  MTU  HOP  REACHABLE  RETRANSMIT  IPSEC
INDX NAME                                     STATE  STATE                                     LMT  TIME    TIME      STATE
-----
1348 GlobalRouter CLIPv6-5 --00:00:00:00:00:05 enable  up    ETHER 1500 64 30000    1000    disable
=====
                        Loopback IPv6 Addresses
=====
IPv6 ADDRESS/PREFIX LENGTH    LOOPBACK-ID  TYPE    ORIGIN    STATUS    NAME
-----
2001:db8:0:0:0:0:ffff/32      C-5          UNICAST MANUAL    PREFERRED INF  INF  EXTREMESERVER_2
=====
Legend: NA - Information not available

1 out of 1 Total Num of Interface Entries displayed.
1 out of 1 Total Num of Address Entries displayed.
```

Variable Definitions

Use the data in the following table to use the **ipv6** commands.

Variable	Value
<code>WORD<1-256></code>	Specifies the CLIP interface ID.
<code>WORD<0-255></code>	Specifies the IPv6 address.
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>WORD<0-64></code>	Specifies the I-SID name associated with the IPv6 address.

Enable IPv6 ECMP

About This Task

Use the following procedure to enable IPv6 ECMP globally or on a specific VRF. IPv6 ECMP is disabled by default.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enter the following command to enable IPv6 ECMP:

```
ipv6 ecmp enable
```

3. Set the default value, IPv6 ECMP is disabled by default.

```
default ipv6 ecmp enable
```

4. Disable IPv6 ECMP globally:

```
no ipv6 ecmp enable
```

Variable Definitions

Use the data in the following table to use the **ipv6 ecmp** command.

Variable	Value
<i>enable</i>	Enables IPv6 ECMP globally. Note: <ul style="list-style-type: none">• Enabling IPv6 ECMP sets the maximum number of paths configured to its default value, which is 8.• Disabling IPv6 ECMP sets the maximum number of paths configured to 1.

Configure Maximum Number of ECMP Paths

Before You Begin

Enable ECMP on the switch before configuring the max-path value. For more information, see [Enable IPv6 ECMP](#) on page 2212.

About This Task

Use the following procedure to configure the maximum number of ECMP paths.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the maximum number of ECMP paths.

```
ipv6 ecmp max-path <1-8>
```

Example

Enable ECMP on the switch and configure the maximum number of ECMP paths to 4:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ipv6 ecmp enable
Switch:1(config)#ipv6 ecmp max-path 4
```

Variable Definitions

The following table defines parameters for the **ipv6 ecmp max-path** command.

Variable	Value
<1-8>	Specifies the maximum number of ECMP paths. The default is 1.

Disable IPv6 ICMP Multicast

On disabling the ICMP multicast processing, ICMPv6 Echo Request packets sent to IPv6 multicast addresses are dropped when they reach the control plane.

About This Task

Use this procedure to disable the IPv6 ICMP multicast on the global router. By default, the IPv6 ICMP multicast feature is enabled.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Disable IPv6 ICMP multicast, enter:

```
no ipv6 icmp echo-multicast-request
```

3. View the IPv6 ICMP multicast state:

```
show ipv6 global [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Configure IPv6 Fragmented ICMP Packet Filtering

About This Task

Use this task to enable Fragmented ICMP packet filtering on IPv6 network globally or on a specific VRF.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable Fragmented ICMP packet filtering:

```
ipv6 icmp drop-fragments
```

View IPv6 ICMP Statistics

View IPv6 ICMP statistics on an interface for ICMP messages sent over a particular interface.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. View IPv6 ICMP statistics:

```
show ipv6 interface icmpstatistics [gigabitethernet <slot/port[/sub-
port]> | loopback <1-256> | tunnel <1-2000> | vlan <1-4059> ] [vrf
WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View ICMP statistics:

```
Switch:1>show ipv6 interface icmpstatistics
=====
Icmp Stats
=====

Icmp stats for IfIndex = 192

IcmpInMsgs: 0
IcmpInErrors: 0
IcmpInDestUnreachs : 0
IcmpInAdminProhibs : 0
IcmpInTimeExcds : 0
IcmpInParmProblems : 0
IcmpInPktTooBigs : 0
IcmpInEchos : 0
IcmpInEchoReplies : 0
IcmpInRouterSolicits : 0
IcmpInRouterAdverts : 0
InNeighborSolicits : 0
InNbrAdverts : 0
IcmpInRedirects : 0
IcmpInGroupMembQueries : 0
IcmpInGroupMembResponses : 0
```

Variable Definitions

Use the data in the following table to use the **show ipv6 interface icmpstatistics** command

Variable	Value
<1-4059>	Shows ICMP statistics for the specific interface index. If you do not specify an interface index, the command output includes all IPv6 ICMP interfaces. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>gigabitethernet</i> { <i>slot/port[/sub-port]</i> }	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>loopback</i> <1-256>>	Identifies a loopback interface.
<i>tunnel</i> <1-2000>	Identifies a 6in4 tunnel ID.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

Enable Stateless Address Autoconfiguration

Enable IPv6 Stateless Address Autoconfiguration to generate addresses using a combination of locally available information and information advertised by routers.

The default is disabled.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable IPv6 autoconfiguration:
`ipv6 autoconfig`

Example

```
Switch:1(config)#ipv6 autoconfig
```

The following example shows a sample output for the **show ipv6 global** command.

```
Switch:1#show ipv6 global
  forwarding                : disable
  default-hop-cnt            : 64
  number-of-interfaces       : 2
  icmp-error-interval        : 1000
  icmp-error-quota           : 50
  icmp-unreach-msg           : disable
  icmp-addr-unreach-msg      : enable
  icmp-port-unreach-msg      : enable
  icmp-echo-multicast-request : enable
  static-route-admin-status  : enable
  alternative-route          : enable
  ecmp                       : disable
  ecmp-max-path              : 1
  source-route               : disable
  host-autoconfig            : enable
```

View IPv6 Default Routers

View the table of default routers learned from router advertisement messages.

A maximum of four routers are visible in the default routers list.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show IPv6 default routers:
`show ipv6 default-routers`

Example

```
Switch:1#show ipv6 default-routers
=====
                        Default Routers
NET ADDRESS                VLAN    LIFETIME    IS ACTIVE
-----
fe80::211:58ff:fe2b:fc00    mgmt      1778        YES
fe80::215:e8ff:fe6e:2800    mgmt      1657        NO
```

Configure an IPv6 Prefix List

Use IPv6 prefix lists to allow or deny specific IPv6 route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

About This Task

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create an IPv6 prefix list:

```
ipv6 prefix-list <WORD 1-64> <WORD 1-256> [ge <0-128>] [le <0-128>]
[id <1-2147483647>]
```

Use the same command to add additional prefixes to the list.

3. To rename the list:

```
ipv6 prefix-list <WORD 1-64> name <WORD 1-64>
```

4. Display the prefix list:

```
show ipv6 prefix-list <WORD 1-256> [vrf WORD<1-16>] [vrfids
WORD<1-512>] [WORD <1-64>]
```

Example

Create an IPv6 prefix list:

```
Switch:1<config>#ipv6 prefix-list list4 2001:DB8::/32 ge 32 le 32
```

To rename the list:

```
Switch:1<config>#ipv6 prefix-list list4 name list5
```


Variable Definitions

Use the data in the following table to use the **ipv6 prefix-list** command.

Variable	Value
<WORD 1-64>	Specifies the IPv6 prefix-list name.
<WORD 1-256>	Specifies the IPv6 prefix and length.
ge <0-128>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
id <1-2147483647>	Specifies the Prefix list ID.
le <0-128>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
name <WORD 1-64>	Names the prefix list. The default value is none.

Use the data in the following table to use the **show ipv6 prefix-list** command.

Variable	Value
<WORD 1-256>	Specifies the IPv6 prefix and length.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0-512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Clear IPv6 Statistics

Clear all IPv6 statistics if you do not require previous statistics.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Clear all the IPv6 statistics:
clear ipv6 statistics all [vrf WORD<1-16> | vrfids WORD<0-512>]
3. Clear interface statistics:
clear ipv6 statistics interface [general|icmp] [gigabitethernet {slot/port[/sub-port]}| loopback <1-256> | tunnel <1-2000> | vlan <1-4059>]
[vrf WORD<1-16> | vrfids WORD<0-512>]
4. Clear TCP statistics:
clear ipv6 statistics tcp [vrf WORD<1-16> | vrfids WORD<0-512>]
5. Enter the following command to clear UDP statistics:
clear ipv6 statistics udp [vrf WORD<1-16> | vrfids WORD<0-512>]

Variable Definitions

Use the information in the following table to use the **clear ipv6 statistics** commands.

Variable	Value
<i>gigabitethernet</i> {slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>loopback</i> <1-256>>	Identifies a loopback interface.
<i>tunnel</i> <1-2000>	Identifies a 6in4 tunnel ID.
<i>vlan</i> <1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>vrf</i> WORD<1-16>	Specifies the VRF name.
<i>vrfids</i> WORD<0-512>	Specifies the VRF ID.

View IPv6 Statistics on an Interface

View IPv6 statistics to view information about the IPv6 datagrams on an interface.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View statistics:

```
show ipv6 interface statistics [gigabitethernet {slot/port[/sub-port]}
| loopback <1-256> | tunnel <1-2000> | vlan <1-4059>] [vrf WORD<1-16>]
[vrfids WORD<0-512>]
```

Variable Definitions

Use the data in the following table to use the **show ipv6 interface statistics** command

Variable	Value
<i>gigabitethernet</i> {slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>loopback</i> <1-256>>	Identifies a loopback interface.
<i>tunnel</i> <1-2000>	Identifies a 6in4 tunnel ID.

Variable	Value
<code>vlan <1-4059></code>	Shows statistics for the specific interface index. If you do not specify an interface index, the command output includes all IPv6 interfaces. Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

Viewing IPv6 Connections Using CLI

Use the procedures in this section to view IPv6 connections using CLI.

View TCP and UDP Information

Perform this procedure to view the TCP and UDP configuration information for IPv6.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display IPv6 TCP connection information:
`show ipv6 tcp connections [vrf WORD<1-16> | vrfids WORD<0-512>]`
3. Display IPv6 TCP listener information for the specified IPv6 address:
`show ipv6 tcp listener [vrf WORD<1-16> | vrfids WORD<0-512>]`
4. Display IPv6 TCP properties
`show ipv6 tcp properties [vrf WORD<1-16> | vrfids WORD<0-512>]`
5. Display IPv6 TCP statistics
`show ipv6 tcp statistics [vrf WORD<1-16> | vrfids WORD<0-512>]`
6. Display IPv6 UDP information:
`show ipv6 udp endpoints [vrf WORD<1-16> | vrfids WORD<0-512>]`

Example

```
Switch:1>show ipv6 tcp connections
```

```
=====
                        TCP connection table info - GlobalRouter
=====
LOCALPORT  LOCALADDR      REMOTEPORT  REMOTEADDR      STATE
-----
21          0:0:0:0:0:0:0  0           0:0:0:0:0:0:0  listen
23          0:0:0:0:0:0:0  0           0:0:0:0:0:0:0  listen
```

```
Switch:1>show ipv6 tcp listener
```

```
=====
```

```

TCP listener table info - GlobalRouter
=====
LOCALPORT  LOCALADDR
-----
21          0:0:0:0:0:0:0:0
23          0:0:0:0:0:0:0:0
80          0:0:0:0:0:0:0:0
443         0:0:0:0:0:0:0:0
513         0:0:0:0:0:0:0:0

Switch:1#show ipv6 tcp properties
=====
TCP Global Properties - GlobalRouter
=====
RtoAlgorithm    constant
RtoMin          5002 milliseconds
RtoMax          60128 milliseconds
MaxConn         127

Switch:1#show ipv6 tcp statistics
=====
TCP Global Statistics - GlobalRouter
=====
ActiveOpens:    0
PassiveOpens:   5
AttemptFails:   0
EstabResets:    0
CurrEstab:      0
InSegs:         0
OutSegs:        0
RetransSegs:    0
InErrs:         0
OutRsts:        0
HcInSegs:       0
HcOutSegs:      0

Switch:1>show ipv6 udp endpoints
=====
UDP endpoint table info - GlobalRouter
=====
LOCALPORT  LOCALADDR  REMOTEPORT  REMOTEADDR  PROCESS  INSTANCE
-----
69          0:0:0:0:0:0:0:0
0           0:0:0:0:0:0:0:0
161         0:0:0:0:0:0:0:0
0           0:0:0:0:0:0:0:0
1220866096  0
1220867644  0

```

Troubleshoot IPv6 Connectivity Loss

If the switch experiences loss of IPv6 connectivity, use the following procedure to troubleshoot the issue.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```

2. Through the command line interface, make sure the required routes are in place and the corresponding neighbor entries are resolved (that is, in REACHABLE, PROBE, DELAY or STALE state).
3. INCOMPLETE neighbor state indicates a problem if the corresponding neighbor is used by some of the IPv6 routes. This applies to neighbor entries with link-local addresses.

**Note**

Global addresses are not normally used as next hops. Having a global IPv6 neighbor entry as INCOMPLETE does not usually lead to a connectivity issue.

4. If the corresponding route is not in place then this is a routing issue. If the neighbor is not present or is INCOMPLETE, then further debugging is needed on the network level (that is, the state of other nodes needs to be examined).
5. Disabling and re-enabling IPv6 on the VLAN often recovers connectivity.
6. Display the RSMLT and MLT status:

```
show ip rsmlt
```

```
show mlt
```

Make sure the RSMLT peer MAC is learned and the IST state is `ist`.

IPv6 Basic Configuration using EDM

Use the procedures in this section to configure IPv6 basics using EDM.

Enable the IPv6–Mode Boot Config Flag

Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits.

This flag is disabled by default. Use this procedure to enable (set to true) the IPv6-mode boot config flag.

When the IPv6-mode boot config flag is enabled, the maximum number of IPv4 routing table entries decreases. For scaling information, see *Fabric Engine Release Notes*

**Note**

This procedure does not apply to 4220 Series and 7830 Series.

.

Procedure

1. In the navigation pane, expand **Configuration > Edit**,
2. Select **Chassis**.
3. Select the **Boot Config** tab.
4. Select the **EnableIpv6Mode** check box.
5. Select **Apply**.

6. Save the configuration, and then reboot the switch for the change to the IPv6-mode boot config flag to take effect.

Configure IPv6 Globally

Global configuration includes the following:

- IPv6 alternative routes: To avoid traffic interruption, enable alternative routes globally on the switch, to replace the best route with the next-best route if the best route becomes unavailable. By default, this feature is enabled.

Before You Begin

Change the VRF instance as required to configure IPv6 globally on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Globals** tab.
4. In the **DefaultHopLimit** box, enter the preferred number of hops before packets drop.
5. To enable ICMP network unreachable messages, select **IcmpNetUnreach**.
6. In the **IcmpErrorInterval** box, enter the preferred interval for sending ICMPv6 error messages.
7. In the **IcmpErrorQuota** box, enter the preferred number of ICMP error messages that the system can send during the ICMP error interval.
8. To enable IPv6 multicast, select **IcmpMulticastRequestEnable**.
9. To enable IPv6 ICMP address unreachable messages, select **IcmpAddrUnreach**.
10. To enable IPv6 ICMP port unreachable messages, select **IcmpPortUnreach**.
11. To enable IPv6 Fragmented ICMP packet filtering, select **IcmpDropFragmentsEnable**.
12. To enable IPv6 autoconfiguration, select **Autoconfig**.
13. To enable IPv6 static routes globally, select **StaticRouteGlobalAdminEnabled**.
14. To enable IPv6 Source Routing, select **SourceRouteEnable**.
15. To clear all IPv6 static routes, select **RouteStaticClear**.
16. To enable IPv6 alternative routes, select **AlternativeRouteEnable**.
17. To configure IPv6 ECMP globally, select **enable** or **disable**, in the **EcmpEnable** option box.
18. In the **EcmpMaxPath** box, enter the preferred number of equal-cost paths.
19. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Forwarding	Configures whether this switch is an IPv6 router with respect to the forwarding of datagrams received by, but not addressed to, the switch. Select forwarding for the switch to act as a router for IPv6 traffic. Select notForwarding to not act as a router for IPv6 traffic. The default is forwarding . You must enable forwarding to use Telnet or Ping with IPv6.
DefaultHopLimit	Configures the hop limit. The default is 64.
Interfaces	Shows the number of interfaces.
IfTableLastChange	Shows the date of the last interface table change.
IcmpNetUnreach	Enables ICMP network unreachable messages. The default is disabled.
IcmpErrorInterval	Configures the interval (in milliseconds) for sending ICMPv6 error messages. The default is 1000. An entry of 1 seconds results in no sent ICMPv6 error messages.
IcmpErrorQuota	Configures the number of ICMP error messages that the system can send during the ICMP error interval. A value of zero specifies not to send any. The default value is 50.
IcmpMulticastRequestEnable	Globally enables or disables the IPv6 ICMP echo multicast request feature. This is enabled by default.
IcmpAddrUnreach	Enables or disabled ICMP address unreachable messages. This is enabled by default.
IcmpPortUnreach	Enables or disables ICMP port unreachable messages. This is enabled by default.
IcmpDropFragmentsEnable	Enables or disables the Fragmented ICMP packet filtering globally. The default is disabled.
Autoconfig	Enables or disables stateless address autoconfiguration. This is disabled by default.

Name	Description
StaticRouteGlobalAdminEnabled	Enables IPv6 static routes globally. If you disable static routes globally, the system removes all enabled static routes from the RTM and does not add new static routes to the RTM. The default is enabled.
RouteStaticClear	Clears all IPv6 static routes.
SourceRouteEnable	Globally enables or disables the IPv6 Source Routing feature. It is disabled by default.
PrefixListTableSize	Displays the prefix list table size.
RoutePrefTableSize	Displays the route preference table size.
AlternativeRouteEnable	Globally enables or disables the IPv6 alternative route feature. By default, this feature is enabled.
EcmpEnable	Enables or disables the IPv6 ECMP globally. By default, it is disabled.
EcmpMaxPath	Globally configures the maximum number of ECMP paths. You cannot configure this feature unless ECMP is enabled globally.

View ICMP Statistics

View ICMP statistics for ICMP configuration information.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select **Interfaces** tab.
4. Select the interface on which you want to view the ICMP statistics.
5. Select **ICMPstats** option from the menu.

ICMP stats Field Descriptions

Use the data in the following table to use the **ICMP stats** tab.

Name	Description
InMsgs	Specifies the total number of ICMP messages which the entity received. Note: This counter includes all those counted by icmpInErrors.
InErrors	Specifies the number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
InDestUnreachs	Specifies the number of ICMP Destination Unreachable messages received by the interface.
InAdminProhibs	Specifies the number of ICMP destination unreachable/communication administratively prohibited messages received by the interface.
InTimeExcds	Specifies the number of ICMP Time Exceeded messages by the interface.
InParmProblems	Specifies the number of ICMP Parameter Problem messages received by the interface.
InPktTooBigs	Specifies the number of ICMP Packet Too Big messages received by the interface.
InEchos	Specifies the number of ICMP Echo (request) messages received by the interface.
InEchoReplies	Specifies the number of ICMP Echo Reply messages received by the interface.
InRouterSolicits	Specifies the number of ICMP Router Solicit messages received by the interface.
InRouterAdvertisements	Specifies the number of ICMP Router Advertisement messages received by the interface.
InNeighborSolicits	Specifies the number of ICMP Neighbor Solicit messages received by the interface.
InNeighborAdvertisements	Specifies the number of ICMP Neighbor Advertisement messages received by the interface.
InRedirects	Specifies the number of ICMP Redirect messages received by the interface.
InGroupMembQueries	Specifies the number of ICMPv6 Group Membership Query messages received by the interface.

Name	Description
InGroupMembResponses	Specifies the number of ICPv6 Group Membership Response messages received by the interface.
InGroupMembReductions	Specifies the number of ICMPv6 Group Membership Reduction messages received by the interface.
OutMsgs	Specifies the total number of ICMP messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
OutErrors	Specifies the number of ICMP messages which this interface did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IPv6 to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
OutDestUnreachs	Specifies the number of ICMP Destination Unreachable messages sent by the interface.
OutAdminProhibs	Specifies the number of ICMP destination unreachable/communication administratively prohibited messages sent.
OutTimeExcds	Specifies the number of ICMP Time Exceeded messages sent by the interface.
OutParmProblems	Specifies the number of ICMP Parameter Problem messages sent by the interface.
OutPktTooBigs	Specifies the number of ICMP Packet Too Big messages sent by the interface.
OutEchos	Specifies the number of ICMP Echo (request) messages sent by the interface.
OutEchoReplies	Specifies the number of ICMP Echo Reply messages sent by the interface.
OutRouterSolicits	Specifies the number of ICMP Router Solicitation messages sent by the interface.
OutRouterAdvertisements	Specifies the number of ICMP Router Advertisement messages sent by the interface.
OutNeighborSolicits	Specifies the number of ICMP Neighbor Solicitation messages sent by the interface.

Name	Description
OutNeighborAdvertisements	Specifies the number of ICMP Neighbor Advertisement messages sent by the interface.
OutRedirects	Specifies the number of Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
OutGroupMembQueries	Specifies the number of ICMPv6 Group Membership Query messages sent.
OutGroupMembResponses	Specifies the number of ICMPv6 Group Membership Response messages sent.
OutGroupMembReductions	Specifies the number of ICMPv6 Group Membership Reduction messages sent.

Configure an IPv6 Interface

You must configure an IPv6 interface for a VLAN or brouter port before you can assign an IPv6 address to the interface.

Before You Begin

- You must configure a VLAN before you can give the VLAN an interface identifier or an IPv6 address. The switch supports port-based, protocol-based, and MAC-source-based VLANs.

For information about how to configure VLANs, see [VLAN Configuration using EDM](#) on page 4043.

- Change the VRF instance as required to configure an IPv6 interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

You can also configure an IPv6 interface for a brouter port through the **Edit > Port > IPv6** navigation path, and for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can create both types of interfaces.

Procedure

- In the navigation pane, expand **Configuration > IPv6**.
- Select **IPv6**.
- Select the **Interfaces** tab.
- Select **Insert**.
- In the **Interface** field, select **Port** or **VLAN**.
- Select a port or VLAN.
- Select **OK**.
- Select the **AdminStatus** field to activate the interface.
- Configure the remaining parameters as required.

10. Select **Insert**.

11. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Type	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500. Different hardware platforms support different MTU values.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.

Name	Description
MacOffset	Requests a particular MAC for an IPv6 VLAN. You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.
RSMLTEnable	Shows whether RSMLT is enabled on the interface. The default value is disabled (false).
ProcessRedirect	Shows whether ND Redirect messages processing is enabled or disabled on this interface. The default value is disabled (false).

View IPv6 Statistics for an Interface

View IPv6 statistics to view information about the IPv6 datagrams on an interface.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Interfaces** tab.
4. Select an interface.
5. Select **IfStats**.
6. (Optional) Select one or more values, and then select on the type of graph to graph the data.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.

Name	Description
InReceives	Shows the total number of input datagrams received by the interface, including those received in error.
InHdrErrors	Shows the number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, and errors discovered in processing the IPv6 options.
InTooBigErrors	Shows the number of input datagrams that could not be forwarded because their size exceeded the link MTU of the outgoing interface.
InNoRoutes	Shows the number of input datagrams discarded because no route could be found to transmit them to their destination.

Name	Description
InAddrErrors	Shows the number of input datagrams discarded because the IPv6 address in the IPv6 header destination field was not a valid address to be received at this entity. This count includes invalid addresses, for example, ::0, and unsupported addresses, for example, addresses with unallocated prefixes. For entities which are not IPv6 routers and do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
InUnknownProtos	Shows the number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This counter is incremented at the interface to which these datagrams were addressed, which is not always the input interface for some of the datagrams.
InTruncatedPkts	Shows the number of input datagrams discarded because the datagram frame did not carry enough data.
InDiscards	Shows the number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded, for example, for lack of buffer space. This counter does not include datagrams discarded while awaiting re-assembly.
InDelivers	Shows the total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed which is not always the input interface for some of the datagrams.
OutForwDatagrams	Shows the number of output datagrams which this entity received and forwarded to their final destinations. In entities which do not act as IPv6 routers, this counter includes only those packets which were Source-Routed using this entity, and the Source-Route processing was successful. For a successfully forwarded datagram the counter of the outgoing interface is incremented.

Name	Description
OutRequests	Shows the total number of IPv6 datagrams which local IPv6 user-protocols (including ICMP) supplied to IPv6 in requests for transmission. This counter does not include datagrams counted in OutForwDatagrams .
OutDiscards	Shows the number of output IPv6 datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded, for example , for lack of buffer space. This counter includes datagrams counted in OutForwDatagrams if such packets met this (discretionary) discard criterion.
OutFragOKs	Shows the number of IPv6 datagrams that have been successfully fragmented at this output interface.
OutFragFails	Shows the number of IPv6 datagrams that have been discarded because they needed to be fragmented at this output interface but could not be.
OutFragCreates	Shows the number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
ReasmReqds	Shows the number of IPv6 fragments received which needed to be reassembled at this interface. This counter is incremented at the interface to which these fragments were addressed, which is not always the input interface for some of the fragments.
ReasmOKs	Shows the number of IPv6 datagrams successfully reassembled. This counter is incremented at the interface to which these datagrams were addressed, which is not always the input interface for some of the fragments.
ReasmFails	Shows the number of failures detected by the IPv6 re- assembly algorithm). This value is not necessarily a count of discarded IPv6 fragments because some algorithms can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed, which is not always the input interface for some of the fragments.

Name	Description
InMcastPkts	Shows the number of multicast packets received by the interface.
OutMcastPkts	Shows the number of multicast packets transmitted by the interface.

Configure an IPv6 Brouter Port Interface

You must configure an IPv6 interface for a brouter port before you can assign an IPv6 address to the interface.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select the **IPv6 Interface** tab.
5. Select **Insert**.
6. Enter the interface identifier.
7. Select the **AdminStatus** field to activate the interface.
8. Configure the remaining parameters as required.
9. Select **Insert**.
10. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Type	Shows the interface type.

Name	Description
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500. Different hardware platforms support different MTU values.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN. You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range.
RSMLTEnable	Shows whether RSMLT is enabled on the interface. The default value is disabled (false).
ProcessRedirect	Shows whether ND Redirect messages processing is enabled or disabled on this interface. The default value is disabled (false).

Configure an IPv6 VLAN Interface

You must configure an IPv6 interface for a VLAN before you can assign an IPv6 address to the interface.

Before You Begin

- You must configure a VLAN before you can give the VLAN an interface identifier or an IPv6 address. The switch supports port-based, protocol-based, and MAC-source-based VLANs.

For information about how to configure VLANs, see [VLAN Configuration using EDM](#) on page 4043.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IPv6**.
6. Select the **IPv6 Interface** tab.
7. Select **Insert**.
8. Enter the interface identifier.
9. Select the **AdminStatus** field to activate the interface.
10. Configure the remaining parameters as required.
11. Select **Insert**.
12. Select **Apply**.

IPv6 Interfaces Field Descriptions

Use the data in the following table to use the **IPv6 Interfaces** tab.

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Type	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.

Name	Description
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN. You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range. Different hardware platforms support different MAC offset ranges.
RSMLTEnable	Shows whether RSMLT is enabled on the interface. The default value is disabled (false).
ProcessRedirect	Shows whether ND Redirect messages processing is enabled or disabled on this interface. The default value is disabled (false).

Assign IPv6 Addresses to Interfaces

Assign IPv6 addresses to interfaces to configure IPv6 routing for the interface.

You can assign an IPv6 address to a VLAN or brouter port.

Before You Begin

Change the VRF instance as required to assign IPv6 addresses to interfaces on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

To create MLT and LAG interfaces with IPv6, you must configure VLAN-based connections and you cannot use brouter ports.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Addresses** tab.
4. Select **Insert**.
5. In the **Interface** field, Select **Port** or **VLAN**.
6. Select the interface.

7. Select **OK**.
8. Type the IPv6 address and prefix length.
9. Select **Insert**.
10. Select **Apply**.

Addresses Field Descriptions

Use the data in the following table to use the **Addresses** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
Addr	<p>Specifies the IPv6 address to which this entry applies.</p> <p>Important: If the IPv6 address exceeds 116 octets, the object identifiers (OIDs) of instances of columns in this row are more than 128 subidentifiers and you cannot use SNMPv1, SNMPv2c, or SNMPv3 to access them.</p>
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Type	Specifies the type of address. The default is unicast.
Origin	<p>Specifies the origin of the address. The following list shows the possible origins:</p> <ul style="list-style-type: none"> • other • manual • dhcp • linklayer • random
Status	<p>Specifies the status of the address, describing whether the address is used for communication. The following list shows the possible statuses:</p> <ul style="list-style-type: none"> • preferred (default) • deprecated • invalid • inaccessible • unknown • tentative • duplicate

Name	Description
Created	Specifies the sysUpTime of the creation of this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
LastChanged	Specifies the value of sysUpTime of the last modification to this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
ValidLifetime	Shows how long, in seconds, the address is valid.
PrefLifetime	Shows how long, in seconds, the address is in use.

Assign IPv6 Addresses to a Brouter Port Interface

Assign IPv6 addresses to interfaces to configure IPv6 routing for the interface.

You can assign an IPv6 address to a VLAN or brouter port.

Before You Begin

Change the VRF instance as required to assign IPv6 addresses to a brouter port interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

To create MLT and LAG interfaces with IPv6, you must configure VLAN-based connections and you cannot use brouter ports.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IPv6**.
3. Select the **Ipv6 Addresses** tab.
4. Select **Insert**.
5. Select **Insert**.
6. Select **Apply**.

IPv6 Addresses *Field Descriptions*

Use the data in the following table to use the **IPv6 Addresses** tab.

Name	Description
Interface	Specifies the port.
Addr	Specifies the IPv6 address to which this entry applies. Note: If the IPv6 address exceeds 116 octets, the object identifiers (OIDs) of instances of columns in this row are more than 128 subidentifiers and you cannot use SNMP-v1, SNMPv2c, or SNMPv3 to access them.
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Type	Specifies the type of address. The default is unicast.
Origin	Specifies the origin of the address.
Status	Specifies the status of the address, describing whether the address is used for communication.
Created	Specifies the sysUpTime of the creation of this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
LastChanged	Specifies the value of sysUpTime of the last modification to this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.

Assign an IPv6 Address to a VLAN

Assign an IPv6 address to a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN interface.
5. Select **IPv6**.
6. Select the **IPv6 Addresses** tab.

7. Select **Insert**.
8. Type the IPv6 address and length in the fields.
9. Select **Insert**.
10. Select **Apply**.

IPv6 Addresses Field Descriptions

Use the data in the following table to use the **IPv6 Addresses** tab.

Name	Description
Interface	Identifies the address to which the address is assigned.
Addr	Specifies an IP address that is associated with a VLAN.
AddrLen	Specifies the prefix address length value for this address.
Type	Specifies the type of address: either unicast or anycast.
Origin	Specifies the origin of the address as one of the following: <ul style="list-style-type: none"> • other • manual • dhcp • linklayer • random
Status	Shows the status of the address and if it can be used for communication.
Created	Shows the time this address entry was created.
LastChanged	Shows the time this address entry was last updated.

Create IPv6 Static Routes

Perform the steps in this task to:

- Create static routes for data traffic in either the GRT or a specific VRF context for any platform.
- Create static routes for a VRF associated with a Segmented Management Instance CLIP or VLAN (routing mode) interface.
- Configure an IPv6 discard static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.

Static routes for the management OOB, if supported, and management VLAN (non-routing mode) must use the Segmented Management Instance. The management CLIP can use the Segmented Management Instance or routes in the associated

VRF routing table manager (RTM). The management VLAN (routing mode) uses the Fabric Engine routing table or routes in the associated VRF.

Before You Begin

- Enable IPv6 forwarding.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. Not all parameters are configurable on non-default VRFs.
- Configure an IPv6 discard static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.
- If an IPv6 discard static route already exists for the destination, you must delete it before you add a regular IPv6 static route.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Globals** tab.
4. Select the **StaticRouteGlobalAdminEnabled** check box.
If you disable static routes globally, the system removes all enabled static routes from the RTM and does not add new static routes to the RTM. The default is enabled.
5. Select **Apply**.
6. Select the **Static Routes** tab.
7. Select **Insert**.
8. In **Dest**, type the IPv6 address.
For a default static route, type 0.0.0.0.
9. In **PrefixLength**, type the IPv6 address.
10. In **NextHop**, type the IPv6 address of the router through which the specified route is accessible.
11. (Optional) Select the **Discard** check box to create a discard static route.
The **NextHop** field autocompletes with
`ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff`.
12. Beside **Interface**, select **Port** or **Vlan** or **Tunnel**.
13. Select the interface, and then select **OK**.
14. In **Cost**, type a number for the distance.
15. Select **Enable**.
16. Select **Insert**.

Static Routes *Field Descriptions*

Use the data in the following table to use the **Static Routes** tab.

Name	Description
Dest	Specifies the IPv6 destination network address. The prefix value must match the value in the PrefixLength field.
PrefixLength	Specifies the number bits you want to advertise from the prefix. The prefix value must match the value in the Dest field. The range is 0 to 128.
NextHop	<p>Specifies the IPv6 address of the next hop on this route. You do not need to specify the next hop if the devices directly connect to one another. Configure the next hop if the two nodes do not share the same network prefix but reside on the same link.</p> <p>If the IPv6 next hop is intended for a discard static route, the field autocompletes with <code>ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff</code>. This value cannot be modified. This address indicates the traffic to this route is discarded.</p>
Discard	<p>Enables or disables IPv6 discard static routes.</p> <p>The default is disabled.</p>
Interface	Specifies the interface to which this entry applies. You must specify the port or VLAN if the next hop is a link-local address.
Cost	Specifies the cost or distance ratio to reach the destination for this node. The range is 1-65535. The default value is 1.
Name	Specifies the name for the static route.
Enable	Enables the static route on the port. The default value is enable.

Name	Description
Status	<p>Shows the status of the static route as one of the following:</p> <ul style="list-style-type: none">• notReachable: The route is not reachable and no neighbor request entry is built to resolved the next-hop. The system displays this status if no route or neighbor exists to reach the next-hop of the static route.• tryToResolve: The route is not reachable but a neighbor request entry is built to resolve the next-hop. The system displays this status if a local equivalent route exists in the system to reach the next-hop but the neighbor is not learned.• reachableNotInRtm: The static route is reachable but it is not in RTM. The system displays this status if the static route is reachable, but it is not the best among alternative static routes.• reachableInRtm: The static route is reachable and it is in RTM. The system displays this status if the static route is reachable, and it is the best among alternative static routes to be added into RTM.
Preference	<p>Specifies the routing preference of the destination IPv6 address. The range is 1-255. The default value is 5.</p>

Configure IPv6 Route Preferences

Change IPv6 route preferences to force the routing protocols to prefer one route over another. Configure route preferences to override default route preferences and give preference to routes learned for a specific protocol.

Before You Begin

Change the VRF instance as required to configure IPv6 route preferences on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task



Important

Changing route preferences is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. If you want to change default preferences for routing protocols, do so before you enable the protocols.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **RoutePref** tab.
4. In the **ConfiguredValue** column, double-click a parameter to change the preference for the given protocol.
5. Select **Apply**.

RoutePref Field Descriptions

Use the data in the following table to use the **RoutePref** tab.

Name	Description
DefaultValue	Specifies the default preference value for the specified protocol.
Protocol	Specifies the protocol name.
ConfiguredValue	Configures the preference value for the specified protocol.

View Route Information

View routes information to view the current configuration.

Before You Begin

Change the VRF instance as required to view route information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

IPv6 host routes created for the IPv6 local interfaces do not display in the routing table.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Routes** tab.

Routes Field Descriptions

Use the data in the following table to use the **Routes** tab.

Name	Description
Dest	Specifies the IPv6 destination network address. The prefix value must match the PrefixLength.
PfxLength	Specifies the number bits you want to advertise from the prefix. The prefix value must match the value in the Dest field.

Name	Description
Index	Specifies the unique value that identifies the route among the routes to the same network layer destination.
Interface	Specifies the interface to which this entry applies.
NextHop	Specifies the IPv6 address of the next hop on this route.
Protocol	Specifies the routing protocol, such as OSPF.
Metric	Specifies the metric assigned to this interface. The default value of the metric is the reference bandwidth or ifSpeed. The value of the reference bandwidth is configured by the rcOspfV3ReferenceBandwidth object. For more information about reference bandwidth, see Globals Field Descriptions on page 2896.
NextHopId	Identifier of the next-hop, hostname, or mac address.
Age	Specifies the number of seconds since the route was last updated or is last active.
Type	Specifies the type of route.
PathType	Specifies the type of path.
SrcVrfId	Specifies the source VRF instance.
Pref	Specifies the preference.

View IPv6 Default Routers

View the table of default routers learned from router advertisement messages.

A maximum of four routers are visible in the default routers list.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Default Routers** tab.

Default Routers Field Descriptions

Use the data in the following table to use the **Default Routers** tab.

Name	Description
Address	Specifies the learned router address for an IPv6 default routers entry.
IfIndex	Specifies the interface number for an IPv6 default routers entry.
Lifetime	Specifies the remaining router lifetime.
Active	Specifies if the default router is active for an IPv6 default routers entry.

Configure a Circuitless IPv6 Interface

Before You Begin

Change the VRF instance as required to configure a Circuitless IPv6 interface on a specific VRF instance.

About This Task

You can use a circuitless IPv6 (CLIPv6) interface to provide uninterrupted connectivity to your system.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Circuitless IP** tab.
4. Select **Insert**.
5. In the **Interface** field, assign a CLIP interface number.
6. Type the IPv6 address and prefix length.
7. Select **Insert**.

Circuitless IPv6 Field Descriptions

Use the data in the following table to use the **Circuitless IPv6** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
Addr	Specifies the IPv6 address to which this entry applies.

Name	Description
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.
Name Note:	Specifies the name assigned to the IPv6 CLIP address.

Configure a IPv6 Prefix List

Use IPv6 prefix lists to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. When there is a match, the route is used.

Before You Begin

- Change the VRF instance as required to configure a prefix list on a specific VRF instance.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Policy**.
3. In the **IPv6-Prefix List** tab, select **Insert**.
4. Edit the parameters as required.
5. Select **Insert**.

IPv6-Prefix list Field Descriptions

Use the data in the following table to use the **IPv6-Prefix list** tab.

Name	Description
Id	Specifies the list identifier. The range is 1 to 2147483647.
Prefix	Specifies the prefix IPv6 address.
PrefixMaskLen	Specifies the length of the prefix mask. You must enter the full 128-bit mask to exact a full match of a specific IPv6 address (for example, when creating a policy to match the next-hop).
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name can be from 1 to 64 characters in length.

Name	Description
MaskLenFrom	Specifies the lower bound on the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Specifies the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Configure an IPv6 Route Policy

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Policy**.
3. Select the **Route Policy** tab.
4. Select **Insert**.
5. Enter the appropriate information for your configuration in the Insert Route Policy dialog box.
6. Select **Insert**.

Route Policy Field Descriptions

Use the data in the following table to use the **Route Policy** tab.

Name	Description
Id	Specifies the ID of an entry in the Prefix list table.
SequenceNumber	Specifies a policy within a route policy group.
Name	Specifies the name of the policy. This command changes the name field for all sequence numbers under the given policy.
Enable	Indicates whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored. The default is disabled.
Mode	Specifies the action to take when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route). The default is permit.

Name	Description
MatchProtocol	Selects the appropriate protocol. If configured, matches the protocol through which the route is learned. This parameter applies only for RIP Announce purposes. The default is to enable all match protocols.
MatchNetwork	Specifies if the system matches the destination network against the contents of the specified prefix list.
MatchIpRouteSource	<p>Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.</p> <p>Select the ellipsis and choose from the list in the Match Route Source dialog box. You can select up to four entries. To clear an entry, use the ALT key.</p> <p>You can also change this field in the Route Policy tab of the Policy dialog box.</p>
MatchIpRouteDest	Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
MatchNextHop	<p>Specifies if the system matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.</p> <p>Select the ellipsis and choose from the list in the Match Next Hop dialog box. You can select up to four entries. To clear an entry, use the ALT key.</p>
MatchInterface	<p>Specifies if the system matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route.</p> <p>Select the ellipsis and choose from the list in the Match Interface dialog box. You can select up to four entries. To clear an entry, use the ALT key.</p>
MatchRouteType	<p>Configures a specific route type to match (applies only to OSPF routes).</p> <p>Externaltypel and Externaltypel2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes. The default is any.</p>

Name	Description
MatchMetric	Specifies if the system matches the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored. The default is 0.
MatchMetricTypeIsis	Specifies the match metric type field in the incoming IS-IS routes in accept policy.
MatchAsPath	Configures if the system matches the BGP autonomous system path. This parameter applies to BGP and overrides the BGP neighbor filter list information.
MatchCommunity	Filters incoming and outgoing updates based on a Community List. This parameter applies to BGP. The default is disable.
MatchCommunityExact	Indicates if the match must be exact (that is, all of the communities specified in the path must match). This parameter applies to BGP. The default is disabled.
MatchTag	Specifies a list of tags used during the match criteria process. This parameter applies to BGP.
MatchVrf	Identifies the source VRFs that leaks routes to the local VRF (applies only to RIP routes).
MatchLocalPref	Specifies the local preference value to be matched.
NssaPbit	Configures or resets the P bit in specified type 7 link state advertisements (LSA). By default, the Pbit is always configured in case the user configures the Pbit to a disable state for a particular route policy other than all type 7. LSAs associated with that route policy have the Pbit cleared. With this intact, not so stubby area (NSSA) area border router (ABR) does not perform translation of these LSAs to type 5. The default is enable.
SetRoutePreference	Configures a preference value. The default value is 0. If the default is configured, the global preference value is used. When configured to a value greater than zero, specifies the route preference value assigned to the routes that matches the policy. This feature applies to accept policies only.
SetMetricTypeInternal	Identifies the MED value for routes advertised to BGP numbers to the Interior Gateway Protocol (IGP) metric value. The value must be 0 or 1. The default is 0.

Name	Description
SetMetricTypeIsis	Sets the metric type IS-IS.
SetMetric	Configures the system to use the metric value for the route while announcing or redistributing. The default-import-metric is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or the default value is used (applies to IS-IS routes also). The default is 0.
SetMetricType	Configures the metric type for the routes to announce into the OSPF routing protocol that matches this policy. Applicable to OSPF protocol only. The default is type 2. This field is applicable only for OSPF announce policies. The default is type2.
SetNextHop	Configures the IP address of the next-hop router. This parameter applies to BGP. The default is 0.0.0.0.
SetInjectNetList	Configures the destination network of the route that matches this policy with the contents of the specified prefix list. Select the ellipsis and choose from the list in the Set Inject NetList dialog box.
SetMask	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.
SetAsPath	Indicates the AS path value to use whether the SetAsPathMode field is Tag or Prepend. This parameter applies to BGP. Note: Prepend is not applicable to an internal BGP (iBGP) peer with outbound route policy.
SetAsPathMode	Configures if the system converts the tag of a route into an AS path. This parameter applies to BGP. The mode is either Tag or Prepend tag. The value is applicable only while redistributing routes to BGP. The default is prepend. Note: Prepend does not apply to an iBGP peer with outbound route policy.
SetAutomaticTag	Enables the automatic tag feature. This parameter applies to BGP. The default is disable.

Name	Description
SetCommunityNumber	Configures the community number for BGP advertisements. This value can be a number (1 to 42949672000) or no-export or no-advertise.
SetCommunityMode	Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged. <ul style="list-style-type: none"> Unchanged—keeps the community attribute in the route path as it is. None—removes the community in the route path additive. Append—adds the community number specified in SetCommunityNumber to the community list attribute.
SetExtCommunity	Configures a BGP community.
SetExtCommunityMode	Configures the extended-community mode. The value can be append, unchanged, or overwrite. The default value is unchanged. <ul style="list-style-type: none"> append — creates another community string. unchanged — keeps the community attribute as it is. overwrite — changes the current value.
SetOrigin	Configures the origin for BGP to IGP, EGP, incomplete, or unchanged. If not configured, the system uses the route origin from the IP routing table (protocol). The default is unchanged.
SetLocalPref	Configures the local preference for BGP. The system uses this value during the route decision process for the BGP protocol. The default is 0.
SetOriginEgpAs	Indicates the remote autonomous system number for BGP. The default is 0.
SetWeight	Configures the weight value for the routing table for BGP. This field must be used with the match as-path condition. For BGP, this value overrides the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. The default is 0.
SetTag	Configures a tag. The default is 0. This parameter applies to BGP.
Ipv6SetNextHop	Specifies the address of the IPv6 next hop router.

Configure IPv6 Route Redistribution Policies

Before You Begin

- Ensure VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required.

About This Task

Configure route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF or BGP. Use a route policy to control the redistribution of routes. You can redistribute routes for the Global router (VRF 0) and within a user-defined VRF but not between different VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Policy**.
3. Select the **Route Redistribution** tab.
4. Select **Insert**.
5. Choose the protocol and route source.
6. Select **Enable**.
7. Choose the route policy to apply to the redistributed routes.
8. Configure other parameters as required.
9. Select **Insert**.

Route Redistribution Field Descriptions

Use the data in the following table to use the **Route Redistribution** tab.

Name	Description
DstVrfId	Specifies the destination VRF ID.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrfId	Specifies the source VRF ID.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution. The default is disabled.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine whether the system advertises a specific route to the specified protocol.

Name	Description
Metric	Specifies the metric announced in advertisements. The default is 0.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone. The default is type2.

Viewing IPv6 Connections using EDM

Use the procedures in this section to view IPv6 connections using EDM.

View TCP Global Information

View TCP and UDP information

Before You Begin

Change the VRF instance as required to view TCP global information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to view information about the handshake (SYN) configuration and the maximum number of TCP connections you can create on your system. When you initiate a TCP connection, both end points send handshake information to create the channel. The retransmission algorithm and fields display the configured timeout value and minimum and maximum retransmission times that your system uses to terminate a connection attempt that falls outside your specified parameters.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **TCP/UDP**.
3. Select the **TCP Globals** tab.

TCP Global Field Descriptions

Use the data in the following table to use the **TCP Global** tab.

Name	Description
RtoAlgorithm	Determines the timeout value used for retransmitting unacknowledged octets.
RtoMin	Displays the minimum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.

Name	Description
RtoMax	Displays the maximum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
MaxConn	Displays the maximum connections for the device.

View TCP Connections Information

Before You Begin

Change the VRF instance as required to view TCP connections information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to view the status of connections that traverse your switch.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **TCP/UDP**.
3. Select the **TCP Connections** tab.

TCP Connections Field Descriptions

Use the data in the following table to use the **TCP Connections** tab.

Name	Description
LocalAddressType	Displays the local address type as IPv6.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
RemAddressType	Displays the remote address type as IPv6.
RemAddress	Displays the IPv6 address for the remote TCP connection.
RemPort	Displays the remote port number for the TCP connection.

Name	Description
State	Displays an integer that represents the state for the connection: <ul style="list-style-type: none"> • closed • listen • synSent • synReceived • established • finWait1 • finWait2 • closeWait • lastAck(9) • closing • timeWait • deleteTCB
Process	Displays the process ID for the system process associated with the TCP connection.

View TCP Listeners Information

View TCP listener information.

Before You Begin

Change the VRF instance as required to view TCP listeners information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to view a detailed list of systems that are in the listening state.

When a connection is in the listen state an end point system is waiting for a returning handshake (SYN). The normal listening state should be very transient, changing all of the time. Two or more systems going to a common system in an extended listening state indicates the need for further investigation. End systems in an extended listening state can indicate a broken TCP connection or a DOS attack on a resource. This type of DOS attack, known as a SYN attack, results from the transmission of SYNs with no response to return replies. While many systems can detect a SYN attack, the TCP listener statistics can provide additional forensic information.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **TCP/UDP**.
3. Select the **TCP Listeners** tab.

TCP Listeners *Field Descriptions*

Use the data in the following table to use the **TCP Listeners** tab.

Name	Description
LocalAddressType	Displays the local address type as IPv6.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
Process	Displays the process ID for the system process associated with the TCP connection.

View UDP Endpoint Information

Before You Begin

Change the VRF instance as required to view UDP endpoint information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to display local and remote UDP activity.

Since UDP is a protocol used to establish connectionless network sessions, you need to monitor local and remote UDP activity and to know which applications are running over UDP. You can determine which applications are active by checking the port number. Processes are further identified with a UDP session to allow for the multiplexing of a port mapping for UDP.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **TCP/UDP**.
3. Select the **UDP Endpoints** tab.

UDP Endpoints *Field Descriptions*

Use the data in the following table to use the **UDP Endpoints** tab.

Name	Description
LocalAddressType	Displays the local address type as IPv6.
LocalAddress	Displays the local IPv6 address.
LocalPort	Displays the local port number.
RemoteAddressType	Displays the remote address type as IPv6.
RemoteAddress	Displays the remote IPv6 address.
RemotePort	Displays the remote port number.

Name	Description
Instance	Distinguishes between multiple processes connected to the UDP endpoint.
Process	Displays the ID for the UDP process.



IPv6 Neighbor Discovery

[Neighbor Discovery](#) on page 2260

[Host autoconfiguration](#) on page 2265

[IPv6 Router Advertisement Options for DNS Support](#) on page 2266

[Neighbor Discovery Configuration using CLI](#) on page 2266

[Neighbor Discovery Configuration using EDM](#) on page 2278

The following sections provide concepts and procedures to complete IPv6 neighbor discovery configuration.

Neighbor Discovery

IPv6 nodes (routers and hosts) on the same link use neighbor discovery (ND) to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services for IPv4 with the Address Resolution Protocol (ARP) and router discovery. In IPv6 ND performs a function similar to ARP (Address Resolution Protocol) in IPv4.

Hosts use ND to discover the routers in the network that you can use as the default routers, and to determine the link-layer address of neighbors attached to local links. Routers also use ND to discover neighbors and link-layer information. ND updates the neighbor database with valid entries, invalid entries, and entries migrated to various locations.

The ND protocol provides the following services:

- address and prefix discovery

Hosts determine the set of addresses that are on-link for the given link. Nodes determine which addresses or prefixes are locally reachable or remote with address and prefix discovery.

- router discovery

Hosts discover neighboring routers with router discovery. Hosts establish neighbors as default packet-forwarding routers.

- parameter discovery

Hosts and routers discover link parameters such as the link MTU or the hop-limit value placed in outgoing packets.

- address autoconfiguration

Hosts configure an address for an interface with address autoconfiguration.

- duplicate address detection

Hosts and nodes determine if an address is assigned to another router or a host.

- address resolution

Hosts determine link-layer addresses (MAC for Ethernet) of the local neighbors (attached on the local network), provided the IP address is known.

- next-hop determination

Hosts determine how to forward local or remote traffic with next-hop determination. The next hop can be a local or remote router.

- neighbor unreachability detection

Hosts determine if the neighbor is unreachable, and if address resolution must be performed again to update the database. For neighbors you use as routers, hosts attempt to forward traffic through alternative default routers.

- redirect

Routers inform the host of more efficient routes with redirect messages.

Neighbor discovery uses three components:

- host-router discovery

Host-router discovery performs the following functions:

- router discovery
 - prefix discovery
 - parameter discovery
 - address autoconfiguration
- host-host communication

Host-host communication performs the following functions:

- address resolution
- next-hop determination
- neighbor unreachability detection
- duplicate address detection

- route redirect



Note

When a neighbor transitions to the STALE state, to initiate Neighbor Unreachability detection (NUD), a duplicate copy of the traffic destined to this neighbor is sent to the switch Control Processor (CP) on a low priority queue (queue 0). The original packet is forwarded to this neighbor. After NUD is initiated, the hardware records are updated and the traffic is no longer sent to the CP. When a high rate of such traffic is sent to the CP, the switch can drop some of these packets due to the built-in CP rate limiting feature, which protects the CP from DOS attacks.

Use the command **show qos cosq-stats cpu-port** to view drop statistics on the CPU queue. This design does not result in loss of traffic.

Use the command **ipv6 nd reachable-time <0-3600000>** to increase the default value of 3000 milliseconds, which in turn delays the scenario of data path sending STALE neighbor destined packets to the CP.

As a best practice, configure a reachable time value of 180000 and retransmit interval of 5000.

ND messages

The following table compares the ICMP message types.

Table 168: IPv4 and IPv6 neighbor comparison

IPv4 Function	IPv6 Function	Description
ARP request message	Neighbor solicitation message	A node sends this message to determine the link-layer address of a neighbor or to verify that a neighbor is still reachable through a cached link-layer address. You can also use neighbor solicitations for duplicate address detection.
ARP reply message	Neighbor advertisement	A node sends this message either in response to a received neighbor solicitation message or to communicate a link-layer address change.
ARP cache	Neighbor cache	The neighbor cache contains information about neighbor types on the network.

Table 168: IPv4 and IPv6 neighbor comparison (continued)

IPv4 Function	IPv6 Function	Description
Gratuitous ARP	Duplicate address detection New VRRP master interface announcement	A host or node sends a request with its own IP address to determine if another router or host uses the address. If the sender receives a reply, then there is a device with a duplicate address. Both hosts and routers use this function. Gratuitous ARP can also be used to announce the new VRRP master interface so that all switches can adjust their MAC tables.
Router solicitation message (optional)	Router solicitation message (required)	The host sends this message after it detects a change in a network interface operational state. The message includes a request for routers to generate router advertisement immediately rather than at the scheduled time.
Router advertisement message (optional)	Router advertisement message (required)	Routers send this message to advertise their presence with various links and Internet parameters either periodically or in response to a router solicitation message. Router advertisements contain prefixes that you use for on-link determination or address configuration, and a suggested hop limit value.
Redirect message	Redirect message	Routers send this message to inform hosts of a better first hop for a destination.

Neighbor discovery cache

The neighbor discovery cache lists information about neighbors in the network and can contain the following types of neighbors:

- static: a configured neighbor
- local: a device on the local system
- dynamic: a discovered neighbor

The following table describes the states in the neighbor cache.

Table 169: Neighbor cache states

State	Description
Incomplete	<p>Address resolution is in progress and the system has not yet determined the link-layer address of the neighbor.</p> <p>The neighbor cache may also enter the Incomplete state when the switch cannot confirm subsequent reachability during the ND process for router neighbors. By contrast, the system deletes host neighbors, rather than enter the Incomplete state, if ND fails to confirm reachability.</p> <p>Tip:</p> <p>Router neighbors: when the R bit is set in the received neighbor advertisement</p> <p>Host neighbors: when the R bit is not set in the received neighbor advertisement</p>
Reachable	A node receives positive confirmation within the last reachable time period.
Stale	<p>Reachability of the neighbor is unknown.</p> <p>Until the system sends traffic to the neighbor, make no attempt to verify its reachability.</p>
Delay	<p>A time period longer than the reachable time period passes since the node received the last positive confirmation, and a packet was sent within the last DELAY_FIRST_PROBE_TIME period.</p> <p>If no reachability confirmation is received within the DELAY_FIRST_PROBE_TIME period after entering the DELAY state, neighbor solicitation is sent and the state changes to probe.</p>
Probe	Reachability confirmation is sought from the device every retransmit timer period.

The following events involve Layer 2 and Layer 3 interaction during processing and affect the neighbor cache:

- flushing the virtual LAN (VLAN) MAC
- removing a VLAN or brouter port
- performing an action on all VLANs
- removing a port from a VLAN
- removing a port from a spanning tree group (STG)
- removing a multilink trunk (MLT) group from a VLAN
- removing an MLT port from a VLAN
- removing an MLT port from an STG
- performing an action that disables a VLAN, such as removing all ports from a VLAN
- disabling a tagged port that is a member of multiple routable VLANs

Router discovery

IPv6 nodes discover routers on the local link with router discovery.

IPv6 nodes discover routers on the local link with router discovery.

Router advertisement

Configured interfaces on an IPv6 router send router-advertisement messages. Interfaces also send router advertisements in response to router-solicitation messages from IPv6 nodes on the link.

Router solicitation

An IPv6 host without a configured unicast address sends router solicitation messages.

Host autoconfiguration

The switch can automatically configure a host (node), and assign IPv6 addresses automatically. This process is called stateless address autoconfiguration (SLAAC). The neighbor discovery (ND) protocol performs autoconfiguration.

Stateless autoconfiguration enables serverless basic configuration of IPv6 nodes and renumbering from a mathematical perspective.

Stateless autoconfiguration uses the following equation:

Stateless autoconfiguration = *network prefix (router advertisement) + IPv6 interface identifiers*

Stateless autoconfiguration uses the network prefix information in the router advertisement and integrates this with the interface ID to form the node global address(es).



Note

The switch cannot autoconfigure an IPv6 address local to itself because IPv6 routers do not process router advertisements in the same manner as hosts. That is, routers check only for consistency in information advertised in IPv6 Router Advertisements on the same link.



Tip

You must manually assign all addresses/prefixes local to the switch.

Assuming an EUI-64 based interface ID is used, the IPv6 interface address is created from the 48-bit (6-byte) MAC address as follows:

1. EUI-64 hexadecimal digits 0xff-fe are inserted between the third and fourth bytes of the MAC address to obtain the EUI-64.
2. The universal or local bit, the second lower-order bit of the first byte of the MAC address, is complemented.

For example, the IPv6 identifier for host A uses the MAC address 00-AA-00-3F-2A-1C.

To automatically assign an address, the following occurs:

1. Convert to EUI-64 format

00-AA-00-FF-FE-3F-2A-1C

2. Complement the Universal/Local (U/L) bit.

The first byte in binary form is 00000000. When the seventh bit is complemented, it becomes 00000010 (0x02).

The result is 02-AA-00-FF-FE-3F-2A-1C or 2AA:FF:FE3F:2A1C.

Host A with MAC address 00-AA-00-3F-2A-1C, combined with network prefix 2001::/64 provided by router advertisement, uses an IPv6 address 2001::2AA:FF:FE3F:2A1C.

A host generates a link-local address with the prefix FE80 regardless of whether an IPv6 router is present or not.

The link-local address for a node with the MAC address 00-AA-00-3F-2A-1C is FE80::2AA:FF:FE3F: 2A1C.

The following list explains the states of an autoconfiguration address:

- Tentative: the address is being verified as unique (link-local address)
- Valid: an address from which unicast traffic can be sent and received; can be in one of two states—either preferred or deprecated
- Preferred: an address for which uniqueness was verified for unrestricted use; can be in one of three states—either tentative, preferred, or deprecated
- Deprecated: an address that remains valid but is withheld for new communication
- Invalid: an address for which a node can no longer send or receive unicast traffic

IPv6 Router Advertisement Options for DNS Support

IPv6 Router Advertisement options for DNS (RFC 8106) provide a mechanism for routers to send DNS information using ND messages so the host can acquire DNS configuration dynamically without requiring DHCPv6.

The following Router Advertisement options provide DNS configuration details:

- Recursive DNS Server (RDNSS)—includes one or more IPv6 addresses of recursive DNS servers.
- DNS Search List (DNSSL)—includes a list of domain names of DNS suffixes.

Neighbor Discovery Configuration using CLI

Use the procedures in this section for neighbor discovery configuration using CLI.

Configure an IPv6 Discovery Prefix

Configure the discovery prefixes to send in router advertisement.

About This Task

Hosts on the link use router advertisements to perform IPv6 autoconfiguration.

The discovery prefix controls which IPv6 addresses will be automatically configured, and for how long they are valid.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create neighbor discovery prefixes for an interface:

```
ipv6 nd prefix-interface WORD<0-255> [eui <1-3>] [no-advertise] [no-  
autoconfig] [no-onlink]
```

3. Modify an existing neighbor discovery prefix:

```
ipv6 nd prefix WORD<0-255> infinite [no-advertise] [preferred-life  
<0-4294967295>] [valid-life <0-4294967295>]
```

Example

Create a new neighbor discovery prefix:

```
Switch:1(config-if)#ipv6 nd prefix-interface fd48:bfb6:4c09:9499::1/64
```

Variable Definitions

Use the data in the following table to use the **ipv6 nd prefix** and **ipv6 nd prefix-interface** commands.

Variable	Value
<i>eui <1-3></i>	<p>Configures the EUI address. The values are:</p> <ul style="list-style-type: none"> • (1) EUI not used • (2) EUI with Universal/Local bit (U/L) complement enabled • (3) EUI used without U/L <p>Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both global- and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the prefix length is 64 or less. The default is EUI not used.</p> <p>If you select EUI not used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. If you select either EUI used with UL complement or EUI used without UL complement, an associated IPv6 address is created by concatenating the specified prefix with the EUI-64 interface ID.</p>
<i>infinite</i>	Configures the prefix valid lifetime so it never expires. The default is disabled, which means the prefix expires.
<i>no-advertise</i>	Removes the prefix from the neighbor advertisement. The default is disabled, which means the prefix is advertised.
<i>no-autoconfig</i>	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. The value is a 1-bit flag. The default is enabled.
<i>no-onlink</i>	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. The value is a 1-bit flag. The default is enabled.

Variable	Value
<code>preferred-life <0-4294967295></code>	<p>Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.</p> <p>The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.</p> <p>The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.</p>
<code>valid-life <0-0-4294967295></code>	<p>Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.</p> <p>A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address.</p> <p>The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.</p>
<code>WORD <0-255></code>	Specifies the IPv6 address and prefix.

Use the data in the following table to use the **interface** command.

Variable	Value
<code><1-4059></code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p>
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure Route Advertisement

Configure route advertisement in IPv6 for neighbor discovery (ND).

About This Task

IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the number of neighbor solicitation messages from duplicate address detection:

```
ipv6 nd dad-ns <0-600>
```

3. Configure the hop limit sent in router advertisements:

```
ipv6 nd hop-limit <0-255>
```

4. Enable managed address configuration (M-bit) on the router:

```
ipv6 nd managed-config-flag
```

5. Configure the MTU for router advertisements:

```
ipv6 nd mtu <0-9500>
```

6. Enable other stateful configuration (O-bit) on the router:

```
ipv6 nd other-config-flag
```

7. Configure the router lifetime included in router advertisement:

```
ipv6 nd ra-lifetime <0-9000>
```

8. Configure the neighbor reachable time:

```
ipv6 nd reachable-time <0-3600000>
```

9. Configure the time between neighbor solicitation messages:

```
ipv6 nd retransmit-timer <0-4294967295>
```

10. Configure the maximum time allowed between sending unsolicited multicast router advertisements:

```
ipv6 nd rtr-advert-max-interval <4-1800>
```

11. Configure the minimum time allowed between sending unsolicited multicast router advertisements:

```
ipv6 nd rtr-advert-min-interval <3-1350>
```

12. Enable periodic router advertisement messages:

```
ipv6 nd send-ra
```

13. (Optional) Configure an IPv6 Recursive DNS Server entry:

```
ipv6 nd dns server [suppress] [WORD<0-255> [lifetime <12-4294967294> |  
infinity] [sequence <1-8>]]
```

14. (Optional) Configure an IPv6 DNS Search List entry:

```
ipv6 nd dns search-list [suppress] [WORD<0-111> [lifetime  
<12-4294967294> | <infinity>] [sequence <1-8>]]
```

Example

Configure the maximum time between sending unsolicited router advertisements:

```
Switch:1(config-if)#ipv6 nd rtr-advert-max-interval 700
```

Configure the minimum time between sending unsolicited router advertisements:

```
Switch:1(config-if)#ipv6 nd rtr-advert-min-interval 500
```

Enable periodic router advertisement messages:

```
Switch:1(config-if)#ipv6 nd send-ra
```

Configure an IPv6 Recursive DNS Server entry:

```
Switch:1(config-if)#ipv6 nd dns server 2001:DB8:0:0:0:0:1 lifetime infinity sequence 3
```

Configure an IPv6 DNS Search List entry:

```
Switch:1(config-if)#ipv6 nd dns search-list mimi.com
```

Variable Definitions

Use the data in the following table to use the **ipv6 nd** commands.

Variable	Value
<i>dad-ns</i> <0-600>	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
<i>hop-limit</i> <0-255>	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a hop-limit value. The default is 64.

Variable	Value
<i>managed-config-flag</i>	<p>Enables the system to configure the M-bit, or managed address configuration flag, in the router advertisements</p> <p>When set, the M-bit flag indicates that addresses are available through DHCPv6.</p> <p>If the M flag is set, the O flag is redundant because DHCPv6 returns all available configuration information.</p> <p>If neither the M nor O flags are set, no information is available through DHCPv6.</p> <p>The default is disabled.</p>
<i>mtu <0-9500></i>	<p>Shows the MTU value sent in router advertisements on this interface.</p> <p>A value of zero indicates that the system sends no MTU options.</p> <p>The default is 0.</p>
<i>other-config-flag</i>	<p>Enables the O-bit, or other stateful configuration, flag in the router advertisement.</p> <p>Other stateful configuration autoconfigures received information without addresses.</p> <p>When set, the O-bit flag indicates that other configuration information is available through DHCPv6; for example, DNS-related information or information about other servers within the network.</p> <p>If neither the M nor O flags are set, no information is available through DHCPv6.</p> <p>The default is disabled.</p>
<i>ra-lifetime <0-9000></i>	<p>Specifies a value placed in the router lifetime field of router advertisements sent from this interface.</p> <p>This value must be either 0, or 4 to 9000 seconds.</p> <p>A value of zero indicates that the system is not a default router.</p> <p>The default is 1800.</p>
<i>reachable-time <0-3600000></i>	<p>Specifies a value (in milliseconds) placed in the router advertisement message sent by the router.</p> <p>The value zero means unspecified (by this system).</p> <p>Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event.</p> <p>The default is 0.</p>

Variable	Value
<i>retransmit-timer</i> <0-4294967295>	Specifies a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. The default is 0.
<i>rtr-advert-max-interval</i> <4-1800>	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The default is 600.
<i>rtr-advert-min-interval</i> <3-1350>	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The default is 200.
<i>send-ra</i>	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is enabled.

Use the data in the following table to use the **ipv6 nd dns** command.

Variable	Value
<i>lifetime</i> <12-4294967294> <i>infinity</i> >	Specifies the time in which IP addresses can be used for name resolution. The default value is 1800.
<i>sequence</i> <1-8>	Specifies the order of priority in which the RDNSS or DNSSL entries are listed in Router Advertisement options. Priority 1, listed first in the Router Advertisement option list is the highest priority while priority 8 is listed last and is the lowest priority. The default value is 8.
<i>suppress</i>	Suppresses RDNSS or DNSSL information in the Router Advertisement message.
<i>WORD</i> <0-255>	Specifies the IPv6 Server address.
<i>WORD</i> <0-111>	Specifies the IPv6 DNS domain name.

Use the data in the following table to use the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

System Interface Values Versus Advertised Values

There are differences in the relationship between the system interface values and advertised values related to Neighbor Discovery (ND). The information in this section describes differences and similarities and provides examples for important IPv6 interface and IPv6 ND commands.

Comparison of Default Values per Interface and as Advertised

The following table compares the default behavior of values per interface and advertised values.

Default values per interface	Default advertised values
hop-limit 64	hop-limit 64
mtu 1500	mtu 0 (unspecified)
reachable-time 30000 ms	reachable-time 0 (unspecified)
retransmit-timer 1000 ms	retransmit-timer 0 (unspecified)

What Happens When You Change the per Interface Value and the Advertised Value?

When you change per-interface values from default to non-default values, the system changes the advertised values to match the interface values.

For example, when you enter the `ipv6 interface mtu 1300` command the values become

- interface mtu 1300
- advertised mtu 1300

Then, when you enter the `show ipv6 nd interface` command, the system marks the mtu value with an (i) which signifies that the ND advertised value is inherited from the interface configuration.

Example: Changing both values

```
Switch:1(config-if)#ipv6 interface mtu 1300
show ipv6 nd interface GigabitEthernet 1/1
=====
Port Ipv6 Nd
=====
=====
IFID BTR RTR- MAX- MIN-  LIFE- MANAG  OTHER DAD-NS MTU  HOP  REACH- RETRANS-
      ADV INT  INT  TIME  TIME  LIMIT TIME  TIME  ABLE  MIT
      FLAG  CONF
-----
320  1/1 True 600  200    0    False False  1   1300(i) 64(d) 0(d)  0(d)
Note: (s) = Set, (d) = Default, (i) = inherit
-----
```

What Happens When You Change the per Interface Value but do not Change the Advertised Value?

To change the per-interface value from the default value to a non-default value but retain the advertised value of 0 (unspecified), you must enter two commands.

For example, to set the reachable-time to 60000 but retain the advertised value of the reachable-time parameter at 0, enter the following commands:

```
ipv6 interface reachable-time 6000
```

```
ipv6 nd reachable-time 0
```

When you enter the `show ipv6 nd interface` command, the system marks the reachable-time value with an (s) to signify that this value is explicitly set by the ND configuration.

Example: Changing only the per interface value

```
Switch:1(config-if)#ipv6 interface reachable-time 60000
```

```
Switch:1(config-if)#ipv6 nd reachable-time 0
```

```
Switch:1(config-if)#
show ipv6 nd interface GigabitEthernet 1/1
=====
Port Ipv6 Nd
=====
=====
IFID BTR RTR- MAX- MIN-  LIFETIME MANAG  OTHER DAD-NS MTU  HOP  REACH- RETRANS-
      ADV INT  INT  CONF  FLAG  LIMIT TIME  TIME  ABLE  MIT
-----
320  1/1 True 600  200    0    False False  1   1500(i) 64(d) 0(s)  0(d)
Note: (s) = Set, (d) = Default, (i) = inherit
-----
-
```

Configure the Neighbor Cache

Configure the address translation table used to map IPv6 addresses to physical addresses. You can manually add static neighbors to the cache.



Note

IPv6 static neighbors are not supported on SMLT.

About This Task

Neighbor cache in IPv6 is similar to the IPv4 Address Resolution Protocol (ARP) table.

The neighbor cache is a set of entries for individual neighbors to which traffic was recently sent.

You make entries on the neighbor on-link unicast IP address, including information such as the link-layer address.

A neighbor cache entry contains information used by the Neighbor Unreachability Detection algorithm, including the reachability state, the number of unanswered probes, and the time the next Neighbor Unreachability Detection event is scheduled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a static neighbor:

```
ipv6 neighbor WORD<0-128> port {slot/port[sub-port]} mac  
<0x00:0x00:0x00:0x00:0x00:0x00> [vlan <1-4059> ]
```

When you create a static neighbor, it always remains in the reachable state. This differs from the general neighbor cache behavior where, among other things, timers and neighbor unreachability detection events can be generated.

Example

Create a static neighbor:

```
Switch:1(config)#ipv6 neighbor 3000::3 port 1/11 mac 00-1A-4B-8A-FB-6B
```

Variable Definitions

Use the data in the following table to use the **ipv6 neighbor** command.

Variable	Value
<i>mac</i> <code><0x00:0x00:0x00:0x00:0x00:0x00></code>	Specifies the MAC address.
<i>{slot/port[/sub-port]}</i>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>vlan</i> <code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>WORD</i> <code><0-128></code>	Specifies the IPv6 address in hexadecimal colon format.

View Cached Destination Information

View the destination cache to see next-hop addresses for destinations.

The destination cache is only populated or updated when IPv6 packets originate locally on the central processor of the switch.

The main purpose of the destination cache is to store, on a per-destination basis, the dynamic Path MTU value currently used when transmitting packets from the local system to the remote destination.

The system uses the PMTU value to calculate how many bytes can fit into an individual packet before fragmentation should be applied.

About This Task

The command output shows the following information:

- the IPv6 destination address
- the IPv6 address for the next hop to the destination
- the path maximum transmission unit (MTU) for the destination
- the time, in seconds, since an ICMPv6 packet-too-big message was received

Not all parameters are available in non-default VRFs.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the destination cache for all interfaces:

```
show ipv6 dcache [vrf WORD<1-16> | vrfids WORD<0-512>]
```
3. View the destination cache for a brouter port:

```
show ipv6 dcache gigabitethernet {slot/port[/sub-port]}
```
4. View the destination cache for a specific tunnel ID:

```
show ipv6 dcache tunnel <1-2000>
```
5. View the destination cache for a VLAN:

```
show ipv6 dcache vlan <1-4059>
```
6. Clear the destination cache:

```
clear ipv6 dcache [gigabitethernet {slot/port[/sub-port]}][tunnel <1-2000>][vlan <1-4059>] [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Example

```
Switch:1(config-if)#show ipv6 dcache
```

```
=====
```

```
IPv6 Destination Cache Information - GlobalRouter
```

```
=====
```

Destination Address	NEXT HOP	VID/BID/TID	IF_TYPE	IF_DATA	PMTU	PMTU_AGE
ff02:0:0:0:0:0:1	0:0:0:0:0:0:0	V-22	real	-	1500	0

```
=====
```

```
1 out of 1 Total Num of Destinaton Cache Entries displayed.
```

Variable Definitions

Use the data in the following table to use the **show ipv6 dcache** and **clear ipv6 dcache** commands.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>srbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

Neighbor Discovery Configuration using EDM

Use the procedures in this section for neighbor discovery configuration using EDM.

Configure an IPv6 Discovery Prefix

Before You Begin

Change the VRF instance as required to configure an IPv6 discovery prefix on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to configure the discovery prefixes to send in router advertisement. Hosts on the link use router advertisements to perform IPv6 autoconfiguration. The discovery prefix controls what IPv6 addresses to autoconfigure and how long they are valid.

You can also configure an IPv6 interface for a brouter port through the **Edit > Port > IPv6** navigation path, and for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can configure both types of interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Discovery Prefix** tab.
4. Select **Insert**.
5. Beside the **Interface** field, select **Port** or **VLAN**.
6. Select a port or VLAN.
7. Select **OK**.
8. Specify the prefix and prefix length.
9. Select **Insert**.
10. Select **Apply**.

Discovery Prefix Field Descriptions

Use the data in the following table to use the **Discovery Prefix** tab.

Name	Description
Interface	Shows a read-only value that indicates an IPv6 interface. For the brouter port, it is the ifindex of the port and, in the case of the VLAN, it is the ifindex of the VLAN.
Prefix	Configures the prefix to create an IPv6 prefix entry as either advertised or suppressed.
PrefixLen	Configures the mask to create an IPv6 address in the IPv6 interface table.

Name	Description
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
ValidLifetime	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000. A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address. The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
PreferedLifetime	Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800. The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address. The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.
Infinite	Configures the prefix valid lifetime so it never expires. The default is false.
OnLinkFlag	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.

Name	Description
AutoFlag	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.
AddressEui	Configures the EUI address. Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both global and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used. If you select eui-not-used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. If you select either eui-used-with-ul-complement or eui-used-without-ul-complement, an associated IPv6 address is created by concatenating the specified prefix with the EUI-64 interface ID.
NoAdvertise	Configures if the prefix is included in the router advertisement. Select true to not include the prefix in the router advertisement. The default is false.

Configure an IPv6 Discovery Prefix Port

About This Task

Perform this procedure to configure the discovery prefixes to send in router advertisement. Hosts on the link use router advertisements to perform IPv6 autoconfiguration. The discovery prefix controls what IPv6 addresses to autoconfigure and how long they are valid.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IPv6**.
3. Select **Discovery Prefix**.
4. Select **Insert**.
5. Select **OK**.
6. Specify the prefix and prefix length.
7. Select **Insert**.
8. Select **Apply**.

IPv6 Discovery Prefix Field Descriptions

Use the data in the following table to use the **IPv6 Discovery Prefix** tab.

Name	Description
Interface	Shows a read-only value that indicates an IPv6 interface. For the brouter port, it is the ifindex of the port and, in the case of the VLAN, it is the ifindex of the VLAN.
Prefix	Configures the prefix to create an IPv6 prefix entry as either advertised or suppressed.
PrefixLen	Configures the mask to create an IPv6 address in the IPv6 interface table.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
ValidLifetime	Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000. A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address. The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.

Name	Description
PreferredLifetime	<p>Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.</p> <p>The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.</p> <p>The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.</p>
Infinite	Configures the prefix valid lifetime so it never expires. The default is false.
OnLinkFlag	Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.
AutoFlag	Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.
AddressEui	<p>Configures the EUI address. Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both global and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used.</p> <p>If you select eui-not-used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. If you select either eui-used-with-ul-complement or eui-used-without-ul-complement, an associated IPv6 address is created by concatenating the specified prefix with the EUI-64 interface ID.</p>
NoAdvertise	Configures if the prefix is included in the router advertisement. Select true to not include the prefix in the router advertisement. The default is false.

Configure an IPv6 Discovery Prefix on a VLAN

About This Task

Perform this procedure to configure the discovery prefixes to send in router advertisement. Hosts on the link use router advertisements to perform IPv6 autoconfiguration. The discovery prefix controls what IPv6 addresses to autoconfigure and how long they are valid.

Procedure

1. In the navigation pane, expand **VLAN > VLANs**.
2. Select the **Basic** tab.
3. Select an interface row.
4. Select **IPv6**.
5. Select the **IPv6 Discovery Prefix** tab.
6. Select **Insert**.
7. Specify the prefix and prefix length.
8. Configure the remaining parameters, as required.
9. Select **Insert**.
10. Select **Apply**.

IPv6 Discovery Prefix Field Descriptions

Use the data in the following table to use the **IPv6 Discovery Prefix** tab.

Name	Description
Interface	Shows a read-only value that indicates an IPv6 interface. For the brouter port, it is the ifindex of the port and, in the case of the VLAN, it is the ifindex of the VLAN.
Prefix	Configures the prefix to create an IPv6 prefix entry as either advertised or suppressed.
PrefixLen	Configures the mask to create an IPv6 address in the IPv6 interface table.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.

Name	Description
ValidLifetime	<p>Configures the valid lifetime in seconds that indicates the length of time this prefix is advertised. The default is 2592000.</p> <p>A valid lifetime is the length of time of the preferred and depreciated state of an autoconfiguration address.</p> <p>The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.</p>
PreferedLifetime	<p>Configures the preferred lifetime in seconds that indicates the length of time this prefix is advertised. The default value is 604800.</p> <p>The preferred lifetime is the length of time for the tentative, preferred, and depreciated state of an autoconfiguration address.</p> <p>The preferred lifetime value must be less than the valid lifetime value. If you must configure the valid lifetime value to a value lower than the current preferred lifetime, you must lower the preferred lifetime value first.</p>
Infinite	<p>Configures the prefix valid lifetime so it never expires. The default is false.</p>
OnLinkFlag	<p>Configures the prefix for use when determining if a node is online. This value is placed in the L-bit field in the prefix information option. It is a 1-bit flag. The default is true.</p>
AutoFlag	<p>Configures the prefix for use as the autonomous address configuration. This value is placed in the autoflag field in the prefix information option. It is a 1-bit flag. The default is true.</p>

Name	Description
AddressEui	Configures the EUI address. Use an EUI-64 interface ID in the address field. If enabled, use EUI, or use EUI-64 and the complement Universal/Local (U/L) bit. This operation provides for both global and link-local addresses. After you create the entry, you cannot modify this value. This value is valid for use only when the PrefixLength is 64 or less. The default is eui-not-used. If you select eui-not-used, this configuration creates an IPv6 ND prefix but no associated IPv6 address on the router. If you select either eui-used-with-ul-complement or eui-used-without-ul-complement, an associated IPv6 address is created by concatenating the specified prefix with the EUI-64 interface ID.
NoAdvertise	Configures if the prefix is included in the router advertisement. Select true to not include the prefix in the router advertisement. The default is false.

Configure Route Advertisement

Before You Begin

Change the VRF instance as required to configure route advertisement on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to configure route advertisement in IPv6 for neighbor discovery (ND). IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.



Note

You only use the ND level configuration when you want to create advertised values that differ from the interface values for reachable-time, retransmit-timer, mtu, or hop-limit.

You can also configure an IPv6 interface for a brouter port through the **Edit > Port > IPv6** navigation path, and for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 navigation path where you can configure both types of interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.

2. Select **IPv6**.
3. Select the **Route Advertisement** tab.
4. Double-click a parameter to change the current value.
You cannot modify the parameters in gray shading.
5. Select **Apply**.

Route Advertisement *Field Descriptions*

Use the data in the following table to use the **Route Advertisement** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
SendAdverts	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
MaxInterval	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 4 seconds and 1800 seconds. The default is 600.
MinInterval	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 3 seconds and 0.75 x max-interval. The default is 200.
ReachableTime	Shows a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
RetransmitTime	Shows a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.

Name	Description
DefaultLifeTime	Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between 4 and 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
ManagedFlag	Enables the system to configure the M-bit or managed address configuration in the router advertisements. The default is false.
DadNsNum	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	Shows the MTU value sent in router advertisements on this interface. A value of zero indicates that the system sends no MTU options.
OtherConfigFlag	Enables the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. The default is disabled.

Configure Route Advertisement on an IPv6 Interface for a Brouter Port

About This Task

Perform this procedure to configure route advertisement in IPv6 for neighbor discovery (ND). IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.



Note

You only use the ND level configuration when you want to create advertised values that differ from the interface values for reachable-time, retransmit-timer, mtu, or hop-limit.

You can also configure an IPv6 interface for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IPv6**.
3. Select the **Route Advertisement** tab.
4. Double-click a parameter to change the current value.
You cannot modify the parameters in gray shading.
5. Select **Apply**.

Route Advertisement *Field Descriptions*

Use the data in the following table to use the **Route Advertisement** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
SendAdverts	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
MaxInterval	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 4 seconds and 1800 seconds. The default is 600.
MinInterval	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 3 seconds and 0.75 x max-interval. The default is 200.
ReachableTime	Shows a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.

Name	Description
RetransmitTime	Shows a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
DefaultLifeTime	Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between 4 and 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
ManagedFlag	Enables the system to configure the M-bit or managed address configuration in the router advertisements. The default is false.
DadNsNum	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	Shows the MTU value sent in router advertisements on this interface. A value of zero indicates that the system sends no MTU options.
OtherConfigFlag	Enables the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. The default is disabled.

Configure Route Advertisement on an IPv6 Interface for a VLAN

About This Task

Perform this procedure to configure route advertisement in IPv6 for neighbor discovery (ND). IPv6 nodes on the same link use ND to discover link-layer addresses and to obtain and advertise various network parameters and reachability information. ND combines

the services provided by Address Resolution Protocol (ARP) and router discovery for IPv4.

**Note**

You only use the ND level configuration when you want to create advertised values that differ from the interface values for reachable-time, retransmit-timer, mtu, or hop-limit.

Procedure

1. In the navigation pane, expand **VLAN > VLANs**.
2. Select the **Basic** tab.
3. Select an interface row, and then select **IPv6**.
4. Select **Route Advertisement**.
5. Double-click a parameter to change the current value.
6. Select **Apply**.

Route Advertisement *Field Descriptions*

Use the data in the following table to use the **Route Advertisement** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
SendAdverts	Specifies whether the router sends periodic router advertisements and responds to router solicitations on this interface. The default is True.
UseDefaultVal	Specifies one included value to use the default value, or use all bits to configure all options to their default value.
MaxInterval	Specifies the maximum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 4 seconds and 1800 seconds. The default is 600.
MinInterval	Specifies the minimum interval (in seconds) at which the transmission of route advertisements occurs on this interface. The value must be between 3 seconds and 0.75 x max-interval. The default is 200.

Name	Description
ReachableTime	Shows a value (in milliseconds) placed in the router advertisement message sent by the router. The value zero means unspecified (by this system). Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
RetransmitTime	Shows a value (in milliseconds) placed in the retransmit timer field in the router advertisement message sent from this interface. The value zero means unspecified (by this system). The value configures the amount of time the system waits for the transmission to occur. You cannot modify this parameter; use the Interfaces tab to change the value for the interface.
DefaultLifeTime	Specifies a value placed in the router lifetime field of router advertisements sent from this interface. This value must be either 0 or between 4 and 9000 seconds. A value of zero indicates that the system is not a default router. The default is 1800.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
ManagedFlag	Enables the system to configure the M-bit or managed address configuration in the router advertisements. The default is false.
DadNsNum	Specifies the number of neighbor solicitation messages for duplicate address detection (DAD). A value of 0 disables the DAD process on this interface. A value of 1 sends one advertisement without retransmissions.
LinkMTU	Shows the MTU value sent in router advertisements on this interface. A value of zero indicates that the system sends no MTU options.
OtherConfigFlag	Enables the O-bit (other stateful configuration) in the router advertisement. Other stateful configuration autoconfigures received information without addresses. The default is disabled.

Name	Description
RdnssSupress	Suppresses IPv6 Recursive DNS Server (RDNSS) information in the Router Advertisement message.
DnsslSupress	Suppresses IPv6 DNS Search List (DNSSL) information in the Router Advertisement message.

Configure the Neighbor Cache

Before You Begin

Change the VRF instance as required to configure neighbor cache on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform the following procedure to configure the address translation table used to map IPv6 addresses to physical addresses. You can manually add static neighbors to the cache.



Note

IPv6 static neighbors are not supported on SMLT.

Neighbor cache in IPv6 is similar to the IPv4 Address Resolution Protocol (ARP) table. The neighbor cache is a set of entries for individual neighbors to which traffic was sent recently. You make entries on the neighbor on-link unicast IP address, including information such as the link-layer address. A neighbor cache entry contains information used by the Neighbor Unreachability Detection algorithm, including the reachability state, the number of unanswered probes, and the time the next Neighbor Unreachability Detection event is scheduled.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Neighbors** tab.
4. Select **Insert**.
5. Beside the **Interface** field, select **Port** or **Port in Vlan**.
6. Select a port or VLAN.
7. Configure the remaining parameters as required.
8. Select **Insert**.
9. Select **Apply**.

Neighbors Field Descriptions

Use the data in the following table to use the **Neighbors** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
NetAddress	Specifies the IP address of the media-dependent physical address.
PhyAddress	Specifies the MAC address.
Interface	Specifies a physical port ID or an MLT port ID.
LastUpdated	Specifies the value of sysUpTime of the last modification to this entry. If the entry was created prior to the last reinitialization of the local management subsystem, the object contains a zero value.
Type	Specifies the mapping type from manually configured entries. While the selection of either dynamic, static, or local is allowed; static is currently the only valid selection.
State	Specifies the Neighbor Unreachability Detection state for the interface after the address mapping in this entry is used. If Neighbor Unreachability Detection is not in use (for example, for IPv4), this object is always unknown. The options include the following: <ul style="list-style-type: none"> • reachable: confirmed reachability • stale: unconfirmed reachability • delay: waiting for reachability confirmation before entering the probe state • probe: actively probing • invalid: an invalidated mapping • unknown: state cannot be determined. • incomplete: address resolution is being performed
BMac	Specifies the backbone MAC address.
Cvid	Specifies the customer VID.

Configure an IPv6 Recursive DNS Server Entry

Before You Begin

Configure one or more IPv6 Router Advertisement prefix options.

About This Task

Perform the following procedure to add an IPv6 Recursive DNS Server (RDNSS) entry.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select **Reverse DNS Servers**.
4. Select **Insert**.
5. Enter the appropriate information for your configuration in the Insert Reverse DNS Servers dialog box.
6. Select **Insert**.

Reverse DNS Servers Field Definitions

Use the data in the following table to use the **Reverse DNS Servers** tab.

Name	Description
IfIndex	Specifies the interface number for an IPv6 Recursive DNS Server (RDNSS) entry.
Sequence	Specifies the order of priority in which the RDNSS entries are listed in Router Advertisement options. Priority 1, listed first in the Router Advertisement option list is the highest priority while priority 8 is listed last and is the lowest priority. The default value is 8.
Lifetime	Specifies the length of time in seconds to use the Reverse DNS server. The default value is 1800. The 4294967295 value represents infinity.
Server	Displays the IPv6 Server address.

Configure an IPv6 DNS Search List Entry

Before You Begin

Configure one or more IPv6 Router Advertisement prefix options.

About This Task

Perform the following procedure to add an IPv6 DNS Search List (DNSSL) entry.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select **DNS Search List**.
4. Enter the appropriate information for your configuration in the Insert DNS Search List dialog box.
5. Select **Insert**.

DNS Search List Field Definitions

Use the data in the following table to use the **DNS Search List** tab.

Name	Description
IfIndex	Specifies the interface number for an IPv6 DNS Search List (DNSSL) entry.
Sequence	Specifies the order of priority in which the DNSSL entries are listed in Router Advertisement options. Priority 1, listed first in the Router Advertisement option list is the highest priority while priority 8 is listed last and is the lowest priority. The default value is 8.
Lifetime	Specifies the time in which the DNSSL IP addresses can be used for name resolution. The default value is 1800.
Domain	Specifies the IPv6 DNS domain name.

View Cached Destination Information

Before You Begin

Change the VRF instance as required to view cached destination information on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to view the next-hop addresses in the destinations cache.

The destination cache updates or populates when IPv6 packets locally originate on the central processor of the switch. The main purpose of the destination cache is to store, on a per-destination basis, the dynamic Path MTU value currently used when transmitting packets from the local system to the remote destination. The Path MTU value itself is used to calculate how many bytes can fit into an individual packet before fragmentation should be applied.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Destination Cache** tab.

Destination Cache Field Descriptions

Use the data in the following table to use the **Destination Cache** tab.

Name	Description
DestAddr	Shows the IPv6 destination address.
Interface	Shows the interface number that is used to reach the destination.
NextHop	Shows the IPv6 address for the next hop to the destination.
IfType	Specifies the interface type (tunnel, VLAN, or brouter) or virtual circuit (VRRP, RSMLT).
IfData	Displays additional information about virtual circuits. For instance, for a VRRP or RSMLT the virtual router ID displays. If the interface type is tunnel, VLAN, or brouter, no additional information displays.
Pmtu	Shows the path maximum transmission unit (MTU) for the destination.
PmtuAge	Shows the time, in seconds, since an ICMPv6 packet too big message was received.



IPv6-in-IPv4 Tunneling

[Tunneling](#) on page 2298

[Tunneling Configuration using CLI](#) on page 2300

[Tunneling Configuration using EDM](#) on page 2303

[IPv6 Tunnel Configuration Example](#) on page 2308

Table 170: IPv6-in-IPv4 tunnels product support

Feature	Product	Release introduced
IPv6-in-IPv4 tunnels	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Not supported on 5320-16P-2MXT-2X
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Tunneling

Tunneling provides a mechanism to transfer IPv6 traffic through an IPv4-only network.

How Tunneling Works

IPv6 tunneling encapsulates IPv6 packets in IPv4 packets for delivery across an IPv4 infrastructure.

At the tunnel source, or head end, the system encapsulates an IPv6 packet into an IPv4 packet and sends it to the remote tunnel destination.

The tunnel destination strips the IPv4 packet header and forwards the original IPv6 packet further into an IPv6 cloud.

These types of tunnels are called dual-stack tunnels because they support both IPv4 and IPv6.

Manually Configured Tunnels

Manually configured tunnels can provide communication between two isolated IPv6 domains over an IPv4 network.

Manually configured tunnels are point-to-point.

You can configure tunnel endpoints to create a point-to-point connection between two isolated IPv6 domains by configuring IPv6 and IPv4 addresses at each end of the tunnel.



Note

The router or host at the source and destination ends of the tunnel must support both IPv4 and IPv6 protocol stacks.



Caution

Ensure that all single-homed point-to-point traffic ingresses and egresses a configured tunnel. Otherwise the traffic is dropped.

IPv6 reachability enables tunnel forwarding but tunnel operational status depends on the IPv4 reachability of the tunnel endpoint.

The IPv4 tunnel endpoint configuration must be symmetrical; that is, if you configure a tunnel with a source of 10.10.10.1 and a destination of 11.11.11.1 from switch A, then Switch B must have a source of 11.11.11.1 and a destination of 10.10.10.1.

Tunnel interfaces are logical point-to-point interfaces.

You can enable dynamic routing when you enable a routing protocol, for example OSPFv3, on the tunnel interfaces.

Unicast Routing Protocols: Detect Link Loss and Redirect IPv6 Route Information

There is no explicit signaling protocol applied to IPv6-in-IPv4 configured tunnels (refer to RFC 4213).

Therefore, if the remote endpoint of a tunnel that terminates several Layer 3 hops away in the network fails, the local state of the tunnel remains active even though the endpoint has failed.

However, you can enable unicast routing protocols over tunnels, for example OSPFv3. These unicast routing protocols introduce their own protocol-specific signaling and, when a unicast routing protocol is present over the tunnel link, the routing protocol can detect link loss and re-direct the IPv6 route information to use an alternate, reachable nexthop.

Operational Events That Trigger Tunnel State Transition

The switch must be able to locally detect operational events that can trigger a tunnel state transition.

These events include:

- deletion of local IPv4 interface
- change or loss of the IPv4 route to the remote tunnel endpoint
- change in the nexthop of the IPv4 route to the remote tunnel endpoint
- loss of the ARP entry for the nexthop router that is used to reach the IPv4 tunnel endpoint

Tunnels and MTU

You cannot configure the MTU for tunnels.

The default MTU value for tunnels is 1280.

Packets are forwarded through the tunnel using the line card network processing units (NPUs) only. Since the packets are not forwarded through the central processing unit (CPU) they do not impact the CPU load.

Tunnels and BGP+

You must configure an IPv6 tunnel and static routes on BGP+ peers when you use BGP+. For more information on IPv6 tunnel configuration for BGP+, see [IPv6 Tunnel Configurations for BGP+](#) on page 654.

Restrictions

The following list identifies tunnel configuration restrictions.

- You cannot configure IPv6 CLIP addresses for IPv6-in-IPv4 tunnels. Also, you cannot configure an IPv6 CLIP interface as the source or destination endpoint of an tunnel.
- You cannot configure SMLT on the switch terminating a tunnel.
- Termination of tunnels on vIST peers is not supported.

Tunneling Configuration using CLI

Use the procedures in this section to configure Tunnel using CLI.

Configure a Tunnel

Configure a tunnel for IPv6 VLANs or brouter ports to communicate through an IPv4-only network. Create a point-to-point connection between the two isolated IPv6 devices by configuring the tunnel endpoints.

Do not create tunnels in a native IPv6 network.

Before You Begin

- The router or host at the source and destination of the tunnel must support both IPv4 and IPv6 protocol stacks.

About This Task

Manual tunnels are point-to-point, so you configure both source and destination addresses. You must configure both IPv6 and IPv4 addresses for both source and destination devices. The IPv6 addresses must represent the same network, for example 6666::1/96 and 6666::2/96.

Tunnel interfaces are automatically configured with a link-local address in the format fe80::<local_ipv4_source_address>.

You cannot configure the maximum transmission unit (MTU) for tunnels. The default MTU value for tunnels is 1280.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

- Create a tunnel:

```
ipv6 tunnel <1-2000> source {A.B.C.D} address WORD<0-46> destination {A.B.C.D}
```

Example

Create tunnel 2:

```
Switch:1(config)#ipv6 tunnel 2 source 11.11.11.1 address 3000:0:0:0:0:0:1/64 destination 12.12.12.2
```

Variable Definitions

Use the data in the following table to use the **ipv6 tunnel** command.

Variable	Value
<1-2000>	Configures the ID for the tunnel.
address WORD<0-46>	Assigns an IPv6 address and prefix to the tunnel.
destination {A.B.C.D}	Configures the address of the remote endpoint of the tunnel.
source {A.B.C.D}	Configures the address of the local endpoint of the tunnel.

View Tunnel Interfaces

View tunnel interfaces to verify the current configuration and operational status of IPv6 tunnels.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Show IPv6 tunnel information:
show ipv6 tunnel [<1-2000>] [detail] [local {A.B.C.D}] [remote {A.B.C.D}]

Example

```
Switch:1#show ipv6 tunnel detail
```

```
=====
                        Tunnel Interface Information
=====
ID          LOCAL ADDRESS  REMOTE ADDRESS  OPER STATUS  TYPE
-----
2           211.1.55.2          44.1.55.1      active       manual
1           211.1.55.2          44.1.55.43     active       manual
210         211.1.60.2           47.1.60.1      active       manual

3 out of 3 Total number of entries displayed.

=====
                        Address Information
=====
IPV6 ADDRESS/PREFIX LENGTH      TYPE    ORIGIN    STATUS
-----
43:210:0:0:0:0:0:2/64           UNICAST MANUAL    PREFERRED
fe80:0:0:0:0:0:d301:3702/64     UNICAST LINKLAYER PREFERRED
44:211:0:0:0:0:0:2/64           UNICAST MANUAL    PREFERRED
fe80:0:0:0:0:0:d301:3702/64     UNICAST LINKLAYER PREFERRED
```

Variable Definitions

Use the data in the following table to use the **show ipv6 tunnel** command.

Variable	Value
<1-2000>	Shows information for a specific tunnel ID.
detail	Shows detailed address information for the tunnel.
local {A.B.C.D}	Shows information for a specific local address (the local endpoint of the tunnel).
remote {A.B.C.D}	Shows information for a specific remote address (the remote endpoint of the tunnel).

Modify Tunnel Hop Limits

Modify tunnel hop limits to update hop-limit values on previously configured tunnels.

About This Task

The tunnel hop limit configures the value of the time-to-live (TTL) for IPv4 packets.

Procedure

- 1. Enter Global Configuration mode:
enable

configure terminal
- 2. Modify the hop limit:
ipv6 tunnel <1-2000> hop-limit <0-255>

Example

Modify the hop limit for tunnel ID 5:

```
Switch:1(config)#ipv6 tunnel 5 hop-limit 200
```

Variable Definitions

Use the data in the following table to use the **ipv6 tunnel** command.

Variable	Value
<0-255>	Configures the maximum number of hops in the tunnel. The default value is 255.
<1-2000>	Specifies the tunnel ID.

Tunneling Configuration using EDM

Use the procedures in this section to configure Tunnel using EDM.

Configure a Tunnel

Configure a tunnel for IPv6 VLANs or brouter ports to communicate through an IPv4-only network. Create a point-to-point connection between the two isolated IPv6 devices by configuring the tunnel endpoints.

Do not create tunnels in a native IPv6 network.

Before You Begin

- The router or host at the source and destination of the tunnel must support both IPv4 and IPv6 protocol stacks.

About This Task

Manual tunnels are point-to-point, so you configure both source and destination addresses. You must configure both IPv6 and IPv4 addresses for both source and destination devices. The IPv6 addresses must represent the same network, for example 6666::1/96 and 6666::2/96.

Tunnel interfaces are automatically configured with a link-local address in the format fe80::<local_ipv4_source_address>.

You cannot configure the maximum transmission unit (MTU) for tunnels. The default MTU value for tunnels is 1280.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Tunnel**.
3. Select the **Tunnel Config** tab.
4. Select **Insert**.
5. Beside **LocalAddress**, select the button, and then select the IPv4 address for the local VLAN or brouter port.
6. In **RemoteAddress**, type the IPv4 address for the destination VLAN or brouter port.
7. In **ID**, type a number to represent the tunnel.
8. In **IPv6AddressAddr**, type the IPv6 address for the tunnel VLAN or brouter port.
9. In **IPv6AddressPrefixLength**, type the number of bits to advertise in the IPv6 address.
10. Select **Insert**.

Tunnel Config *Field Descriptions*

Use the data in the following table to use the **Tunnel Config** tab.

Name	Description
AddressType	Shows the address type over which the tunnel encapsulates packets.
LocalAddress	Configures the address of the local endpoint of the tunnel.
RemoteAddress	Configures the address of the remote endpoint of the tunnel.
EncapsMethod	Configures the tunnel mode, which is manual for manually configured tunnels.
ID	Configures the ID for the tunnel.
IfIndex	Shows the value of ifIndex that corresponds to the tunnel interface. A value of 0 indicates that the interface index has not yet been assigned. The system displays this field only on the Tunnel Config tab.
Ipv6AddressAddr	Specifies the IPv6 address for the local VLAN or brouter port. The system displays this field only on the Insert Tunnel Config dialog box.
Ipv6AddressPrefixLength	Specifies the number of bits to advertise in the IPv6 address. The system displays this field only on the Insert Tunnel Config dialog box.

Modify Tunnel Hop Limits

Modify tunnel hop limits to update hop-limit values on previously configured tunnels.

Use this procedure to modify the hop limits for multiple tunnels simultaneously.

About This Task

The tunnel hop limit configures the value of the time-to-live (TTL) for IPv4 packets.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Tunnel**.
3. Select the **Tunnel Interface** tab.
4. Double-click the **HopLimit** value to modify the information as required.
5. Select **Apply**.

Tunnel Interface Field Descriptions

Use the data in the following table to use the **Tunnel Interface** tab.

Name	Description
Index	Identifies the tunnel interface internally. The value is derived from the tunnel ID.
EncapsMethod	Displays the encapsulation method for the tunnel: manual for manually configured tunnels and 6to4 for automatically configured tunnels.
HopLimit	Configures the maximum number of hops in the tunnel. The default value is 255.
Security	Indicates the type of security on the tunnel interface.
TOS	Displays the method used to configure the high 6 bits (the differentiated services codepoint) of the IPv4 type of service (TOS) or IPv6 traffic class in the outer IP header. A value of -1 indicates that the bits are copied from the payload header. A value of -2 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB module. A value from 0 to 63 indicates that the bit field is configured to the indicated value.
FlowLabel	Displays the method used to configure the IPv6 flow label value. This object is not required where AddressType indicates the tunnel is not over IPv6. A value of -1 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB. Any other value indicates that the flow label field is configured to the indicated value.

Name	Description
AddressType	Displays manual for a manually configured tunnel, or sixToFour for autoconfigured tunnels.
LocalIpAddress	Identifies the local endpoint address of the tunnel.
RemoteIpAddress	Identifies the remote endpoint of the tunnel.
EncapsLimit	Displays the maximum number of additional encapsulations permitted for packets undergoing encapsulation at this node. A value of -1 indicates that no limit exists, except as a result of the packet size.

Modify Tunnel Hop Limits for a Specific Tunnel

Modify tunnel hop limits to update hop-limit values on previously configured tunnels.

Use this procedure to modify the hop limits for a specific tunnel interface.

About This Task

The tunnel hop limit configures the value of the time-to-live (TTL) for IPv4 packets.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Tunnel**.
3. Select the **Tunnel Config** tab.
4. Select the tunnel row.
5. Select **Tunnel Interface**.
6. Double-click the **HopLimit** value to modify the information as required.
7. Select **Apply**.

Tunnel Interface Field Descriptions

Use the data in the following table to use the **Tunnel Interface** tab.

Name	Description
Index	Identifies the tunnel interface internally. The value is derived from the tunnel ID.
EncapsMethod	Displays the encapsulation method for the tunnel: manual for manually configured tunnels and 6to4 for automatically configured tunnels.
HopLimit	Configures the maximum number of hops in the tunnel. The default value is 255.

Name	Description
Security	Indicates the type of security on the tunnel interface.
TOS	Displays the method used to configure the high 6 bits (the differentiated services codepoint) of the IPv4 type of service (TOS) or IPv6 traffic class in the outer IP header. A value of -1 indicates that the bits are copied from the payload header. A value of -2 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB module. A value from 0 to 63 indicates that the bit field is configured to the indicated value.
FlowLabel	Displays the method used to configure the IPv6 flow label value. This object is not required where AddressType indicates the tunnel is not over IPv6. A value of -1 indicates that a traffic conditioner is invoked and more information can be available in a traffic conditioner MIB. Any other value indicates that the flow label field is configured to the indicated value.
AddressType	Displays manual for a manually configured tunnel, or sixToFour for autoconfigured tunnels.
LocalInetAddress	Identifies the local endpoint address of the tunnel.
RemoteInetAddress	Identifies the remote endpoint of the tunnel.
EncapsLimit	Displays the maximum number of additional encapsulations permitted for packets undergoing encapsulation at this node. A value of -1 indicates that no limit exists, except as a result of the packet size.

View IPv6 Addresses on a Tunnel

View a tunnel for IPv6 addresses.

You can assign an IPv6 address to a VLAN or brouter port.

About This Task

To create MLT and LAG interfaces with IPv6, you must configure VLAN-based connections and you cannot use brouter ports.

You can also assign an IPv6 address through the **Edit > Port > IPv6** navigation path, and through the **VLAN > VLANs > Basic > IPv6** navigation path.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Tunnel**.
3. Select the **Tunnel Config** tab.
4. Select **IPv6 Address**.

Tunnel Config Field Descriptions

Use the data in the following table to use the **Tunnel Config** tab.

Name	Description
AddressType	Shows the address type over which the tunnel encapsulates packets.
LocalAddress	Configures the address of the local endpoint of the tunnel.
RemoteAddress	Configures the address of the remote endpoint of the tunnel.
EncapsMethod	Configures the tunnel mode, which is manual for manually configured tunnels.
ID	Configures the ID for the tunnel.
IfIndex	Shows the value of ifIndex that corresponds to the tunnel interface. A value of 0 indicates that the interface index has not yet been assigned. The system displays this field only on the Tunnel Config tab.

IPv6 Tunnel Configuration Example

This section shows examples of manually configured tunnels between router ports and VLANs.

Between Router Ports

The following figure shows the tunnel configuration between router ports.




Figure 163: Tunnel configuration between brouter ports

You must configure static routes, RIP, or OSPF on both the source (Switch 1) and destination (Switch 2) IPv4 interfaces to communicate on the IPv4 network. You must configure IPv4 addresses on the source and destination.

Configuring Switch 1

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 10 type port-mstprstp 0
vlan mlt 10 4
vlan members 10 1/1 portmember
interface vlan 10
ipv6 interface
ipv6 interface enable
ipv6 interface address 4000::0:0:0:0:0:0:1/64
exit
```

Create an IPv4 brouter port and enable OSPF on the port.

```
interface GigabitEthernet 1/30
brouter port 1/30 vlan 1000 subnet 172.21.80.1/255.255.255.0 mac-offset
6
```

Create the tunnel from the source to the destination.

```
ipv6 tunnel 1 source 172.21.80.1 address 2500::0000:0000:0000:0000:0000:0000:0001/64
destination 192.168.20.1
```

Configure a static route on the source.

```
ipv6 route 4000::0:0:0:0:0:0:2/64 cost 1 tunnel 1
```

Optionally, you can create an OSPFv3 interface through the tunnel.

```
router ospf ipv6-enable
router ospf
ipv6 tunnel 1 area 0.0.0.0
ipv6 tunnel 1 enable
exit
```

Configuring Switch 2

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 40 type port-mstprstp 0
vlan mlt 40 4
vlan members 40 1/2 portmember
interface vlan 40
ipv6 interface
ipv6 interface enable
ipv6 interface address 4000:0:0:0:0:0:0:2/64
exit
```

Create an IPv4 router port and enable OSPF on the port.

```
interface GigabitEthernet 1/30
router port 1/30 vlan 2000 subnet 192.168.20.1/255.255.255.0 mac-offset 6
```

Create the tunnel from the destination to the source.

```
ipv6 tunnel 1 source 192.168.20.1 address 2500:0000:0000:0000:0000:0000:0000:0002/64
destination 172.21.80.1
```

Configure a static route on the destination.

```
ipv6 route 4000:0:0:0:0:0:0:1/64 cost 1 tunnel 1
```

Optionally, you can create an OSPFv3 interface through the tunnel.

```
router ospf ipv6-enable
router ospf
ipv6 tunnel 1 area 0.0.0.0
ipv6 tunnel 1 enable
exit
```

Between VLANs

The following figure shows the tunnel configuration between VLANs.




Figure 164: Tunnel configuration between VLANs

You must configure static routes, and either RIP or OSPF on both the source (Switch 1) and destination (Switch 2) IPv4 interfaces to communicate on the IPv4 network. You must configure IPv4 addresses on the VLANs.

Configuring Switch 1

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 10 type port-mstprstp 0
vlan mlt 10 4
vlan members 10 1/1 portmember
interface vlan 10
ipv6 interface
ipv6 interface enable
ipv6 interface address 4000:0:0:0:0:0:1/64
exit
```

Create an IPv4 VLAN, add ports to the VLAN, and enable OSPF on the VLAN.

```
vlan create 20 type port-mstprstp 0
vlan mlt 20 4
vlan members 20 1/30 portmember
interface vlan 20
ip address 172.21.80.1 255.0.0.0
ip ospf enable
exit
```

Create the tunnel from the source to the destination.

```
ipv6 tunnel
1 source 172.21.80.1 address 2500:0000:0000:0000:0000:0000:0001/64
destination 192.168.20.1
```

Configuring Switch 2

Create an IPv6 VLAN and add ports to the VLAN.

```
vlan create 40 type port-mstprstp 0
vlan mlt 40 4
vlan members 40 1/2 portmember
interface vlan 40
ipv6 interface
ipv6 interface enable
ipv6 interface address 4000:0:0:0:0:0:2/64
exit
```

Create an IPv4 VLAN , add ports to the VLAN, and enable OSPF on the VLAN.

```
vlan create 30 type port-mstprstp 0
vlan mlt 30 4
vlan members 30 1/30 portmember
interface vlan 30
ip address 192.168.20.1 255.0.0.0
ip ospf enable
exit
```

Create the tunnel from the destination to the source.

```
ipv6 tunnel 1 source 192.168.20.1 address 2500:0000:0000:0000:0000:0000:0002/64
destination 172.21.80.1
```

Verification

Use the following show command to verify tunnel creation on the source device:

```
Switch:1(config)#show ipv6 tunnel 1 detail
```

```
=====
```

Tunnel Interface Information				
ID	LOCAL ADDRESS	REMOTE ADDRESS	OPER STATUS	TYPE
1	172.21.80.1	192.168.20.1	active	manual

```
1 out of 1 Total number of entries displayed.
```

```
=====
```

Address Information			
IPV6 ADDRESS	TYPE	ORIGIN	STATUS
2500:0:0:0:0:0:1	UNICAST	MANUAL	PREFERRED
fe80:0:0:0:0:0:ac15:5001	UNICAST	LINKLAYER	PREFERRED

```
2 out of 2 Total number of entries displayed.
```

Use the following show command to verify tunnel creation on the destination device:

```
Switch:1(config)#show ipv6 tunnel 1 detail
```

```
=====
```

Tunnel Interface Information				
ID	LOCAL ADDRESS	REMOTE ADDRESS	OPER STATUS	TYPE
1	192.168.20.1	172.21.80.1	active	manual

```
1 out of 1 Total number of entries displayed.
```

```
=====
```

Address Information			
IPV6 ADDRESS	TYPE	ORIGIN	STATUS
2500:0:0:0:0:0:2	UNICAST	MANUAL	PREFERRED
fe80:0:0:0:0:0:c0a8:1401	UNICAST	LINKLAYER	PREFERRED

```
2 out of 2 Total number of entries displayed.
```



Key Health Indicators (KHI)

[Key Health Indicators Using the CLI](#) on page 2313

[Key Health Indicators Using EDM](#) on page 2323

Table 171: Key Health Indicator product support

Feature	Product	Release introduced
Key Health Indicator (KHI)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The Key Health Indicators (KHI) feature provides a subset of condition information that allows for quick assessment of the overall operational state of the device.



Note

KHI was not designed to provide a comprehensive debugging solution. Instead, KHI identifies key information that could lead support personnel towards discovery of a specific failure. After the technician assesses the KHI information, further debugging is required to determine the specific reason for the fault.

Capture KHI information during normal operations to provide a baseline for support personnel when detecting fault situations.

Key Health Indicators Using the CLI

Use the procedures in this section to display Key Health Indicator (KHI) information using the CLI.

Display KHI Performance Information

Use the following commands to display KHI information about the performance of the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display buffer performance and utilization statistics:

```
show khi performance buffer-pool [{slot[-slot] [,...]}]
```
3. Show current utilization, 5-minute average utilization, and 5-minute high water mark with date and time of event:

```
show khi performance cpu [{slot[-slot] [,...]}]
```
4. Display memory performance and utilization statistics on the specified slot or all slots:

```
show khi performance memory [history | {slot[-slot] [,...]}]
```
5. Display process performance and utilization statistics on the specified slot or all slots:

```
show khi performance process [{slot[-slot] [,...]}]
```
6. Display thread performance and utilization statistics on the specified slot or all slots or sort the process output:

```
show khi performance pthread [{slot[-slot] [,...]}]
```
7. Display the queue performance and utilization statistics on the switch:

```
show khi performance rx-queue
```
8. Display internal memory management resource performance and utilization statistics on the specified slot or all slots:

```
show khi performance slabinfo [{slot[-slot] [,...]}]
```
9. Alternatively, you can sort the process and thread output using one of the following commands:
 - ```
show khi performance process sort <process-name | vmsize>
```
  - ```
show khi performance pthread sort <cpu | thread-name>
```

Example

```
Switch:1>show khi performance buffer-pool 1
Slot:1
  CPP:
    UsedFBuffs: 12
    FreeFBuffs: 3060
    RxQ0FBuffs: 0
    RxQ1FBuffs: 0
    RxQ2FBuffs: 0
    RxQ3FBuffs: 0
    RxQ4FBuffs: 0
    RxQ5FBuffs: 0
    RxQ6FBuffs: 0
    RxQ7FBuffs: 0
    TxQueueFBuffs: 0
    NoFbuff: 0
  Network stack system:
    UsedMbuf: 244
    FreeMbuf: 47606
    SocketMbuf: 19
  Network stack data:
```



```

        UsedMbuf: 4
        FreeMbuf: 10748
    Letter API message queue:
        QHigh: 0
        QNormal: 0
        FreeQEntries: 51200

Switch:1>show khi performance cpu 1
  Slot:1
    Current utilization: 6
    5-minute average utilization: 5
    5-minute high water mark: 15 (07/20/23 08:15:22)

Switch:1>show khi performance memory 1
  Slot:1
    Used: 6259174 (KB)
    Free: 1933970 (KB)
    Current utilization: 76 %
    5-minute average utilization: 76 %
    5-minute high water mark: 80 (%)
    10-minute average utilization: 76 %
    10-minute high water mark: 80 (%)
    1-Hour average utilization: 76 %
    1-Hour high water mark: 79 (%)
    1-Day average utilization: 76 %
    1-Month average utilization: 0 %
    1-Year average utilization: 0 %

SMLT socket buffer memory usage 0 KB on 0 blocks

```

The following example shows partial output for the **show khi performance memory history** command and depending on the hardware platform, you can also display virtual memory history.

```

Switch:1>show khi performance memory history
  Slot:1
Values indicate VMSize(VMRSS) in KB

-----
-----
Pid      Pname                5-Min      10-Min      1-Hour      1-Day      1-
Month      1-Year
-----
-----
2289  logger                2 (1)       2 (1)       2 (1)       2 (0)       --
(--      --(--))
2305  namServer             851 (7)     851 (7)     851 (7)     851 (0)     --
(--      --(--))
2306  sockserv              5 (2)       5 (2)       5 (2)       5 (0)       --
(--      --(--))
2307  oom95                 1178 (403)  1178 (403)  1178 (403)  1178 (0)    --
(--      --(--))
2308  oom90                 1178 (403)  1178 (403)  1178 (403)  1178 (0)    --
(--      --(--))
2309  imgsync.x             788 (7)     788 (7)     788 (7)     788 (0)     --
(--      --(--))
2380  logServer             794 (12)    794 (12)    794 (12)    794 (0)     --
(--      --(--))
2381  trcServer             787 (7)     787 (7)     787 (7)     787 (0)     --
(--      --(--))
2382  oobServer             787 (7)     787 (7)     787 (7)     787 (0)     --
(--      --(--))
2383  nickServer            724 (8)     724 (8)     724 (8)     724 (0)     --
(--      --(--))

```

2384	nickClient	724 (7)	724 (7)	724 (7)	724 (0)	--
(--)	--(--)					
3363	dhclient-ipv4-oob	81 (9)	81 (9)	81 (9)	80 (0)	--
(--)	--(--)					
2385	hwsServer	1688 (8)	1688 (8)	1688 (8)	1688 (0)	--
(--)	--(--)					
2386	redis-server	20 (5)	20 (5)	20 (5)	20 (0)	--
(--)	--(--)					
3405	hiveagent	13 (5)	13 (5)	13 (5)	12 (0)	--
(--)	--(--)					
2387	bundle-recovery.py	--(--)	--(--)	--(--)	--(--)	--
(--)	--(--)					
2388	cbcp-main.x	2570 (482)	2570 (482)	2570 (482)	2569 (0)	--
(--)	--(--)					
2392	rssServer	787 (6)	787 (6)	787 (6)	787 (0)	--
(--)	--(--)					
2393	dbgServer	787 (7)	787 (7)	787 (7)	787 (0)	--
(--)	--(--)					
2394	dbgShell	717 (6)	717 (6)	717 (6)	717 (0)	--
(--)	--(--)					
2396	khiCollection	1110 (7)	1110 (7)	1110 (7)	1110 (0)	--
(--)	--(--)					
2398	coreManager.x	788 (6)	788 (6)	788 (6)	788 (0)	--
(--)	--(--)					
2399	ovsdbServerCtl	--(--)	--(--)	--(--)	--(--)	--
(--)	--(--)					
2411	ssio	2916 (345)	2916 (345)	2916 (345)	2915 (0)	

The following example shows partial output for the **show khi performance process** command.

```
Switch:1>show khi performance process 1
Slot:1
```

```
-----
PID  PPID  PName          VmSize  VmLck  VmRss  VmData  VmStk  VmExe  VmLib
-----
1     0     init           2152    0      1488   168     132    32     1692
2     0     kthreadd       0        0        0       0       0       0       0
3     2     ksoftirqd/0    0        0        0       0       0       0       0
9819  2     kworker/u4:0    0        0        0       0       0       0       0
5     2     kworker/0:0H   0        0        0       0       0       0       0
7     2     rcu_sched      0        0        0       0       0       0       0
8     2     rcu_bh         0        0        0       0       0       0       0
9     2     migration/0    0        0        0       0       0       0       0
10    2     lru-add-drain  0        0        0       0       0       0       0
11    2     cpuhp/0        0        0        0       0       0       0       0
12    2     cpuhp/1        0        0        0       0       0       0       0
13    2     migration/1    0        0        0       0       0       0       0
14    2     ksoftirqd/1    0        0        0       0       0       0       0
15    2     kworker/1:0    0        0        0       0       0       0       0
16    2     kworker/1:0H   0        0        0       0       0       0       0
17    2     kdevtmpfs      0        0        0       0       0       0       0
```

```
--More-- (q = quit)
```

The following example shows partial output for the **show khi performance pthread** command.

```
Switch:1>show khi performance pthread 1
Slot:1
```

```

-----
---
TID   PID   PName          CPU(%) 5MinAvg CPU(%) 5MinHiWater CPU(%(time stamp))
-----
---
1     1     init           0.0    0.0
2     2     kthreadd       0.0    0.0
3     3     ksoftirqd/0    0.1    0.0
9967  9967  kworker/u4:0    0.0    0.0
5     5     kworker/0:0H    0.0    0.0
10005 2551  restweb_voss.py 0.1    0.2
7     7     rcu_sched       0.0    0.0
8     8     rcu_bh          0.0    0.0
9     9     migration/0     0.0    0.0
10    10    lru-add-drain   0.0    0.0
11    11    cpuhp/0         0.0    0.0
12    12    cpuhp/1         0.0    0.0
13    13    migration/1     0.0    0.0
14    14    ksoftirqd/1     0.0    0.0
15    15    kworker/1:0     0.0    0.0

--More-- (q = quit)

```

The following example shows output for the **show khi performance rx-queue** command.

```

Switch:1>show khi performance rx-queue
-----
CPP COUNTERS
-----

FBUF COUNTERS
InUseFBuffs:          0
FreeFBuffs:           3072
Stolen Rx Packets:    0
TxQueueFBuffs:        0
NoFBuffs:             0

PACKET COUNTERS:
totalCppEnetRxPkts:   0
totalRxPkts:          0
OamPktsRcvd:          0
IoPktsRcvd:           0
IoCopPktsRcvd:        0
PcapPktsRcvd:         0
Ipv6PktsRcvd:         2796
RxTestPkts:           0
RxCpHbPkts:           0
RxOopPkts:            0
RxMacMgmtPkts:        0
RxIpfixPkts:          0
RxOtherPkts:          0
RxCreditsAdded:       0
MyFramesReceived:     0
totalQueuedPkts:      920329
QueuedOamPkts:        0
QueuedIoPkts:         0
QueuedIoCopPkts:      0
QueuedTestPkts:       0
QueuedOtherPkts:      0
totalDequeuedPkts:    920329
DequeuedOamPkts:      0
DequeuedIoPkts:       920329

```

```

DequeuedIoCopPkts:      0
DequeuedTestPkts:      0
DequeuedOtherPkts:      0
totalPktsProcessed:     920329
OamPktsProcessed:       0
IoPktsProcessed:        920329
IoCopPktsProcessed:     0
OtherPktsProcessed:     0

minPktsRcvdAtOneTime:   0
maxPktsRcvdAtOneTime:   0
numCppInterrupts:       0
numTimesWeReceivedPkts: 0
numTimesTmainProcessedPkts: 816006

TxPktAttempts:          0
TxQueuedPkts:           0
TxDequeuedPkts:         0
TxPktsOk:               0
TxTestPkts:             0
TxIpv6Pkts:             0
TxIpfixPkts:            0
TxLsmPktsOk:           0
TxPktsRecovered:        0
totalLldpPktsRcvd       81683
totalDroppedUpnpFilterPKts 0

avgNumPktsRcvdAtOneTime: 0
avgNumPktsProcessedAtOneTime: 0

PACKET ERROR/DISCARD COUNTERS (non-zero counters only)
DroppedGlbSpbmDisPkts:   3713
DroppedAllRxPkts:        3713

NODE COUNTERS
pRxNodeList count:       0
pRxFreeList count:       0
pRxDirtyList count:      0
OutOfRxNodes count:      0
errorFindingRxBufCount:  0

pTxNodeList count:       0
pTxFreeList count:       0
pTxDirtyList count:      0
OutOfTxNodes count:      0
errorFindingTxBufCount:  0

CPP BUDGET COUNTERS
cppHardBudgetCount: 2
cppSoftBudgetCount: 2043
cppTicAbsenceTimeCount: 1
cppMimCount: 0
lastIntPid: 29
numSyncRxFrameEvents: 816006

CPP QUEUE STATS:
CPP Priority Queue Num 0  Total Rx Queue Count: 66848
CPP Priority Queue Num 0  Current Rx Queue Count: 0
CPP Priority Queue Num 0  Max Rx Queue Count: 12

CPP Priority Queue Num 1  Total Rx Queue Count: 2286
CPP Priority Queue Num 1  Current Rx Queue Count: 0
CPP Priority Queue Num 1  Max Rx Queue Count: 1

```

```

CPP Priority Queue Num 2  Total Rx Queue Count:    66
CPP Priority Queue Num 2  Current Rx Queue Count:   0
CPP Priority Queue Num 2  Max Rx Queue Count:       2

CPP Priority Queue Num 3  Total Rx Queue Count:     0
CPP Priority Queue Num 3  Current Rx Queue Count:   0
CPP Priority Queue Num 3  Max Rx Queue Count:       0

CPP Priority Queue Num 4  Total Rx Queue Count:    13
CPP Priority Queue Num 4  Current Rx Queue Count:   0
CPP Priority Queue Num 4  Max Rx Queue Count:       1

CPP Priority Queue Num 5  Total Rx Queue Count:   33957
CPP Priority Queue Num 5  Current Rx Queue Count:   0
CPP Priority Queue Num 5  Max Rx Queue Count:     277

CPP Priority Queue Num 6  Total Rx Queue Count:  608854
CPP Priority Queue Num 6  Current Rx Queue Count:   0
CPP Priority Queue Num 6  Max Rx Queue Count:     26

CPP Priority Queue Num 7  Total Rx Queue Count:  208305
CPP Priority Queue Num 7  Current Rx Queue Count:   0
CPP Priority Queue Num 7  Max Rx Queue Count:     11

CPP Tx Queue Count:      0
CPP Tx Max Queue Count:  0

Processed 10320 pkts over 2505497 milliseconds = 4 pkts/second

```

The following example shows partial output for the **show khi performance slabinfo** command.

```

Switch:1>show khi performance slabinfo
Slot:1

```

```

-----
---
Name                Active  Num    Objsize Objper  Pageper  Active  Num
                   Objs   Objs          slab    slab    slab    Slabs  Slabs
-----
---
nf_contrack         306    306     224    18      1       17     17
jffs2_refblock      32     32     248    16      1        2      2
jffs2_i             20     20    408    20      2        1      1
ip6-frags            0      0    136    30      1        0      0
UDIPv6              60     60    800    20      4        3      3
tw_sock_TCPv6        0      0    184    22      1        0      0
request_sock_TCPv6   0      0    232    17      1        0      0
TCPv6               40     40   1600    20      8        2      2
sgpool-128          12     12   2560    12      8        1      1
sgpool-64           12     12   1280    12      4        1      1
sgpool-16           12     12    320    12      1        1      1
cfq_queue            23     23    176    23      1        1      1
mqueue_inode_cache  30     30    544    15      2        2      2
nfs_direct_cache     0      0    184    22      1        0      0
nfs_commit_data      18     18    448    18      2        1      1

--More-- (q = quit)

```

Display KHI Control Processor Information

Display information about the type of packets and protocols received on a port. This command helps debug high CPU utilization issues.

About This Task

You can use the packets-per-second information in the output to identify where the bulk of packets destined for the CPU enter the switch.

For MLTs, statistics in the **show khi cpp port-statistics** command output increment equally on all MLT member ports.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display statistics for control packets that go to the control processor:

```
show khi cpp port-statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```
3. Display port statistics for SPBM internal ports:

```
show khi cpp port-statistics spbm-internal-ports [rx | tx]
```

Example

```
Switch:1>show khi cpp port-statistics
```

```
=====
```

```
      KHI CPP Details - Port Statistics
```

```
=====
```

Port	Packet Type	Rx Packets	Rx Diff	Rx Pps	Tx Packets	Tx Diff	Tx Pps
347 seconds since issuing the last KHI command.							
1/2	Ether2_LLDP(3)	0	0	0	522	12	0
1/2	Ether2_ARP_Request(10)	18	0	0	3706	48	0
1/2	Ether2_ARP_Other(11)	104	0	0	18	0	0
1/2	Ether2_Ipv4_Other(19)	20	0	0	0	0	0
1/2	LLC_BPDU(93)	0	0	0	7811	174	0
1/2	LLC_TDP(94)	0	0	0	2594	58	0
1/12	Ether2_LLDP(3)	521	12	0	521	12	0
1/12	LLC_TDP(94)	2598	58	0	2594	58	0
1/12	LLC_ISIS(95)	3105	69	0	3151	69	0
1/13	Ether2_LLDP(3)	521	12	0	521	12	0
1/13	Ether2_ARP_Request(10)	3121	69	0	3126	70	0
1/13	Ether2_ARP_Other(11)	3123	70	0	3121	69	0
1/13	Ether2_IPv4_UDP_Other(16)	3236	69	0	0	0	0
1/13	LLC_TDP(94)	2592	58	0	2594	58	0
1/13	LLC_ISIS(95)	0	0	0	3176	68	0
1/14	Ether2_LLDP(3)	521	12	0	521	12	0
1/14	Ether2_ARP_Request(10)	42	1	0	26	0	0
1/14	Ether2_ARP_Other(11)	24	0	0	0	0	0
1/14	LLC_TDP(94)	2594	60	0	2596	58	0
1/23	Ether2_LLDP(3)	520	12	0	521	12	0
1/23	LLC_TDP(94)	0	0	0	2596	58	0
1/24	Ether2_LLDP(3)	0	0	0	521	12	0
1/24	LLC_TDP(94)	0	0	0	2596	58	0
1/s1	Ether2_ARP_Request(10)	853	20	0	3140	72	0
1/s1	Ether2_ARP_Other(11)	3114	72	0	853	20	0
1/s1	Ether2_IPv4_UDP_Other(16)	3781	92	0	0	0	0
1/s1	LLC_BPDU(93)	0	0	0	7816	179	0
1/s1	LLC_ISIS(95)	0	0	0	3784	93	0

Display Fabric KHI Resource Scaling

About This Task

Perform the following procedure to display Key Health Indicator (KHI) information to determine the resources available for Fabric scaling.

After a resource in this command output reaches 80%, 90%, and 100% in the Usage column, the switch logs a WARNING message, sends an SNMP trap, and sets an alarm. If you use Extreme Cloud management applications, the switch also sends a message to that application. You can use this information to automate resources monitoring in

the Fabric. The switch clears the alarm after the resource usage drops below 80%, and also logs an INFO messages and send an SNMP trap. The following example displays both sample log and alarm text.

```
###Values reach the 80% watermark###
#1 2025-10-06T08:49:21.101Z 5420F-48T-4XE-FabricEngine CP1 - 0x001dc70a - 00770003.11 DYNAMIC SET GlobalRouter ISIS INFO Transparent UNIs
usage has reached 81% (over 80%)

###Values reach the 90% watermark###
#1 2025-10-06T08:49:21.101Z 5420F-48T-4XE-FabricEngine CP1 - 0x001dc70a - 00770003.11 DYNAMIC SET GlobalRouter ISIS INFO Transparent UNIs
usage has reached 90% (over 90%)

###Values reach the 100% watermark###
#1 2025-10-06T08:49:21.101Z 5420F-48T-4XE-FabricEngine CP1 - 0x001dc70a - 00770003.11 DYNAMIC SET GlobalRouter ISIS INFO Transparent UNIs
usage has reached 100% (over 100%)

Switch:1>show alarm database alarm-id 00770003

ALARM          EVENT          ALARM          ALARM          CREATION          UPDATED
CLEARED
SLOT    ID          CODE          TYPE          STATUS          SEVERITY  FREQ          TIME
TIME          TIME          REASON

-----
CP1  00770003.11    0x001dc70a    DYNAMIC      SET             WARNING    1    [10/03/25 11:05:07.833] [10/03/25 11:05:07.833]
[--/--/-- --:--:--:--] Transparent UNIs usage has reached 88% (over 80%)
```

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display Fabric KHI scaling resource information:
show khi resource-scaling

Example

```
Switch:1#show khi resource-scaling
=====
KHI resource-scaling
=====
Item          Maximum          Maximum          Currently          Available          Usage          Shared resource
              (theoretical)    (actual)          used               group
-----
SPBM nodes
-----
SPBM nodes home area          500          500          7          493          1 %    RES1,RES2
SPBM interior nodes home area          1
SPBM virtual nodes home area          2
SPBM boundary nodes home area          4
SPBM nodes remote area          500          500          6          494          1 %    RES1,RES2
SPBM interior nodes remote area          5
SPBM virtual nodes remote area          1
SPBM boundary nodes remote area          0
Total SPBM Nodes          800          695          13          682          2 %    RES1,RES2

BMACs
-----
Nodal SYS Id BMACs          13
VIST virtual BMACs          2
Multi-Area boundary node redirect BMACs          8
MCoSPB_source BMACs (*N1, *N2)          0
Total BMACs          800          800          23          777          3 %    RES1,RES2

ISIS
-----
Physical ISIS interfaces          7
L2FE ISIS interfaces          1
L3FE ISIS interfaces          128          120          0          120          0 %    RES1
Total ISIS interfaces          128          128          8          120          6 %    RES1

Services
-----
Multicast-fib entries          10000          10000          1844          8156          18 %
L2VSNs          3580          3580          465          3115          13 %    RES1
L2VSNs redistributed          2000          1998          463          1535          23 %
Transparent UNIs          53          53          0          53          0 %    RES1,RES3
Switched UNIs          4900          4900          1          4899          1 %    RES1,RES3,RES3
Private VLANs          200          200          1          199          1 %
```

```

L3VSNs                256          256          32          224          12 %    RES1,RES3
** Details:
*N1 - physical nodes and MA virtual nodes
*N2 - BEBs sending traffic to this node

```

View KHI Segmented Management Instance Information

Use the following commands to view Key Health Indicator (KHI) for segmented management instance interface information of the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View KHI management instance statistics:



Note

The type of management interface supported varies by platform. For information about supported Segmented Management Instance interface types, see *Fabric Engine and VOSS Feature Matrix*.

```
show khi mgmt statistics [clip | oob | vlan]
```

Example

The following example displays output from a switch with an OOB and Mgmt-sdwan interfaces configured.



Note

For information on the Mgmt-sdwan interface, see [Fabric Extend \(FE\) States](#) on page 32.

```
Switch:1>show khi mgmt statistics
```

Packet Counters - Interface Mgmt-oob			
Packet Type	Rx Packets	Tx Packets	Rx Packets Dropped
Ssh	1624	1552	0
DhcpBootpClient	0	0	13
DhcpBootpServer	0	0	17425
OtherUdpEphemeral	0	0	83
OtherUdpWellKnown	0	0	560
Icmp-echo-request	227	0	0
Icmp-echo-reply	0	4	0
Arp-reply	25	290	0
Arp-request	38165	25	0
V6-Http	0	0	111
IcmpV6-echo-request	336	0	0
IcmpV6-echo-reply	0	336	0
IcmpV6-dest_unreachable	0	111	0
IcmpV6-nd-nbr-solicit	133	139	0
Icmpv6-nd-nbr-advert	147	133	0
Packet Counters - Interface Mgmt-sdwan			
Packet Type	Rx Packets	Tx Packets	Rx Packets Dropped
Dropped			

Telnet	42	53	
0			
Ssh	31	37	
0			
DhcpBootpServer	0	0	
1071			
Icmp-echo-request	0	1	0
Icmp-echo-reply	1	0	0
Arp-reply	18295	0	0
Arp-request	0	1	
0			
IcmpV6-nd-nbr-solicit	0	1	0

Clear KHI Information

KHI information can be cleared globally across the whole device. Use the command to clear the CPP port statistics or Segmented Management Instance statistics.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Clear CPP statistics:
`clear khi cpp port-statistics`
3. Clear segmented management instance statistics:
`clear khi mgmt statistics`

Key Health Indicators Using EDM

Use the procedures in this section to display KHI information using EDM.

Clear KHI Statistics

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **CPP Stats Control** tab.
4. Select the statistics you want to clear.
5. Select **Apply**.

CPP Stats Control Field Descriptions

Use the data in the following table to use the **CPP Stats Control** tab.

Name	Description
PortStatsClear	Clears port statistics.

Display KHI Port Information

About This Task

Use the following commands to display key health information about the types of control packets and protocols received on a port and sent to the control processor.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **CPP Stats** tab.

CPP Stats Field Descriptions

Use the data in the following table to use the **CPP Stats** tab.

Name	Description
Port	Identifies the slot and port.
Packet	Shows the packet type.
PacketName	Shows the name of the packet.
RxPackets	Indicates the number of received packets on the port for the packet type.
TxPackets	Indicates the number of transmitted packets on the port for the packet type.

View KHI Segmented Management Instance Statistics Information

About This Task

Use the following procedure to view key health information about the types of control packets and protocols received on a port and sent to the Segmented Management Instance interface.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Mgmt Instance**.
2. Select **Stats**.
3. Select the **KHI** tab.
4. (Optional) To clear KHI statistics for a management interface, select an **InstanceId** row, and then select **Clear Stats**.

KHI Field Descriptions

Use the data in the following table to use the **KHI** tab.

Name	Description
InstanceId	Shows the management interface instance.
PacketType	Shows the packet type.
PacketName	Shows the name of the packet.
RxPackets	Shows the number of received packets on the port for the packet type.
TxPackets	Shows the number of transmitted packets on the port for the packet type.
RxDropped	Shows the number of received packets dropped on the port for the packet type.



Layer 2 Security

[Layer 2 Security for IPv4 and IPv6 Deployments](#) on page 2326

[Layer 2 Security Configuration Using the CLI](#) on page 2347

[Layer 2 Security Configuration using EDM](#) on page 2389

[Layer 2 security example scenarios](#) on page 2405

This section provides conceptual and procedural information for Layer 2 security for IPv4 and IPv6 deployments.

Layer 2 Security for IPv4 and IPv6 Deployments

This section describes Layer 2 security concerns and the security features you can use to mitigate them.

Security Features for IPv4 Deployments:

- DHCP Snooping
- Dynamic ARP Inspection
- IP Source Guard for IPv4 addresses

Security Features for IPv6 Deployments:

- First Hop Security (FHS)
- DHCP Snooping and IPv6 Neighbor Discovery Inspection
- IP Source Guard for IPv6 addresses

Dynamic ARP Inspection (DAI)

Table 172: Dynamic ARP Inspection (DAI) product support

Feature	Product	Release introduced
Dynamic ARP Inspection (DAI)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Dynamic ARP Inspection (DAI) is a security feature that validates ARP packets in the network.

Without DAI, a malicious user can attack hosts, switches, and routers connected to the Layer 2 network by poisoning the ARP caches of systems connected to the subnet, and intercepting traffic intended for other hosts on the subnet. DAI prevents these attacks by intercepting, logging, and discarding the ARP packets with invalid IP to MAC address bindings.

The switch dynamically builds the address binding table from the information gathered from the DHCP requests and replies when DHCP Snooping is enabled. The switch pairs the MAC address from the DHCP request with the IP address from the DHCP reply to create an entry in the DHCP binding table. For more information, see [Create DHCP Binding Table Entries](#) on page 2400.

When you enable DAI, the switch filters ARP packets on untrusted ports based on the source MAC and IP addresses seen on the switch port. The switch forwards an ARP packet when the source MAC and IP address matches an entry in the address binding table. Otherwise, the switch drops the ARP packet.



Note

- For DAI to function, you must enable DHCP Snooping globally.
- Configure DAI on a VLAN to VLAN basis.

DAI cannot be enabled on:

- Private VLANs (Etree)
- SPBM B-VLANs
- MLT port members

First Hop Security

First Hop Security (FHS) improves local network security by employing a number of mitigation techniques. This section describes the base set functionality which provides protection from a wide host of rogue or mis-configured users, and this can be extended with additional features for different deployment scenarios. For example, see the following topology.

Sample Topology

In the following topology, Layer 2 switch SW-1 is connected to another Layer 2 switch SW-2. SW-2 is connected to three hosts and SW-1 is connected to two hosts.

In this network, if FHS is enabled only on SW-1, then it can only save the nodes which are directly connected to it. To protect the good node connected to SW-2, the FHS must be enabled on SW-2.




Figure 165: First Hop Security topology

First Hop Security contains the majority of the RIPE 554 mandatory requirements for Layer 2 switches. This includes the following:

- DHCPv6 Guard or DHCPv6 filtering
- RA Guard or Router Advertisement filtering

DHCPv6 Security Concerns

The enhancements in IPv6 provide better security in certain areas, but some of these areas are still open to exploitation by attackers. This section identifies the IPv6 FHS concerns associated with Dynamic Host Configuration Protocol version 6 (DHCPv6).

DHCPv6 (RFC 3315) describes how a host can acquire an IPv6 address and other configuration options from a server that is available on its local link. DHCPv6 is

described as a stateful protocol. In other words, DHCPv6 can operate in a stateless fashion where it provides configuration information to nodes and does not perform address assignments (RFC 3736). In addition, it can operate in a stateful manner, where it assigns IPv6 addresses and configuration information to hosts that request it.

As in IPv4 DHCP, DHCPv6 is susceptible to rogue server attacks. In other words, if DHCPv6 is used to provide IPv6 addresses to the hosts, an attacker that managed to insert a rogue DHCPv6 server in the link can potentially assign addresses and configuration options to the link hosts. In turn, the attacker can deploy adversary-in-the-middle, traffic interception, or discard traffic, similar to those in the stateless address autoconfiguration scenario. Therefore, it is important to use DHCP protections for both IPv4 and IPv6.

DHCPv6 Guard

Table 173: DHCPv6 Guard product support

Feature	Product	Release introduced
DHCPv6 Guard	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

DHCPv6 Guard is a type of security for IPv6 deployments in an enterprise environment, it provides Layer 2 security to DHCPv6 clients by protecting them against rogue DHCPv6 servers. The basic concept of DHCPv6 Guard is that a Layer 2 device filters DHCPv6 messages meant to DHCPv6 clients, based on a number of different criteria. The basic filtering criterion is, the DHCPv6 server generated packets which are received on non-server ports or from an untrusted server will be dropped by the Layer 2 device.

Various levels of granularity are provided. Following are the policies that are supported:

- Port based filtering using device role (server or client)
- Server or relay agent IPv6 address based filtering
- Advertising IPv6 prefix based filtering
- DHCPv6 packet filtering based on Server Preference checks

The following figures are DHCPv6 topology samples:

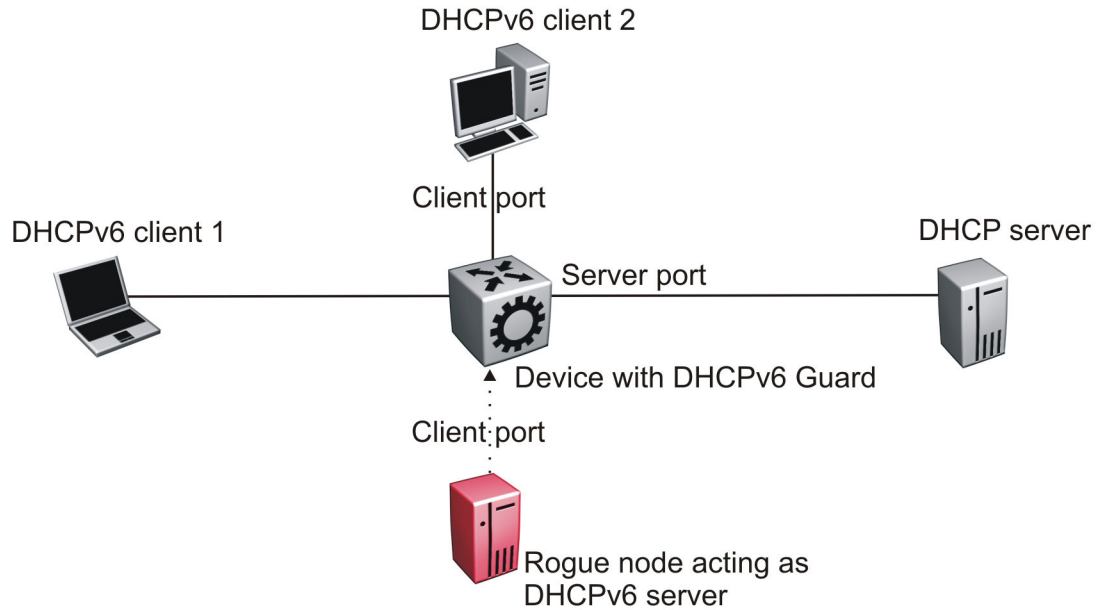


Figure 166: DHCPv6 Topology 1




Figure 167: DHCPv6 Topology 2

DHCPv6 Guard Policies Configuration

You can configure DHCPv6 Guard policies using CLI, SNMP and EDM. The following policies are supported for DHCPv6 Guard.

Port-based Filtering Using Device-role

Port-based filtering using device-role is an interface-level configuration. Only a DHCPv6 server or relay agent can send a DHCPv6 advertisement or reply. By configuring the device-role attached to the port (whether it is a client or server), the rogue server generating DHCPv6 advertisement or reply packets can be blocked if these packets are received on a port configured as a client. Device-role can be applied only on port, and not on MLT, SMLT, or VLAN. If you configure device-role on an MLT, SMLT, or VLAN, you must configure same device-role on all the MLT, SMLT, or VLAN member ports.

In DHCPv6 Guard Topology 1, only DHCPv6 server packets (that is, advertisement, reply) received on a port configured as a Server port accept the packets and process them for security validation and forwarding. The Client port drops the packets if it receives packets generated from a DHCPv6 rogue server.

Server or Relay Agent IPv6 Address Based Filtering

Server or relay agent IPv6 address-based filtering enables the verification of the advertised DHCPv6 server and relay address in messages with the configured authorized server access list. In DHCPv6 Guard Topology 1 and Topology 2, you can configure the access list to accept DHCPv6 server packets from a specific Source IPv6 address such as a DHCPv6 server or DHCPv6 relay IPv6 address.

Advertising IPv6 prefix-based filtering

Advertising IPv6 prefix-based filtering enables verification of the advertised prefixes in DHCPv6 reply messages with the configured authorized prefix list.

Server preference-based filtering

Server preference-based filtering enables verification by checking if the advertised preference (in preference option) is greater than or less than the specified limit.

RA Guard

Table 174: IPv6 Router Advertisement (RA) Guard product support

Feature	Product	Release introduced
IPv6 Router Advertisement (RA) Guard	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

IPv6 hosts can configure themselves automatically when connected to a routed IPv6 network through ICMPv6 router discovery messages. When the host is connected to the network for the first time, it sends a link-local router solicitation multicast request for its configuration parameters. If the host is configured correctly, routers respond to the request with a Router Advertisement (RA) packet. The RA packet contains network-layer configuration parameters.

In addition to filtering RAs, RA Guard introduces the concept of router authorization proxy. Instead of each node on the link analyzing RAs and making an individual

decision, a legitimate adversary-in-the-middle performs the analysis on behalf of all other nodes on the link.

Stateless and statefull RA Guard functions are available. The switch supports only the stateless RA Guard function.

Stateless RA Guard examines incoming RAs and decides whether to forward or block them based on the information found in the message or in the Layer 2 device configuration. The following list identifies the typical information available in the received frames that are used for RA validation:

- Port on which the frame is received
- Source IPv6 address
- Prefix list which RA carries
- Link-Layer address of the sender

After the Layer 2 device successfully validates the RA packet content against the configuration, the RA is forwarded to its destination, whether unicast or multicast. If the validation fails, the RA is dropped at the Layer 2 device.

RA Guard policies description

This section describes the RA Guard policies. The following policies are supported for RA Guard:

- Port-based filtering using device role (host or router)
- Source IPv6 based filtering
- Advertised IPv6 prefix-based filtering
- Source MAC address-based filtering
- RA packet for managed address configuration flag validation
- RA packet for hop count limit validation
- RA packet for Router Preference validation

Port-based Filtering Using Device-role

This configuration is an interface-level configuration. According to Neighbor Discovery (ND) RFC 4861, only the IPv6 router can generate the RA packets. By configuring the device-role attached to the port whether it is a host or router, the rogue host which is generating RA packets can be blocked. Device-role can be applied only on a port, and not on an MLT, SMLT, or VLAN. If you configure device-role on an MLT, SMLT, or VLAN, you must configure the same device-role on all the MLT, SMLT, or VLAN member ports.

In the following topology, the switch is connected to a Layer 3 router and three hosts. Because the router is directly connected to port 1/2, the device-role of the port 1/2 is configured in Router mode. The other hosts are connected to ports 1/3, 1/4, and 1/5, and the device-role of ports 1/3, 1/4, and 1/5 are configured in Host Mode.

The host connected to the port 1/4 is a rogue host and if it is trying to send RA packets, then the switch drops those RA packets received on the interface 1/4 as the device-role of this port is Host Mode.




Figure 168: RA Guard Topology1

IPv6 Source Address Based Filtering

An IPv6 source address based filtering policy enables the source IPv6 address verification of the RA packets against the configured RA source IPv6 list.

The following figure shows the RA packet format. RA Guard policy verifies the IPv6 source address (SrcIP) in the IPv6 Header against the configured RA source IPv6 list.



Figure 169: IPv6 ICMP RA Data Packet Outline

Advertised IPv6 Prefix-based Filtering

Advertised IPv6 prefix-based filtering enables verification of the advertised prefixes in inspected messages against the configured RA prefix list.

The following figure illustrates the IPv6 ICMP RA data packet outline. This RA Guard policy verifies the RA (Prefix Information) in ICMPv6 data against the configured RA prefix list.



Figure 170: IPv6 ICMP RA Data Packet Outline

Source MAC Address-based Filtering

Source MAC address-based filtering enables the source MAC address of the RA packets verification against the configured authorized MAC list.

The following figure illustrates the IPv6 Ethernet packet. This RA Guard policy verifies the received RA packets source MAC address against the configured authorized MAC access list.




Figure 171: IPv6 Ethernet Packet

RA packet for managed address configuration flag validation

In the RA packets, there is an “M” flag (managed address configuration flag) that can be configured to indicate that the address assignments are available through DHCPv6. This means that DHCPv6 takes care of the interface address assignment in that LAN segment. If a filtering policy is enabled, then all the RA packets without an “M” flag are dropped. By default, this validation is not performed.

The following figure illustrates IPv6 ICMP RA data packet outline for managed address configuration.




Figure 172: IPv6 ICMP RA data packet outline

RA packet for hop count limit validation

RA packet for hop count limit validation policy verifies the advertised RA message if the hop count limit is within the configured hop count limit. If the received hop count limit is not within the configured limit, then those RA packets are dropped.

The following figure illustrates IPv6 ICMP RA data packet outline for hop count limit validation.




Figure 173: IPv6 ICMP RA data packet outline

RA packet for router preference validation

The RA packet contains the Router Preference as part of the flags field. This can be high, medium, or low. This filtering policy option verifies if the advertised default router preference parameter value is lower than or equal to a specified limit.

The following figure illustrates IPv6 ICMP RA data packet outline for router preference validation.




Figure 174: IPv6 ICMP RA data packet outline for router preference validation

Capturing and verifying FHS specific packets against the configured policies

First Hop Security filters can be installed only if FHS is enabled globally. The DHCPv6 Guard or RA Guard filters are created as a part of First Hop Security filter with port bit mask "0".

The following list identifies the high-level tasks to capture DHCPv6 packets received on a physical port:

1. Enable FHS globally.
2. Enable DHCPv6 Guard or RA Guard globally.
3. Create a DHCPv6 Guard or an RA Guard policy.
4. Configure RA Guard or DHCPv6 Guard device role on the port.
5. Attach DHCPv6 Guard and/or RA Guard policy to a physical port if needed.

On configuring RA Guard or DHCPv6 Guard device role on the port, the appropriate port bitmask for that port will be updated in the data path filter.

The RA or DHCPv6 sever initiated packets received on trusted ports (router or server ports) will be sent to the local CPU for further validations. If these packets pass the RA Guard and DHCPv6 Guard validation, they will be forwarded towards the intended host or DHCPv6 client; if not, they will be dropped by the switch.

FHS Restrictions

The following list identifies First Hop Security (FHS) restrictions:

- Fragmented RA and DHCPv6 server initiated packets are dropped on the FHS enabled switch.
- DHCPv6 Guard and RA Guard do not work on devices connected on shared media or on tunneled interfaces.
- DHCPv6 Guard or RA Guard policies are not VLAN or MLT based.
- FHS is not supported on the Out Of Band (OOB) port on the switch.
- Packets received on FHS ports with more than one extension header, and if they are destined to link-local unicast or link-scope multicast address, are dropped as they cannot be classified as RA or DHCPv6 reply packets.
- The FHS functionality can be bypassed at the first hop switch, if the malicious packets are destined to global address, and have more than one extension header.
- If the FHS rules and IPv6 filters match for a packet, the IPv6 filter has precedence.
- In a Layer 2 VSN, packets are not filtered based on FHS rules. Enable FHS on the required UNI ports to protect the connected devices from FHS attacks.

FHS Configuration Guidelines

The following list provides details on how certain FHS configurations work and how you should use them:

1. FHS IPv6 Access lists are generic access/prefix lists which can be applied on IPv6 source address or the prefixes advertised in RA or DHCPv6 messages. If you filter on the basis of a particular IPv6 source address, you must configure the access list entry with complete source address with prefix-length value of 128. If you allow a group of source addresses within a prefix range, you must configure the IPv6 ACL entry with an appropriate prefix length and attach this IPv6 ACL to the appropriate match parameters in RA Guard or DHCPv6 Guard policies.

If you filter a particular prefix, you must configure an IPv6 access list entry with appropriate prefix and prefix-lengths. To filter based on prefix, prefix-lengths should be less than 128. Following is an example of IPv6 access list entry:

```
ipv6 fhs ipv6-access-list match_src_allow fe80:0:0:0:0:ff:fe00:113/128
mode allow
```



Note

- a. If no IPv6 ACL is attached to an RA Guard or DHCPv6 Guard policy as a source ACL, then IPv6 source address in the incoming RA packets or packets from DHCP server will not be validated, and such packets will not be dropped due to source address validations.
 - b. If no IPv6 ACL is attached to an RA Guard or DHCPv6 Guard policy as a prefix ACL, then prefix information in incoming RA packets or packets from DHCP server will not be validated and these packets will not be dropped due to prefix validations.
 - c. The FHS access or prefix lists are different from IPv6 prefix lists. For FHS, the switch maintains a separate list (cannot reuse IPv6 prefix lists) as IPv6 prefix lists do not have any action associated with them, whereas FHS has an action associated with each ACL entry.
2. When an IPv6 ACL is attached to an RA Guard or DHCPv6 Guard policy and the address or prefix in the incoming RA Guard or DHCPv6 server packets received on port to which this RA Guard or DHCPv6 Guard policy is attached does not match any of the entries in that IPv6 ACL, the packet will be dropped by default. If you want to change this behavior to default (allow, for IPv6 ACLs), you can add an entry that matches all the packets and set the action as allow. To do this, use the following command:


```
ipv6 fhs ipv6-access-list no_match_src_def_allow 0:0:0:0:0:0:0:0/0
mode allow
```
 3. IPv6 ACL entries with conflicting prefixes within an IPv6 ACLs are not allowed, and such configuration will fail with appropriate error message. Conflicting entries can be present in two or more different IPv6 ACLs.
 4. The entries within an IPv6 ACL will be sorted in increasing order of IPv6 prefixes. If there are two entries with same prefix address within an ACL, then such entries will be ordered with increasing value of their prefix-lengths.
 5. MAC ACL entries are ordered in the increasing order of MAC addresses within a MAC ACL. If none of the entries in the MAC ACL match the source MAC address of RA packet, then the packet will be dropped by default. If no MAC ACL is attached to an RA Guard policy, then the source MAC address of RA packets is not validated.
 6. When matching for a prefix using IPv6 ACL entry, if you advertise a prefix with matching prefix but prefix-length lesser than configured prefix-length, then the packet has to be considered as no match and prefix matching process has to continue with remaining IPv6 ACL entries in that ACL.

The rationale behind this functionality is to avoid wrong configuration of access side devices. This functionality safeguards the devices in an access network if a wrongly configured IPv6 prefix is advertised or a malicious user is sending invalid (wrong) prefixes. For example, consider the following scenario:

Configure the prefix in ACL entry (without ge and le values): **ipv6 fhs ipv6-access-list ipv6_acl_entry_1 2000:0123:4567:89ab::/64 mode allow**

Advertise the prefix in RA packet: **2001:0123:4567:89ab::/48**

This advertised prefix matches the configured IPv6 ACL entry and without this prefix-length check functionality, the packet is allowed to pass through. But, actually it is configuring all access devices in that network with wrong IPv6 configurations in different IPv6 network (**2001:0123:4567::/48**)

With prefix-length check functionality, this configuration is not allowed as advertised prefix length is not equal to configured prefix length. So, the wrong configuration of access devices is avoided.

7. Importance of “ge” and “le” parameters in an IPv6 ACL entry:

A user can optionally configure “ge” (greater than or equal to) and “le” (lesser than or equal to) parameters while configuring an IPv6 ACL entry. If the prefix advertised in a packet matches the configured prefix in an IPv6 ACL entry, and “ge” and “le” values are configured (not default) for that IPv6 ACL entry:

- The packet will be allowed to go through only if the prefix-length in the packet is within the range of configured “ge” and “le” values.
- If prefix lengths in the packet are not within the configured range of “ge” and “le” values (non-default values), then the packets would be considered as no match for that IPv6 ACL entry and search for matching IPv6 ACL entry continues within that IPv6 ACL.
- If no ge and le values are configured, those values by default are set to configured prefix length in that IPv6 ACL entry.
- ge and le values are allowed only if they are greater than configured prefix.
- When both are configured (not default values), ge value should always be smaller than le value.

These configurations provide more control over the advertised prefixes in RA or DHCPv6 packets.

8. As “ge” and “le” values are valid only for advertised prefixes, they will not be applied to IPv6 addresses, which are not prefixes. For such addresses, prefix match is considered as match for that IPv6 ACL entry and the corresponding action of that ACL entry is applied on that packet. “ge” and “le” configurations are irrelevant for the following:

- IPv6 source address in RA packet
- IPv6 source address in packets from DHCPv6 server (like DHCPv6 advertise, DHCPv6 reply)
- IPv6 address (temporary or non-temporary) advertised in packets from DHCPv6 server. For example, IPv6 addresses advertised in IANA option of DHCPv6 reply packets

9. Order of packet validations:

In RA or DHCPv6 packets received at the CP for FHS processing, the following order of processing is carried out:

- a. Packet parsing
- b. Checking for presence of IPv6 fragment header

- c. Checking if packets are RA packets or DHCPv6 packets from server (Advertise, Reply, Reconfigure, Relay-Reply)
- d. Basic validations:
 - Non-Link-Local source IPv6 address (only for RA packets)
 - L4 length validations
 - Checksum validations
- e. If an RA Guard or DHCPv6 Guard policy is attached to a port:
 - MAC ACL validations (if configured) (Only for RA packets)
 - IPv6 source address ACL validation (if configured)
 - IPv6 prefix ACL validations (if configured)
 - Other packet parameter validations like:
 - Managed config flag (RA)
 - ICMP hop limit (RA)
 - Router preference (RA)
 - Server preference (DHCPv6)

If any of these validations fail or if action associated with a match ACL entry indicates to DROP (or default drop if ACL is attached to corresponding policy but packet does not match any ACL entry in that ACL), then the packets are dropped and corresponding statistics are updated. If all these pass or actions related to all matched ACL entries are PERMIT, then the packet is allowed to go through.

- 10. Longest prefix match: If a packet matches multiple entries in an ACL, then the action associated with an entry with longest prefix match would be applied on the packet.
- 11. If a port is configured as untrusted (“host” as device role for RA Guard or “client” as device role for DHCPv6 Guard), all the FHS trusted traffic (RA packets for RA Guard or packets from DHCPv6 server for DHCPv6 Guard) are dropped in data path itself. Also for such drops, statistics are not incremented.

If a port is neither configured as trusted nor untrusted, then the FHS traffic (RA packets or DHCPv6 packets from DHCPv6 server) is switched as if FHS is not present.

- 12. Creation of FHS port policy mappings:

Until, and unless, any of the FHS parameters are configured on a port, port policy mappings are not created and thus with no port to policy mapping configured, the system does not display any entries while listing port policy mappings using the command **show ipv6 fhs port-policy**.

- 13. If a RA Guard or DHCPv6 Guard policy is attached to any of the ports, deletion of such policy is not allowed. In the contrary, to delete an RA Guard or DHCPv6 Guard policy, those policies need to be detached from all the ports in the switch. However, modification of an RA Guard or DHCPv6 Guard policy is allowed even if it is attached to ports.
- 14. If a MAC or IPv6 ACL is attached to an RA Guard or DHCPv6 Guard policy, you cannot delete the ACL itself. You can delete the entries from this policy even if it is attached to any policy. At least one entry needs to exist in a MAC or IPv6 ACL; you cannot delete the last entry in that ACL if that ACL is attached to any RA Guard or DHCPv6

Guard policy. You must detach that ACL from all the policies to delete that ACL. However, you can update the entries in that ACL even if it is attached to a policy.

If a port is configured as trusted ("Server" port for DHCPv6 Guard and "Router" port for RA Guard), then only one can attach a DHCPv6 Guard or RA Guard policy to that port. In the contrary, if any policy is attached to a port, the port role cannot be changed from trusted ("Server" port for DHCPv6 Guard and "Router" port for RA Guard) to other role ("Client" port for DHCPv6 Guard, "Host" port for RA Guard or "None" for both) until that policy is not detached from port.

DHCP Snooping and Neighbor Discovery inspection

Table 175: Dynamic Host Configuration Protocol Snooping and Neighbor Discovery Inspection product support

Feature	Product	Release introduced
DHCP Snooping(IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
DHCP Snooping (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 175: Dynamic Host Configuration Protocol Snooping and Neighbor Discovery Inspection product support (continued)

Feature	Product	Release introduced
Neighbor Discovery Inspection (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

DHCP Snooping

DHCP Snooping is a Layer 2 security feature, that provides network security by filtering untrusted DHCP messages received from the external network causing traffic attacks within the network. DHCP Snooping is based on the concept of trusted versus untrusted switch ports. Switch ports configured as trusted can forward DHCP Replies, and the untrusted switch ports cannot. DHCP Snooping acts like a firewall between untrusted hosts and DHCP servers.

The switch supports:

- DHCP Snooping for both IPv4 and IPv6.
- Neighbor Discovery (ND) inspection for IPv6.

Security is critically important in an access network because various devices can connect to an access network that may not be administratively controlled by a single administrator. Stateless Address Autoconfiguration (SLAAC) and Duplicate Address Detection (DAD) mechanisms used by IPv6 are more vulnerable to attacks from a malicious user. If any person, intentionally or unintentionally, configures an IP address on the device interface wrongly and advertises that IP address as one's own address during DAD mechanism initiated by other device, DAD initiated devices cannot assign this address. If a malicious user replies to all the DAD IP addresses as own address, none of the devices in the access network can assign any IP addresses to their interfaces. Thus, DoS attacks can be easily carried out by the malicious user making the entire network unfunctional. In another kind of attack, a malicious user can try to poison the neighbor cache of a host by sending ND packets with bogus MAC address which is learnt by other hosts into their neighbor table. Due to the infiltration of the bogus MAC address in the host's neighbor table, the packets destined to its neighbor is sent to the bogus MAC address and is eventually dropped or received by an unintended host.

In general, these kinds of attacks are carried out by sending different Neighbor Discovery (ND) packets – either through solicited ND packet exchanges or as a result of unsolicited ND packet exchanges triggered due to an event like the expiry of ND timers. These packets carry interface IP address information and link-layer address information. Other devices use this information to build their neighbor table for forwarding traffic to or through the malicious device. As part of ND inspection mechanism, ND (specifically, NS, NA, and redirect) packets from only trusted hosts are allowed to pass through and the packets from untrusted hosts are dropped in the switch itself. Other network devices can safely use ND mechanisms for correctly assigning IP address to their interfaces resulting in a smooth traffic flow.

For validating the ND packets, the switch must first learn the trusted information by various mechanisms and store the information in a DHCP binding table. If the switch receives ND packets on an untrusted port, the packets are validated against entries in the DHCP binding table. If the ND packets pass the validation, the packets are forwarded. If the packets fail the validation, they are dropped in the switch itself. This process avoids invalid NA packets from propagating beyond the access switch.

DHCP Snooping and ND inspection feature protects the network from the following types of attacks:

- **User misconfigurations:** Host assigns an address which should not be used by the recipient device. ND inspection blocks this address in the access switch because binding entry does not exist for that address for that host.
- **DAD spoofing:** Malicious user claims that the address is taken even if it is not.
- **NUD spoofing:** Malicious host responds to NUD NS packets indicating that the address is still reachable via that host even if that neighbor is actually not reachable.
- **ND cache poisoning:** Malicious user sends different (invalid) link-layer addresses for a target IP address causing other hosts in the network to program bogus MAC for a given IP neighbor, as a result of which, the traffic gets discarded or misused by malicious host.

DHCP Binding Table

DHCP Snooping builds and maintains a binding table, this binding table contains the MAC address, IP address, lease time, binding type, VLAN number, and port information that correspond to the local untrusted ports of the switch. When the switch receives a DHCPRELEASE or DHCPDECLINE broadcast message, DHCP Snooping performs a lookup of the MAC address in the binding table to determine if the port information in the binding table matches the port on which the message was received. If the port information matches, the DHCP packet is forwarded, otherwise it is dropped.

Trust Bindings

A switch enabled with the Neighbor Discovery inspection feature allows NA packets through, if the packets are from a trusted host. To allow or deny Neighbor Advertisement (NA) packets, trust bindings must be established using following methods:

- Configuring the port connected to a device (or host) as trusted.
- Building a DHCP binding table which contains entries from trusted devices (or hosts) only. This DHCP binding table is used for validating NA packets.

This method of trust binding involves 2 processes:

- IP address learning (snooping) process

In this process an IP address is learnt through a trusted means and a DHCP binding table is built. The switch supports the DHCP binding table entry learning by:

- Statically configuring the entries
 - Dynamically learning by DHCP Snooping packets
- NA packet validation (inspection) process

This process uses the DHCP binding table entries which are populated as part of IP address learning process to validate the incoming NA packets.

After the DHCP binding tables are built, the information gathered using trust binding is used to validate the ND packets. If the ND packets cannot be validated using this information, they are considered as packets received from an un-trusted host and are dropped by the switch.

Restrictions

In addition to the FHS restrictions, DHCP Snooping and ND inspection have the following restrictions:

- DHCPv6 Snooping entries are not displayed in an SPB network. The Administrator must add manual entries.
- Link-local address validation is not supported under ND inspection. Thus, an First Hop Security (FHS) enabled switch is vulnerable to attackers who try to attack with link-local addresses.
- As a 5-second timer is used to cleanup expired DHCP binding table entries, the expired DHCP binding table entries can remain in the DHCP binding table for up to 5 seconds after they expire.
- If an FHS-enabled switch is rebooted, all the dynamically-learned binding entries get flushed and those entries need to be re-learned for ND inspection to pass. However, when the switch is rebooted, DHCP clients connected to it do not re-initiate DHCP learning, due to which, the switch cannot learn these assigned IP addresses. As a result, ND inspection fails for these addresses. To overcome this problem either DHCP client must learn the IP address again through DHCP mechanisms or the administrator must add static entries for these addresses.
- For IPv6, DHCP binding table entries learned through DHCP are not removed from the DHCP table on DHCP clients that release these addresses. The administrator must manually remove these entries after the addresses are released.
- A dynamic DHCP binding table entry is learned only using the DHCP mechanism. For other modes of address configuration on the host, a relevant DHCP binding table entry must be configured on the FHS switch so that ND packets from such host are not blocked due to ND inspection processing.
- The switch deletes DHCP-Snooping static bindings if you use CLI to remove the port from the associated default VLAN ID. This deletion also occurs on an EAP port when the port dynamically moves from the default VLAN ID, Guest VLAN, or Fail Open VLAN.

- DHCP Snooping is not supported on:
 - DHCP Relay switch
 - Etree
 - Private VLANs
- DHCP Snooping is supported on Split Multi-Link Trunking (SMLT) only if the SMLT ports are DHCP Snooping trusted.
- As a best practice, do not configure DHCP Snooping on core switches.
- Configure **ip dhcp-snooping trusted** only on the port that connects to the DHCP Server.
- You can configure DHCP Snooping on separate switches in the following scenarios:
 - on the switch where the DHCP server connects on the DHCP server VLAN.
 - on the switch where the DHCP client connects on the DHCP client VLAN(s).
- You can configure DHCP Snooping only on the switch where the DHCP server connects and only on a DHCP server VLAN.
- If DHCP clients connect to the same switch as the DHCP server, you can configure DHCP Snooping on the client VLANs.
- A best practice is to connect the DHCP server and DHCP clients to different switches.

IP Source Guard

Table 176: IP Source Guard product support

Feature	Product	Release introduced
IP Source Guard (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 176: IP Source Guard product support (continued)

Feature	Product	Release introduced
IP Source Guard (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

IP Source Guard (IPSG) is a Layer 2 port-to-port feature that works closely with DHCP Snooping. It prevents IP spoofing by allowing only IP addresses obtained using DHCP Snooping. When you enable IPSG on an untrusted port with DHCP Snooping enabled, an IP filter is automatically created or deleted for that port based on the information stored in the corresponding DHCP Snooping binding table entry. When a connecting client receives a valid IP address from the DHCP server, the filter installed on the port allows traffic only from that assigned IP address.

You can configure IPSG on a port using the command line interface (CLI), the Enterprise Device Manager (EDM), or SNMP.

**Note**

The switch supports configuration of IP Source Guard for both IPv4 and IPv6 addresses.

The following table shows how IPSG works with DHCP Snooping.

Table 177: IP Source Guard and DHCP snooping

IP Source Guard configuration state	DHCP snooping configuration state	DHCP snooping Binding Entry action (untrusted ports)	IP Source Guard action
change from disabled to enabled	enabled	creates a binding entry	creates a filter for the IP address using the IP address from the binding table entry
enabled	enabled	creates a binding entry	creates a filter for the IP address using the IP address from the binding table entry

Table 177: IP Source Guard and DHCP snooping (continued)

IP Source Guard configuration state	DHCP snooping configuration state	DHCP snooping Binding Entry action (untrusted ports)	IP Source Guard action
enabled	enabled	deletes a binding entry	deletes the IP filter and installs a default filter to block all IP traffic on the port
enabled	enabled	deletes binding entries when one of the following conditions occur: <ul style="list-style-type: none"> • a DHCP release packet is received • the port link is down • the lease time has expired • the port is removed from the VLAN • the VLAN is deleted • the port is set as trusted • the binding entries are manually deleted 	deletes the corresponding IP filter and installs a default filter to block all IP traffic
change from enabled to disabled	enabled	not applicable	deletes the installed IP filter for the port
disabled	enabled	creates a binding entry	
disabled	enabled	deletes a binding entry	

IP Source Guard General Guidelines

The following provides a list of guidelines to follow if you use IP Source Guard (IPSG):

- You can enable IP Source Guard (IPSG) only on a port that is DHCP Snooping and Dynamic ARP Inspection untrusted.
- The port must be a member of a VLAN. DHCP Snooping must be enabled globally and on the VLAN. You must also enable Dynamic ARP Inspection on the same VLAN.
- You cannot enable IPSG on MLT, SMLT, DMLT or LAG ports.
- You cannot enable IPSG on a router port.
- You cannot enable IPSG on ports that are members of a private VLAN.

- You cannot remove a port that is IPSG-enabled from a VLAN. Similarly, you cannot delete a VLAN that has at least one port that is IPSG-enabled.
- A maximum of 10 IP addresses are allowed on each IPSG-enabled port. Correspondingly, a maximum of 10 IP filters are automatically created for each of those ports. When this number is reached, no more filters are set up and all traffic is dropped.

For more information about the supported number of IP filters on each platform, see *Fabric Engine Release Notes*.

Layer 2 Security Configuration Using the CLI

Use the following sections to help you configure Layer 2 security features and protect the network by mitigating various types of attacks, using the Command Line Interface (CLI).

For IPv4 deployments, configure:

- DHCP Snooping
- Dynamic ARP Inspection
- IP Source Guard for IPv4 addresses

For IPv6 deployments, configure:

- First Hop Security (FHS)

**Note**

FHS does not solve all cases of denial of services like blocking flooding of the IPv6 messages.

- DHCP Snooping and IPv6 Neighbor Discovery Inspection
- IP Source Guard for IPv6 addresses

DHCP Snooping Configuration Using CLI

The following section provides procedures to configure DHCP Snooping using the CLI.

Enable or Disable DHCP Snooping globally

Use the following procedure to enable DHCP Snooping globally. If DHCP Snooping is globally disabled, the switch forwards DHCP reply packets to all required ports, both trusted or untrusted.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable DHCP Snooping globally:
`ip dhcp-snooping enable`
3. Disable DHCP Snooping globally:
`no ip dhcp-snooping enable`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip dhcp-snooping enable
```

Variable Definitions

The following table defines parameters for the **ip dhcp-snooping** command.

Variable	Value
<i>enable</i>	Enables or disables DHCP Snooping globally. By default, DHCP Snooping is disabled.

Enable or Disable DHCP Snooping on a VLAN

Use the following procedure to configure DHCP Snooping on a specific VLAN. If DHCP Snooping is globally disabled, the switch forwards DHCP reply packets (received on trusted or untrusted ports) to all ports.

If you enable DHCP Snooping globally, the agent determines whether to forward DHCP reply packets based on the DHCP Snooping mode of the VLAN and trusted state of the port.



Note

You cannot enable DHCP Snooping on Private VLANs (E-Tree) and SPBM B-VLANs.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. Enter VLAN Interface Configuration mode:
`enable`
`configure terminal`
`interface vlan <1-4059>`
2. Enable DHCP Snooping on the VLAN:
`ip dhcp-snooping enable`
3. Disable DHCP Snooping on the VLAN:
`no ip dhcp-snooping enable`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 5
Switch:1(config-if)#ip dhcp-snooping enable
```

Variable Definitions

The following table defines parameters for the **ip dhcp-snooping** command.

Variable	Value
<i>enable</i>	Enables or disables DHCP Snooping on the specified VLAN.

Configure Trusted and Untrusted Ports

Use the following procedure to set the trust factor associated with a port for DHCP Snooping. By default, the trust factor is set to untrusted.



Note

For ports that are members of an MLT, DHCP Snooping must be configured using the MLT configuration mode.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the trust factor for the port:

```
ip dhcp-snooping <trusted|untrusted>
```



Note

Configure **ip dhcp-snooping trusted** on a DHCP server switch port only.

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#ip dhcp-snooping trusted
```

Variable Definitions

The following table defines parameters for the **ip dhcp-snooping** command.

Variable	Value
<trusted untrusted>	Specifies the trust factor of the port for DHCP Snooping. Note: Configure ip dhcp-snooping trusted on a DHCP server switch port only.

Display DHCP Snooping Global Configuration

Use the following procedure to display the global status of DHCP Snooping configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the global configuration:
show ip dhcp-snooping

Example

```
Switch:1>show ip dhcp-snooping
=====
                        Dhcp Snooping General Info
=====
Dhcp Snooping                : Enabled
=====
```

Display DHCP Snooping Interface Information

Use the following procedure to view the DHCP Snooping interface information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display DHCP Snooping brouter port information:
show ip dhcp-snooping interface [gigabitEthernet [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} [vrf WORD<1-16> | vrfids WORD<0-512>] | <1-4059> [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} [vrf WORD<1-16> | vrfids WORD<0-512>] | vrf WORD<1-16> | vrfids WORD<0-512>] | vrf WORD<1-16> | vrfids WORD<0-512>]]
3. Display DHCP Snooping VLAN information:
show ip dhcp-snooping vlan <1-4059>
4. Display DHCP Snooping information for specific VRF name:
show ip dhcp-snooping vrf WORD<1-16>
5. Display DHCP Snooping information for specific VRF ID:
show ip dhcp-snooping vrfids WORD<0-512>

Examples

```
Switch:1>show ip dhcp-snooping interface gigabitEthernet
```

```
=====
```

Dhcp Snooping Interface Info		
PORT NUM	PORT CLASS	TRUNK ID
1/1	UNTRUSTED	none
1/2	UNTRUSTED	none
2/1	UNTRUSTED	none
2/2	UNTRUSTED	none

```
=====
```

All 4 out of 4 Total Num of Dhcp Snooping entries displayed

```
Switch:1>show ip dhcp-snooping vlan
```

```
=====
```

Dhcp Snooping Vlan Info			
VLAN ID	VRF NAME	ENABLE	ORIGIN
1	GlobalRouter	false	RADIUS
10	GlobalRouter	false	RADIUS
4051	GlobalRouter	false	RADIUS
4052	GlobalRouter	false	RADIUS

```
=====
```

All 4 out of 4 Total Num of Dhcp Snooping entries displayed

```
Switch:1>show ip dhcp-snooping binding vrfids 0
```

```
=====
```

DHCP Snooping Binding Table							
MAC ADDRESS	IP ADDRESS	PORT NUM	VLAN ID	VRF NAME	LEASE TIME	EXPIRY TIME	ENTRY TYPE
36:63:0e:73:03:fe	192.0.2.8	1/1/2	200	GlobalRouter	86400	83700	Learned
36:63:0e:73:03:ff	192.0.2.9	1/1/2	200	GlobalRouter	86400	83700	Learned

```
Static entries : 0
Learned entries : 2
Total entries : 2
=====
```

All 2 out of 2 Total DHCP Snooping binding entries displayed

```
Switch:1>show ip dhcp-snooping binding vrf vrf100
```

```
=====
```

DHCP Snooping Binding Table							
MAC ADDRESS	IP ADDRESS	PORT NUM	VLAN ID	VRF NAME	LEASE TIME	EXPIRY TIME	ENTRY TYPE
00:00:00:00:01:01	192.0.2.11	1/24	100	vrf100	Infinite	none	Static

```
Static entries : 1
Learned entries : 0
Total entries : 1
=====
```

All 1 out of 3 Total DHCP Snooping binding entries displayed

Add Static Entries to DHCP Snooping Binding Table

Use the following procedure to add devices with a specified MAC address to the DHCP Snooping binding table.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Add the static entry to the DHCP Snooping binding table:

```
ip dhcp-snooping binding <1-4059> 0x00:0x00:0x00:0x00:0x00:0x00 ip
```

```
{A.B.C.D} port {slot/port[/sub-port]} [expiry <0-2147483646>]
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# ip dhcp-snooping binding 1 00-14-22-01-23-45 ip 10.10.10.01 port 1/2
expiry 2
```

Variable Definitions

The following table defines parameters for the **ip dhcp-snooping binding** command.

Variable	Value
<1-4059>	Specifies the VLAN ID.
0x00:0x00:0x00:0x00:0x00:0x00	Specifies the MAC address of the DHCP client.
ip {A.B.C.D}	Specifies the IP address of the DHCP client.
port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}	Specifies the switch port to which the DHCP client connects.
expiry <0-2147483646>	Specifies the expiry time (in seconds) for the DHCP client.

Clear Entries from DHCP Snooping Binding Table

Use the following procedure to clear entries (static or dynamic) from the DHCP Snooping binding table.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. Enter:

```
clear ip dhcp-snooping binding [dynamic|static]
```

Example

```
Switch:1>enable
Switch:1#clear ip dhcp-snooping binding static
```

Variable Definitions

The following table defines parameters for the **clear ip dhcp-snooping binding** command.

Variable	Value
<i>static</i>	Clears static entries from the DHCP Snooping binding table.
<i>dynamic</i>	Clears dynamic entries from the DHCP Snooping binding table.

Display DHCP Snooping Binding Table Information

Use the following procedure to display the DHCP Snooping binding table, you can filter the entries displayed based on the type, port, or VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all binding entries:

```
show ip dhcp-snooping binding
```
3. Display binding entries based on the MAC address or IP address:

```
show ip dhcp-snooping binding address [0x00:0x00:0x00:0x00:0x00:0x00 | {A.B.C.D}]
```
4. Display binding entries configured on the ports:

```
show ip dhcp-snooping binding interface [gigabitEthernet{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```
5. Display binding entries configured on the VLANs:

```
show ip dhcp-snooping binding vlan <1-4059>
```
6. Display binding entries configured on a specific VRF:

```
show ip dhcp-snooping binding vrf WORD<1-16>
```
7. Display binding entries configured on a specific VRF ID:

```
show ip dhcp-snooping binding vrfids WORD<0-512>
```
8. Display a summary of the DHCP Snooping binding table:

```
show ip dhcp-snooping binding summary [<1-4059>] [vrf WORD<1-16>] [vrfids WORD<0-512>] [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```
9. Display binding entries based on the type of entry:

```
show ip dhcp-snooping binding type [dynamic|static]
```

Example

```
Switch:1>show ip dhcp-snooping binding
```

DHCP Snooping Binding Table							
MAC ADDRESS	IP ADDRESS	PORT NUM	VLAN ID	VRF NAME	LEASE TIME	EXPIRY TIME	ENTRY TYPE
23-74-44-33-15-33	192.0.2.40	225	1	13446	0	0	Static
ab-22-44-23-22-11	192.0.2.56	213	34	52341	0	0	Static

bb-22-44-33-af-ab	192.0.2.134	197	234	34345	0	0	Static
bb-22-44-af-af-ab	192.0.2.88	197	999	52342	0	0	Static
fe-92-44-33-22-33	192.0.2.13	211	333	52343	0	0	Static
fe-ab-44-33-22-33	192.0.2.45	197	74	52343	0	0	Static

```

Static entries   : 6
Learned entries  : 0
Total entries    : 6

```

Dynamic ARP Inspection Configuration Using CLI

The following section provides procedures to configure Dynamic ARP Inspection (DAI) using CLI.

Enable or disable Dynamic ARP Inspection on a VLAN

You must enable DAI separately for each VLAN. When you enable DAI on a specific VLAN, the ARP packets are captured and inspected on that VLAN. DAI is disabled by default.



Note

DAI cannot be enabled on Private VLANs (E-Tree) and SPBM B-VLANs.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable DAI on the VLAN:

```
ip arp-inspection enable
```

3. Disable DAI on the VLAN:

```
no ip arp-inspection enable
```

Example

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 5
Switch:1(config-if)#ip arp-inspection enable

```


Variable Definitions

The following table defines parameters for the **ip arp-inspection** command.

Variable	Value
<i>enable</i>	Enables or disables DAI on the specified VLAN.

Configure Trusted and Untrusted Ports

Use the following procedure to set the trust factor associated with a port for DAI. By default, the trust factor is set to untrusted.



Note

For ports that are part of an MLT, DAI must be configured using the MLT configuration mode.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} OR interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Set the trust factor for the port:

```
ip arp-inspection <trusted|untrusted>
```

Example

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/2  
Switch:1(config-if)#ip arp-inspection trusted
```

Variable Definitions

The following table defines parameters for the **ip arp-inspection** command.

Variable	Value
<i><trusted untrusted></i>	Specifies the trust factor of the port for DAI.

Display Dynamic ARP Inspection Interface Information

Use the following procedure to view the DAI interface information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display DAI brouter port information:

```
show ip arp-inspection interface [ gigabitEthernet [ {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] [ vrf WORD<1-16> | vrfids WORD<0-512> ] | <1-4059> [ {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] [ vrf WORD<1-16> | vrfids WORD<0-512> ] | vrf WORD<1-16> | vrfids WORD<0-512> ] ]
```
3. Display DAI VLAN information:

```
show ip arp-inspection vlan <1-4059>
```
4. Display DAI information for specific VRF name:

```
show ip arp-inspection vrf WORD<1-16>
```
5. Display DAI information for specific VRF ID:

```
show ip arp-inspection vrfids WORD<0-512>
```

Example

```
Switch:1>show ip arp-inspection interface gigabitEthernet 1/2
```

Arp Inspection Port Info		
PORT NUM	PORT CLASS	TRUNK ID
1/2	UNTRUSTED	none

All 1 out of 1 Total Num of Arp Inspection entries displayed

```
Switch:1>show ip arp-inspection vlan
```

Arp Inspection Vlan Info			
VLAN ID	VRF NAME	ENABLE	ORIGIN
1	GlobalRouter	false	RADIUS
2	GlobalRouter	false	RADIUS
20	GlobalRouter	false	RADIUS
55	GlobalRouter	true	RADIUS

All 4 out of 4 Total Num of Arp Inspection entries displayed

```
Switch:1>show ip arp-inspection vrfids 5
```

Arp Inspection Vlan Info			
VLAN ID	VRF NAME	ENABLE	ORIGIN

```

10          tt          true          CONFIG
-----

Switch:1>show ip arp-inspection vrf TT

=====
                        Arp Inspection Vlan Info
=====
VLAN      VRF
ID        NAME      ENABLE      ORIGIN
-----
10        tt          true          CONFIG
-----

```

FHS configuration

Configure IPv6 FHS features to enable IPv6 link security and management over the Layer 2 links.

Enable or Disable FHS Globally

About This Task

You must enable First Hop Security globally for RA Guard or DHCPv6 Guard to be operational.

Enabling FHS globally installs the required filters for FHS. Disabling FHS, uninstalls these filters. By default, FHS is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable First Hop Security globally:

```
ipv6 fhs enable
```
3. Disable First Hop Security globally:

```
no ipv6 fhs enable
```

OR

```
default ipv6 fhs enable
```

Manage the FHS IPv6 Access List

About This Task

You can create an FHS IPv6 access list or add IPv6 prefixes to an existing IPv6 access list.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create an FHS IPv6 access list or add IPv6 prefixes to an existing IPv6 access list:

```
ipv6 fhs ipv6-access-list [WORD<1-64>] [WORD<0-46>] [ge<0-128>] [le<0-128>] [mode <allow | deny>]
```
3. Delete an FHS IPv6 access list or delete a particular IPv6 prefix from the IPv6 access list:

```
no ipv6 fhs ipv6-access-list [WORD<1-64>] [WORD<0-46>]
```
4. Set the ge/le values and mode of the FHS IPv6 access list to default value:

```
default ipv6 fhs ipv6-access-list [WORD<1-64>] [WORD<0-46>] [ge|le|mode]
```

Example

```
Switch>enable
Switch#config t
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#ipv6 fhs ipv6-access-list ipv6_acl_1 fe80::221:2fff:fe31:5376/64
Switch(config)#
```

Variable Definitions

The following table defines parameters for the **ipv6 fhs ipv6-access-list** command.

Variable	Description
<i>WORD<1-64></i>	Specifies the IPv6 access list name.
<i>WORD<0-46></i>	Specifies the IPv6 address or the prefix length to be added to the IPv6 access list.
<i>ge <0 -128></i>	<p>Specifies the minimum value of prefix length advertised in prefix information of RA or DHCPv6 packets.</p> <p>By default, the value is equal to the configured prefix length.</p> <p>Note: If you manually configure the value, ensure that it is greater than the configured prefix length. Also ensure, the ge value is always less than the le value.</p>

Variable	Description
<code>le <0 -128></code>	Specifies the maximum value of prefix length advertised in prefix information of RA or DHCPv6 packets. By default, the value is equal to the configured prefix length. Note: If you manually configure the value, ensure that it is greater than the configured prefix length.
<code>mode <allow deny></code>	Specifies the access mode. By default, the value is allow.

Display FHS IPv6 Access List Information

About This Task

Displays the current FHS IPv6 access list information.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display the current FHS IPv6 access list information:
`show ipv6 fhs ipv6-access-list [WORD<1-64>]`

Example

```
Switch:1# show ipv6 fhs ipv6-access-list

=====
                        IPv6 FHS Access List Table Info
=====
                                MASK-RANGE
ACC-LIST-NAME      IPV6-PREFIX      MASK FROM TO      MODE
-----
v6_acl1            1:0:0:0:0:0:1      64   64   64   Allow
v6_acl2            1:0:0:0:0:0:1      64   64   64   Allow
-----

All 2 out of 2 Total Num of ipv6 access list entries displayed
```

Manage the FHS MAC Access List

About This Task

You can create an FHS MAC access list or add MAC addresses to an existing MAC access list.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`

2. Create an FHS MAC access list or add MAC addresses to an existing MAC access list:

```
ipv6 fhs mac-access-list WORD<1-64> <0x00:0x00:0x00:0x00:0x00:0x00>
[mode <allow | deny>]
```
3. Delete an FHS MAC access list or delete a particular MAC address from the MAC access list:

```
no ipv6 fhs mac-access-list WORD<1-64> <0x00:0x00:0x00:0x00:0x00:0x00>
```
4. Set the MAC ACL mode to its default value:

```
default ipv6 fhs mac-access-list WORD<1-64>
<0x00:0x00:0x00:0x00:0x00:0x00> [mode]
```

Variable Definitions

The following table defines parameters for the **ipv6 fhs mac-access-list** command.

Variable	Description
<i>WORD<1-64></i>	Specifies the MAC access list name.
<i><0x00:0x00:0x00:0x00:0x00:0x00></i>	Specifies the MAC address to be added or deleted.
<i>mode <allow deny></i>	Specifies the access mode. By default, the value is Allow

Display FHS MAC Access List Information

About This Task

Displays the current FHS MAC access list information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```
2. Display the current FHS MAC access list information:

```
show ipv6 fhs mac-access-list [WORD<1-64>]
```

Example

```
Switch#show ipv6 fhs mac-access-list
=====
IPv6 FHS Mac Access List Table Info
=====
ACC-LIST-NAME      MAC-ADDRESS      ACL-MODE
-----
List2              10:20:30:40:50:60    Allow
                   00:11:22:33:44:55    Deny
-----
All 1 out of 1 Total Num of MAC access list entries displayed
=====
```

Display Current FHS Configuration

About This Task

Displays the current FHS configuration.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the current FHS configuration:
show ipv6 fhs port-policy {slot/port[-slot/port][,...][slot/all][all]}

Example

```
Switch:1#show ipv6 fhs port-policy
```

```
=====
                        IPv6 FHS Port Policy Info
=====
PORT  DHCPG-DEVICE-ROLE  DHCPG-POLICY          RAG-DEVICE-ROLE  RAG-POLICY
-----
1/1   Server             dhcp_poll             Router           ra_poll
-----
```

```
All 1 out of 1 Total Num of fhs port policy entries displayed
```

DHCPv6 Guard Policy Configuration

DHCPv6 Guard policy blocks DHCPv6 reply and advertisement messages that originate from unauthorized DHCPv6 servers and relay agents that forward DHCPv6 packets from servers to clients.

Enable or Disable DHCPv6 Guard Globally

About This Task

Enabling DHCPv6 Guard globally installs filters on the configured interfaces. By default, DHCPv6 Guard is disabled.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Enable FHS globally:
ipv6 fhs enable
3. Enable DHCPv6 Guard globally:
ipv6 dhcp-guard enable
4. Disable DHCPv6 Guard globally:
no ipv6 dhcp-guard enable
5. Set DHCPv6 Guard to its default value:
default ipv6 dhcp-guard enable

Manage the DHCPv6 Guard Policy

About This Task

Configure or modify the DHCPv6 Guard policy.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create a DHCPv6 Guard policy:

```
ipv6 dhcp-guard policy WORD<1-64>
```
3. Delete a DHCPv6 Guard policy:

```
no ipv6 dhcp-guard policy WORD<1-64>
```



Note

You cannot delete a policy that is already attached to a port.

Variable Definitions

The following table defines parameters for the **ipv6 dhcp-guard policy** command.

Variable	Description
<i>WORD<1-64></i>	Specifies the created or deleted DHCPv6 Guard policy name.

Attach a DHCPv6 Guard Policy to a Port

About This Task

Applies a DHCPv6 Guard policy to a specific interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
```

```
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Apply a DHCPv6 Guard policy.

```
ipv6 fhs dhcp-guard attach-policy WORD<1-64>
```
3. Detach a DHCPv6 Guard policy from an interface.

```
no ipv6 fhs dhcp-guard attach-policy
```

OR

```
default ipv6 dhcp-guard attach-policy
```


4. Enable device role verification attached to the port. By default, router is selected.

```
ipv6 fhs dhcp-guard device-role {client|server} attach-policy
WORD<1-64>
```



Note

A DHCPv6 Guard policy can be attached to a port only if the device-role configured on that port is 'server'.

Variable Definitions

The following table defines parameters for the **ipv6 fhs dhcp-guard attach-policy** and **ipv6 fhs dhcp-guard device-role** command.

Variable	Description
<i>WORD<1-64></i>	Specify the name of the DHCPv6 Guard policy to be attached or detached.
<i>{client server}</i>	Sets the DHCPv6 Guard device role as client or server.

Configure DHCPv6 Guard in dhcp-guard Mode

About This Task

Configures DHCPv6 Guard under dhcp-guard mode.

Procedure

1. Enter DHCP-guard Configuration mode.

```
enable
```

```
configure terminal
```

```
ipv6 fhs dhcp-guard policy WORD<1-64>
```

2. Specify IPv6 access list to verify IPv6 source address of DHCPv6 packets..

```
match server access-list <ipv6-access-list-name>
```

3. Remove DHCPv6 Guard filtering for the sender's IPv6 addresses.

```
no match server access-list
```

OR

```
default match server access-list
```

4. Specify IPv6 prefix list to verify advertised prefixes.

```
match reply prefix-list <ipv6-prefix-list-name>
```

5. Remove DHCPv6 Guard filtering for advertised prefixes.

```
no match reply prefix-list
```

OR

```
default match reply prefix-list
```

6. Specify the minimum limit for verification of the advertised preference.

```
preference min-limit <0-255>
```
7. Set the minimum limit for verification of the advertised preference to its default value.

```
default preference min-limit
```
8. Specify the maximum limit for verification of the advertised preference.

```
preference max-limit <0-255>
```
9. Set the maximum limit for verification of the advertised preference to its default value.

```
default preference max-limit
```

Variable Definitions

The following table defines parameters for the **dhcp-guard** configuration mode commands.

Variable	Description
match server access-list <i><ipv6-access-list-name></i>	<p>Enables verification of the sender's IPv6 address in inspected messages from the configured authorized device source access list specified.</p> <p>Note: If the access-list is not attached, the IPv6 source address in DHCPv6 packet is not validated. If the list is attached and it does not match any entries in IPv6 access list, the switch drops the DHCPv6 packet. If you wish to change this behavior, add an entry with IPv6 prefix "0::0/0" with the Allow option, which changes the default drop to default Allow.</p>
{ no default } match server access-list	Removes the sender's IPv6 address based DHCPv6 Guard filtering.
match reply prefix-list <i><ipv6-prefix-list-name></i>	<p>Enables verification of the advertised prefixes in DHCPv6 reply messages from the configured authorized prefix list. If prefix-list is not configured, this check is bypassed.</p> <p>Note: If the access-list is not attached, the inspection does not occur. If the list is attached and advertised IPv6 address does not match any IPv6 prefixes in the list, the switch drops the DHCPv6 packet. If you wish to change this behavior, add an IPv6 access list entry with prefix 0::0/0 with the Allow option, which changes the default drop to default Allow.</p>
{ no default } match reply prefix-list	Removes the advertised prefix-based DHCPv6 Guard filtering.

Variable	Description
<code>preference min-limit<0-255></code>	Enables validation of advertised preference (in preference option) to check if it is greater than the specified limit. If preference is not specified, this field in the packet is not validated. While changing the preference limit, ensure the maximum limit is greater than the minimum limit.
<code>default preference min-limit</code>	Sets the specified limit to its default value. By default, the value is 0.
<code>preference max-limit<0-255></code>	Enables validation of advertised preference (in preference option) to check if it is less than the specified limit. If preference is not specified, this field in the packet is not validated. Note: The preference value in the packet is not validated if both minimum and maximum values are zero.
<code>default preference max-limit</code>	Sets the specified limit to its default value. By default, the value is 0.

Display the DHCPv6 Guard Policy

About This Task

Displays DHCPv6 Guard policy information for all the configured DHCPv6 Guard policies or a particular policy.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display DHCPv6 Guard policy information:

```
show ipv6 fhs dhcp-guard policy WORD<1-64>
```

Example

```
Switch:1# show ipv6 fhs dhcp-guard policy
```

```
=====
=====
                                IPv6 DHCP Guard Policy Info
=====
=====
POLICY-NAME          SERVER-ACC-LIST      REPLY-PREF-LIST      MIN-RTR-PREF  MAX-RTR-PREF
-----
dhcp_pol1            v6_acl1              v6_acl2                0              0
-----
-----
All 1 out of 1 Total Num of dhcp-guard stats entries displayed
```

Variable Definitions

The following table defines parameters for the `show ipv6 fhs dhcp-guard policy` command.

Variable	Description
<i>WORD<1-64></i>	Displays DHCPv6 Guard policy information for all the configured DHCPv6 Guard policies. Policy name is an optional parameter. If policy name is provided, only the DHCPv6 Guard policy of the specified policy-name is displayed.

RA Guard Configuration

IPv6 RA Guard provides support to the administrator to block or reject unwanted RA Guard messages that arrive at the network switch platform. The routers use Router Advertisements (RAs) to announce themselves on the link. The RA Guard feature analyzes these RAs and filters out bogus RAs sent by unauthorized routers. The RA Guard feature compares configuration information on the Layer 2 device with the information found in the received RA frame. After the Layer 2 device validates the content of the RA packet against the configuration, it forwards the RA to its destination. If the RA packet validation fails, the RA is dropped.

Enable or Disable RA Guard Globally

About This Task

Enables RA Guard globally. By default, RA Guard is disabled.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable FHS globally:
`ipv6 fhs enable`
3. Enable RA Guard globally:
`ipv6 fhs ra-guard enable`
4. Disable RA Guard globally:
`no ipv6 fhs ra-guard enable`
5. Set the RA Guard to its default value:
`default ipv6 fhs ra-guard enable`

Manage the RA Guard Policy

About This Task

Configure or modify RA Guard policy. This command also enables the RA Guard configuration mode.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create the RA Guard policy:
`ipv6 fhs ra-guard policy WORD<1-64>`
3. Delete the RA Guard policy:
`no ipv6 fhs ra-guard policy WORD<1-64>`



Note

You cannot delete a policy that is attached to a port.

Variable Definitions

The following table defines parameters for the **ipv6 fhs ra-guard policy** command.

Variable	Description
<i>WORD<1-64></i>	Specifies the name of the RA Guard policy to be created or deleted. This is a mandatory parameter in this command.

Configure RA Guard on an Interface

About This Task

Attaches or detaches a RA Guard policy on the specific interface.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
`enable`
`configure terminal`
`interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]`
`[,...][slot/all][all]}`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Apply a RA Guard policy.
`ipv6 fhs ra-guard attach-policy WORD<1-64>`



Note

RA Guard device-role on the port has to be configured as 'router' before attaching any RA Guard policy to the port. If device-role on the port is not 'router', this command will fail with an appropriate error message.

- Detach a RA Guard policy from an interface.

```
no ipv6 fhs ra-guard attach-policy
```

OR

```
default ipv6 fhs ra-guard attach-policy
```

- Enable device role verification attached to the port.

```
ipv6 fhs ra-guard device-role {router|host} attach-policy WORD<1-64>
```



Note

A DHCPv6 Guard policy can be attached to a port only if the device-role configured on that port is 'server'.

Variable Definitions

The following table defines parameters for the **ipv6 fhs ra-guard attach-policy** and **ipv6 fhs ra-guard device-role** command.

Variable	Description
<i>WORD<1-64></i>	Specifies the name of the RA Guard policy to be attached or detached.
<i>{host router}</i>	Sets the RA Guard device role as host or router.

Configure RA Guard in RA Guard Mode

About This Task

Configures RA Guard in the RA Guard configuration mode.

Procedure

- Enter RA Guard Configuration mode.

```
enable
```

```
configure terminal
```

```
ipv6 fhs ra-guard policy WORD<1-64>
```

- Configure the filter to match the IPv6 prefixes advertised in RA packets.

```
match ra-prefix-list WORD<1-64>
```

- Remove RA Guard filtering for the advertised prefixes.

```
no match ra-prefix-list
```

OR

```
default match ra-prefix-list
```

- Configure the filter to match the source MAC address of RA packets.

```
match ra-macaddr-list WORD<1-64>
```

5. Remove the source MAC address-based RA Guard filtering.

```
no match ra-macaddr-list
```

OR

```
default match ra-macaddr-list
```
6. Configure the filter to match source IPv6 address of RA packets.

```
match ra-srcaddr-list WORD<1-64>
```
7. Remove the source IPv6 address based RA Guard filtering.

```
no match ra-srcaddr-list
```

OR

```
default match ra-srcaddr-list
```
8. Enable managed address configuration flag verification in the advertised RA packet.

```
managed-config-flag <none | on | off>
```
9. Enable advertised hop count limit verification.

```
hop-limit {maximum | minimum} <0-255>
```
10. Enable the advertised default router-preference parameter value verification.

```
router-preference maximum {none | high | low | medium}
```

Variable Definitions

The following table defines parameters to configure RA Guard policy.

Variable	Description
match ra-prefix-list <i>WORD<1-64></i>	Verifies the advertised prefixes in RA packets against the configured authorized prefix list. Note: RA packet's sender IPv6 address is not validated if no IPv6 source access list is attached to the RA Guard policy. If the list is attached and if RA packet's sender IPv6 address does not match any entry in that IPv6 prefix list, then the RA packet is dropped. To change this behavior, add a entry with ipv6 prefix "0::0/0" with Allow option. The default value changes from Drop to Allow.
{no default} match ra-prefix-list	Removes the advertised prefix-based RA Guard filtering
match ra-macaddr-list <i>WORD<1-64></i>	Verifies sender's source MAC address against the configured mac-access-list. Note: Advertised prefixes in RA packet are not validated if no IPv6 prefix list is attached to the RA Guard policy. If the list is attached and if it does not match any MAC in the list, then the RA packet is dropped.

Variable	Description
{no default} match ra-macaddr-list	Removes the source MAC address-based RA Guard filtering for the specified MAC address access list names.
match ra-srcaddr-list <i>WORD<1-64></i>	Verifies sender's source IPV6 address against the configured list. Note: Inspection is not done if the access-list is not attached. If the list is attached and if it does not match any IPV6 in the list, then the RA packet is dropped. To change the behavior, add a dummy IPV6 "0:0:0:0:0:0" to the list with Allow option. The default value changes from Drop to Allow.
{no default} match ra-srcaddr-list	Removes the source IPV6 address-based RA Guard filtering for the specified IPV6 address access list names.
managed-config-flag <none on off>	Verifies managed address configuration flag in the advertised RA packet. By default, the value is none and check is bypassed.
hop-limit {maximum minimum} <0-255>	Verifies the advertised hop count limit. The limit value range is from 0 to 255. While changing the minimum or maximum value, ensure the maximum value is greater than the minimum value. By default, the minimum and maximum limit are 0. In this case, the hop-limit check is bypassed.
router-preference maximum {none high low medium}	Verifies if the advertised default router-preference parameter value is lower than or equal to a specified limit. By default, the value is none and the check is bypassed.

Display RA Guard Configuration

About This Task

Display configured RA Guard policy information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display configured RA Guard policy information:
`show ipv6 fhs ra-guard policy WORD<1-64>`

Example

```
Switch:1# show ipv6 fhs ra-guard policy
=====
IPv6 Ra Guard Policy Info
=====
POLICY-NAME  RA-SRC-ADDR-LIST  RA-MAC-ADDR-LIST  RA-PREFIX-LIST  LIMIT  LIMIT  MIN-HOP  MAX-HOP  MANAGED
CON-FLAG  PREF
-----
Ra_guard_poll  None             None             acl1             0      0      None     None
-----
All 1 out of 1 Total Num of ra-guard policy entries displayed
```


Variable Definitions

The following table defines parameters for the **show ipv6 fhs ra-guard policy** command.

Variable	Description
<i>WORD</i> <1-64>	Displays the RA Guard policy for the specified policy-name. By default, all the configured RA Guard policies are displayed.

IPv6 Neighbor Discovery inspection configuration

This section describes how to configure ND inspection on the switch and protect the network by mitigating the various types of attacks.



Important

Enable FHS globally before enabling ND inspection.

Enable ND Inspection Globally

Before You Begin

Enable FHS globally for ND inspection to work.

About This Task

Use this procedure to enable Neighbor Discovery (ND) inspection globally.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable ND inspection globally:
`ipv6 fhs nd-inspection enable`

Clear Neighbor Discovery Inspection Statistics

About This Task

Use this procedure to clear the Neighbor Discovery inspection statistics.

Procedure

1. Enter Privileged EXEC mode:
`enable`

2. Clear the Neighbor Discovery inspection statistics:

```
clear ipv6 fhs statistics nd-inspection [{slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]]
```

**Note**

Alternatively, you can use the command `clear ipv6 fhs statistics all` to clear the ND inspection statistics along with RA guard statistics and DHCPv6 Guard statistics.

Variable Definitions

The following table defines parameters for the **clear ipv6 fhs statistics nd-inspection** command.

Variable	Value
<code>{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

*Enable Neighbor Discovery Inspection on a VLAN***Before You Begin**

Enable FHS globally for ND inspection to work.

About This Task

Use this procedure to enable Neighbor Discovery inspection on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable Neighbor Discovery inspection on the VLAN:

```
ipv6 fhs nd-inspection enable
```

*Enable Neighbor Discovery Inspection on a Port***Before You Begin**

Enable FHS globally for ND inspection to work.

About This Task

Use this procedure to enable Neighbor Discovery inspection on a port

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable Neighbor Discovery inspection on the port:

```
ipv6 fhs nd-inspection enable
```

View Neighbor Discovery Inspection Status Globally

About This Task

Use this procedure to view the Neighbor Discovery inspection status globally.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the ND inspection status globally:

```
show ipv6 fhs status
```

View Neighbor Discovery Inspection Status on a Port

About This Task

Use this procedure to view Neighbor Discovery inspection status on a port.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display port-wise ND inspection status:

```
show ipv6 fhs port-policy
```

View Neighbor Discovery Inspection Statistics on a Port

About This Task

Use this procedure to view the Neighbor Discovery inspection statistics on a port or set of ports.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display ND inspection statistics on a port or a set of ports:
show ipv6 fhs statistics nd-inspection {slot/port[/sub-port][-slot/
port[/sub-port]][,...][slot/all][all]}

Variable Definitions

The following table defines parameters for the **show ipv6 fhs statistics nd-inspection** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

IPv6 DHCP Snooping Configuration

This section describes how to configure IPv6 DHCP snooping on the switch and protect the network by mitigating the various types of attacks.



Important

Configure DHCPv6 Guard before enabling IPv6 DHCP snooping. DHCPv6 Guard classifies the ports as trusted or untrusted and extracts DHCPv6 reply packets received on trusted ports to the control path. For more information on how to configure DHCPv6 Guard, see [DHCPv6 Guard Policy Configuration](#) on page 2361.

Create a Static Security Binding Table Entry

Use this procedure to enable learning Security Binding Table (SBT) entries on all the VLANs where IPv6 DHCP snooping is configured.

About This Task

Use this procedure to create a static SBT entry.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal

2. Add a static SBT entry:

```
ipv6 fhs snooping static-binding ipv6-address WORD<0-46> vlan <1-4059>
mac-address 0x00:0x00:0x00 port {slot/port[/sub-port]}
```

**Note**

To delete an SBT entry, use the command `no ipv6 fhs snooping static-binding`.

Example

Add a static SBT entry.

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)#ipv6 fhs snooping static-binding ipv6-address
2001:DB8:89ab:cdef:0123:4567:89ab:cdef vlan 1000 mac-address 00:11:22:33:44:55 port 1/2
```

Variable Definitions

The following table defines parameters for the **ipv6 fhs snooping static-binding ipv6-address** command.

Variable	Value
<i>mac-address</i> <i>0x00:0x00:0x00</i>	Specifies the MAC address of the binding entry.
<i>port {slot/port[/sub-port]}</i>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>vlan <1-4059></i>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>WORD<0-46></i>	Specifies the IPv6 address for the binding entry.

*Clear a Dynamic SBT Entry***About This Task**

Use this procedure to clear all or a particular dynamic SBT entry.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear a dynamic SBT entry:

```
clear ipv6 fhs snooping [vlan <1-4059>][ipv6-address WORD<0-46>]
```

Example

Clear a dynamic SBT entry on a VLAN.

```
Switch:1> enable
Switch:1>clear ipv6 fhs snooping vlan 1000 ipv6-address
2001:DB8:89ab:cdef:0123:4567:89ab:cdef
```

Variable Definitions

The following table defines parameters for the **clear ipv6 fhs snooping** command.

Variable	Value
<i>ipv6-address</i> <i>WORD<0-46></i>	Specifies the IPv6 address for the binding entry to clear. You cannot specify an address without first specifying the VLAN.
<i>vlan <1-4059></i>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. If you do not specify a VLAN, the command clears all entries.

Enable IPv6 DHCP Snooping on a VLAN

Before You Begin

Enable IPv6 DHCPv6 Guard for IPv6 DHCP snooping to work.

About This Task

Use this procedure to configure IPv6 DHCP snooping on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```
2. Configure IPv6 DHCP snooping on the VLAN:

```
ipv6 fhs snooping dhcp enable
```

View IPv6 DHCP Snooping and ND Inspection Status on a VLAN

About This Task

Use this procedure to view IPv6 DHCP snooping and ND inspection status on a VLAN.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

- View the IPv6 DHCP snooping and ND inspection status on a VLAN:

```
show ipv6 fhs status vlan [<1-4059>]
```

Example

View the status for all VLANs.

```
Switch:1#show ipv6 fhs status vlan
```

```
=====
                        IPv6 FHS VLAN Information
=====
VLAN-ID                DHCP-SNOOPING-STATUS      ND-INSPECTION-STATUS
-----
1                      Disabled                Disabled
3                      Disabled                Disabled
4                      Disabled                Disabled
22                     Disabled                Disabled
=====
All 4 out of 4 Total Num of FHS VLAN entries displayed
```

Variable Definitions

The following table defines parameters for the **show ipv6 fhs status vlan** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. If you do not specify a VLAN ID, the command output includes all VLANs.

View SBT Entries

About This Task

Use this procedure to view SBT entries.

Procedure

- Enter Privileged EXEC mode:
`enable`
- View all SBT entries:
`show ipv6 fhs snooping binding`
- View the SBT entries by type:
`show ipv6 fhs snooping binding type {dynamic | static}`
- View the SBT entries by VLAN:
`show ipv6 fhs snooping binding vlan <1-4059>[ipv6-address WORD<0-46>]`

Variable Definitions

The following table defines parameters for the **show ipv6 fhs snooping binding** command.

Variable	Value
<i>ipv6-address</i> <i>WORD<0-46></i>	Specifies the IPv6 address for the binding entry.
<i>type {dynamic static}</i>	Shows only dynamic binding entries or static binding entries.
<i>vlan <1-4059></i>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

IP Source Guard Configuration

The following sections provide procedural information you can use to configure IP Source Guard (IPSG) using the Command Line Interface (CLI).



Note

The switch supports configuration of IP Source Guard for both IPv4 and IPv6 addresses.

Enable IP Source Guard on a Port for IPv4 Addresses

Before You Begin

Ensure that the following conditions are all satisfied, before you enable IPSG on a port. Otherwise, the system displays error messages.

- DHCP Snooping is enabled globally.
- The port is a member of a VLAN that is configured with both DHCP Snooping and Dynamic ARP Inspection.
- The port is an untrusted port enabled with both DHCP Snooping and Dynamic ARP Inspection.
- The port has enough resources allocated, to support the maximum number of 10 IP addresses allowed for IPSG.

About This Task

Enable IP Source Guard (IPSG) on a port to add a higher level of security to the port by preventing IP spoofing. When you enable IPSG on the interface, filters are automatically installed for the IPv4 addresses that are already learned on that interface.



Important

Do not enable IPSG on MLT, DMLT, SMLT, LAG, trunk ports or ports that are a part of private VLANs.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable IPSG on the port:

```
ip source verify enable
```

3. Verify IPSG configuration:

```
show ip source verify interface gigabitEthernet [{slot/port[/sub-port]}
[-slot/port[/sub-port]][,...][slot/all][all]]
```

Example

Configure IPSG on port 1/1.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#ip source verify enable
```

Verify the configuration.

```
Switch:1(config-if)#show ip source verify interface gigabitEthernet 1/1
```

```
=====
Source Guard Port Info
=====
```

PORT NUM	ENABLE	IPSC ORIGIN
1/1	true	RADIUS

```
=====
```

All 1 out of 1 Total Num of Ip Source Guard entries displayed

Variable Definitions

The following table defines parameters for the **ip source verify** command.

Variable	Value
enable	Enables IP Source Guard on the port.

Disable IP Source Guard for IPv4 Addresses

About This Task

Disable IP Source Guard (IPSG) on a port to allow traffic from all IPv4 addresses to go through the port without being filtered.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable IPSG for IPv4 addresses:

```
no ip source verify
```

3. Verify IPSG configuration:

```
show ip source verify interface gigabitEthernet [{slot/port[/sub-port]
[-slot/port[/sub-port]][,...][slot/all][all]}
```

Example

Disable IPSG on port 1/1.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#no ip source verify
```

Verify the configuration.

```
Switch:1(config-if)#show ip source verify interface gigabitEthernet 1/1
```

```
=====
                        Source Guard Port Info
=====
```

PORT NUM	ENABLE	IPSC ORIGIN
1/1	false	RADIUS

```
=====
```

All 1 out of 1 Total Num of Ip Source Guard entries displayed

View IP Source Guard Configuration on a Port

About This Task

View IP Source Guard (IPSG) configuration on a port, with filters for IPv4 addresses.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View IPSG configuration:

show ip source verify interface gigabitEthernet [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]

Example

```
Switch:1>show ip source verify interface gigabitEthernet 1/1

=====
Source Guard Port Info
=====
PORT          IPSC
NUM           ENABLE  ORIGIN
-----
1/1           true    RADIUS
=====

All 1 out of 1 Total Num of Ip Source Guard entries displayed
```

Variable Definitions

The following table defines parameters for the **show ip source verify interface gigabitEthernet** command.

Variable	Value
{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}	<div>Identifies the slot and port in one of the following formats:</div> <div><ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all)</div> <div>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</div>

View IPv4 Address Bindings

About This Task

View the IPv4 address bindings that IP Source Guard (IPSG) allows.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the allowed IPv4 address bindings for a specific interface:

```
show ip source binding [interface gigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]|[vlan <1-4059>]|
[vrf WORD<1-16>]|[vrfids WORD<0-512>]
```
3. View the allowed IPv4 address bindings for a specific IP address:

```
show ip source binding {A.B.C.D}
```

Example

View the allowed IPv4 address bindings for the port 1/1.

```
Switch:1>show ip source binding interface gigabitEthernet 1/1
```

```
=====
IPSG Source Table
=====
```

PORT NUM	IP ADDRESS	VLAN ID	VRF NAME
1/1	192.0.2.1	200	GlobalRouter

```
-----
All 1 out of 1 Total IP Source Guard entries displayed
```

View the IPv4 address bindings for a specific IP address.

```
Switch:1>show ip source binding 192.0.2.1
```

```
=====
IPSG Source Table
=====
```

PORT NUM	IP ADDRESS	VLAN ID	VRF NAME
1/1	192.0.2.1	200	GlobalRouter

```
-----
All 1 out of 1 Total IP Source Guard entries displayed
```

Variable Definitions

The following table defines parameters for the **show ip source binding** command.

Variable	Value
<code>{A.B.C.D}</code>	Identifies the IPv4 address.
<code>interface gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code>vlan <1-4059></code>	Specifies the VLAN ID of the VLAN for which to view IPv4 address bindings.
<code>vrf WORD<1-16></code>	Specifies the VRF name of the VRF for which to view the IPv4 address bindings.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID of the VRF for which to view IPv4 address bindings.

Enable IP Source Guard on a Port for IPv6 Addresses

Before You Begin

Ensure that the following conditions are all satisfied, before you enable IPSPG on a port. Otherwise, the system displays error messages.

- DHCP Snooping is enabled globally.
- The port is a member of a VLAN that is configured with both DHCP Snooping and IPv6 Neighbor Discovery inspection.
- The port is an untrusted port enabled with both DHCP Snooping and IPv6 Neighbor Discovery inspection.
- The port has enough resources allocated, to support the maximum number of 10 IP addresses allowed for IPSPG.

About This Task

Enable IP Source Guard (IPSPG) on a port, to add a higher level of security to the port by preventing IP spoofing. When you enable IPSPG on the interface, filters are installed for IPv6 addresses that are already learned on that interface.



Important

Do not enable IPSPG on MLT, DMLT, SMLT, LAG, trunk ports or ports that are a part of private VLANs.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the maximum number of allowed IPv6 addresses on a port:

```
ipv6 source-guard [max-allowed-addr <2-10>]
```



Note

Ensure that you configure the maximum number of allowed IPv6 addresses on a port, before you enable IPSG on that port.

3. Enable IPSG on the port:

```
ipv6 source-guard enable
```

4. Verify IPSG configuration information on the port:

```
show ipv6 source-guard interface gigabitEthernet [{slot/port[/sub-  
port] [-slot/port[/sub-port]] [,...][slot/all][all]}
```

Example

Enable IPSG on a port.

Configure the maximum allowed IPv6 addresses on port 1/1 as 10 and enable IPSG on that port.

```
Switch:1>enable  
Switch:1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#interface gigabitEthernet 1/1  
Switch:1(config-if)#ipv6 source-guard max-allowed-addr 10  
Switch:1(config-if)#ipv6 source-guard enable
```

Verify the configuration.

```
Switch:1(config-if)#show ipv6 source-guard interface gigabitEthernet 1/1  
Slot/Port  Source Guard  Number of  IPv6  Address  
           Mode          address allowed  overflow count  
=====
```

1/1	Enabled	10	0
-----	---------	----	---

Optionally view all interfaces with IPSG enabled.

```
Switch:1(config-if)#show ipv6 source-guard interface enabled  
Slot/Port  Source Guard  Number of  IPv6  Address  
           Mode          address allowed  overflow count  
=====
```

1/1	Enabled	4	0
1/2	Enabled	9	0

Variable Definitions

The following table defines parameters for the **ipv6 source-guard** command.

Variable	Value
<i>enable</i>	Enables IP Source Guard on a port.
<i>max-allowed-addr</i> <2-10>	Specifies the maximum number of IPv6 addresses allowed to transmit data through the port. The default value is 4. Note: To reset the value to default, IPSPG must be disabled on the interface.

Disable IP Source Guard for IPv6 Addresses

About This Task

Disable IP Source Guard (IPSPG) on a port to allow traffic from all IPv6 addresses to go through the port without being filtered.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable IPSPG for IPv6 addresses on a port:

```
no ipv6 source-guard enable
```

3. Verify IPSPG configuration on the port:

```
show ipv6 source-guard interface gigabitEthernet [{slot/port[/sub-
port] [-slot/port[/sub-port]] [,...][slot/all][all]}
```

Example

Disable IPSPG on port 1/1.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#no ipv6 source-guard enable
```

Verify the configuration.

```
Switch:1(config-if)#show ipv6 source-guard interface gigabitEthernet 1/1
Slot/Port  Source Guard  Number of  IPv6  Address
           Mode      address allowed  overflow count
=====
1/1         Disabled      10          0
```

Clear IP Source Guard Overflow Counters

About This Task

Overflow counters consist of IPv6 addresses that are not added to IP Source Guard (IPSG) due to lack of filter resources. Use this procedure to clear the overflow counters for an IPSG port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Clear the overflow counters:

```
ipv6 source-guard overflow-count clear
```

3. Verify the configuration on the port:

```
show ipv6 source-guard interface gigabitEthernet [{slot/port[/sub-
port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

4. (Optional) View the overflow counters on all IPSG enabled ports:

```
show ipv6 source-guard interface enabled
```

Example

Clear overflow counters on the IPSG port 1/1.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#ipv6 source-guard overflow-count clear
```

Verify the configuration on port 1/1.

```
Switch:1(config-if)#show ipv6 source-guard interface gigabitEthernet 1/1
Slot/Port  Source Guard  Number of  IPv6  Address
           Mode      address allowed  overflow count
=====
1/1         Enabled      10          0
```


Optionally, view the overflow counters on all IPSP enabled ports.

```
Switch:1(config-if)#show ipv6 source-guard interface enabled
Slot/Port  Source Guard  Number of  IPv6  Address
           Mode    address allowed  overflow count
=====
1/1        Enabled      4           0
1/2        Enabled      9           0
```

View IP Source Guard Configuration for IPv6 Addresses

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View IPSP configuration on a specified interface:

```
show ipv6 source-guard interface gigabitEthernet [{slot/port[/sub-
port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```
3. View IPSP configuration on all IPSP enabled interfaces:

```
show ipv6 source-guard interface enabled
```

Example

```
Switch:1#show ipv6 source-guard interface gigabitEthernet 1/1

Slot/Port  Source Guard  Number of  IPv6  Address
           Mode    address allowed  overflow count
=====
1/1        Enabled      4           0

Switch:1#show ipv6 source-guard interface enabled

Slot/Port  Source Guard  Number of  IPv6  Address
           Mode    address allowed  overflow count
=====
1/1        Enabled      4           0
1/2        Enabled      9           0
```

Variable Definitions

The following table defines parameters for the **show ipv6 source-guard interface gigabitEthernet** command.

Variable	Value
enabled	Displays IPSPG configuration on all IPSPG enabled interfaces.
gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}	<p>Displays IPSPG configuration on the specified interface.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View IPv6 Address Bindings

About This Task

View the IPv6 address bindings that IP Source Guard (IPSPG) allows.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the allowed IPv6 address bindings:

```
show ipv6 source-guard binding [WORD<0-46>] [interface gigabitEthernet
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}]
```

Example

View the IPv6 address bindings for port 1/3.

```
Switch:1>show ipv6 source-guard binding interface gigabitEthernet 1/3
Slot/Port   IPv6   Address
=====
1/3         2001::10:10:0:1
1/3         fe80::210:94ff:fe00:550b
-----
```

View the IPv6 address bindings for a specific IPv6 address.

```
Switch:1>show ipv6 source-guard binding fe80::210:94ff:fe00:550b
Slot/Port   IPv6   Address
=====
1/3         fe80::210:94ff:fe00:550b
-----
```

Variable Definitions

The following table defines parameters for the **show ipv6 source-guard binding** command.

Variable	Value
<i>WORD</i> <0-46>	Identifies the IPv6 address.
interface gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Layer 2 Security Configuration using EDM

Use the following sections to help you configure Layer 2 security features and protect the network by mitigating various types of attacks, using Enterprise Device Manager (EDM).

For IPv4 deployments, configure:

- DHCP Snooping
- Dynamic ARP Inspection
- IP Source Guard for IPv4 addresses

For IPv6 deployments, configure:

- First Hop Security (FHS)



Note

FHS does not solve all cases of denial of services like blocking flooding of the IPv6 messages.

- DHCP Snooping and IPv6 Neighbor Discovery Inspection
- IP Source Guard for IPv6 addresses

Configure Dynamic ARP Inspection on VLANs

Use the following procedure to enable or disable DAI on one or more VLANs.



Note

DAI cannot be enabled on Private VLANs (E-Tree) and SPBM B-VLANs.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **ARP Inspection**.
3. Select the **ARP Inspection-VLAN** tab.
4. In the row for the VLAN, double-click **Enabled**, and select **true**.
5. Select **Apply**.

ARP Inspection-VLAN Field Descriptions

Use the data in the following table to use the **ARP Inspection-VLAN** tab.

Name	Description
VlanId	Specifies the VLAN ID.
Enabled	Specifies if DAI is enabled or disabled for the particular VLAN. By default, DAI is disabled.
Origin	Specifies the origin of Address Resolution Protocol (ARP) Inspection configuration on the VLAN. The supported values are: <ul style="list-style-type: none">• config - Set by the user.• radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.

Configure Dynamic ARP Inspection on Ports

Use the following procedure to set the trust factor associated with a port for DAI . By default, the trust factor is set to untrusted.



Note

For ports that are part of an MLT, DAI must be configured using the MLT configuration mode.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. In the navigation pane, expand **Configuration > IP**.

2. Select **ARP Inspection**.
3. Select the **ARP Inspection-port** tab.
4. In the row for the port, double-click **IfTrusted**, and select **trusted** or **untrusted**.
5. Select **Apply**.

ARP Inspection-port Field Descriptions

Use the data in the following table to use the **ARP Inspection-port** tab.

Name	Description
Port	Specifies the port on the switch.
IfTrusted	Specifies the trust factor for DAI on the specific port. By default, it is set as untrusted.

Configure FHS Globals

About This Task

Use this procedure to enable FHS to enable DHCPv6 Guard, RA Guard, and ND-inspection globally, and to configure the lifetime for these policies.

Procedure

1. From the navigation pane, expand **Configuration > IPv6**.
2. Select **FHS**.
3. Select the **Globals** tab.
4. Select FHS global options.
5. Select **Apply** to save the changes.
6. (Optional) Select **Refresh** to update the results.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Admin	Enables or disables the FHS policy.
RAGuardAdmin	Enables or disables the RA Guard policy.
DHCPv6GuardAdmin	Enables or disables the DHCPv6 Guard policy.
NdInspectAdmin	Enables or disables Neighbor Discovery inspection.

Create an IPv6 Access List

About This Task

Use this procedure to create an FHS IPv6 access list or add IPv6 prefixes to the existing IPv6 access list.

Procedure

1. In the navigation pane, expand **Configuration**—> **IPv6**.
2. Select **FHS**.
3. Select the **IPv6 Access List** tab.
4. Select **Insert**.
5. Configure the parameters for the IPv6 access list.
6. Select **Insert**.

IPv6 Access List Field Descriptions

Use the data in the following table to use the **IPv6 Access List** tab.

Name	Description
Name	Specify the IPv6 access list name to create the IPv6 access list.
Prefix	Specify the IPv6 prefix for adding it to the IPv6 access list.
PrefixMaskLen	Specify the prefix length for adding it to the IPv6 access list. The value range is from 0 to 128. By default, the value is 0.
MaskLenFrom	Specify the start mask length for providing the IPv6 range. The value range is from 0 to 128. By default, the value is set to the configured prefix length of the IPv6 access list entry.
MaskLenTo	Specify the end mask length for providing the IPv6 range. The value range is from 0 to 128. By default, the value is set to the configured prefix length of the IPv6 access list entry.
AccessType	Select the access type to allow or deny the entry. By default, the access type is allow.



Note

- **MaskLenFrom** and **MaskLenTo** must always be greater than or equal to the configured **PrefixMaskLen** for this IPv6 access list entry
- The **MaskLenFrom** value must always be less than or equal to the **MaskLenTo** value.

Create MAC Access List

About This Task

Use this procedure to create a MAC access list or add a MAC address to the existing MAC access list.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **FHS**.
3. Select the **MAC Access List** tab.
4. Select **Insert**.
5. Configure the parameters for the MAC access list.
6. Select **Insert**.

MAC Access List Field Descriptions

Use the data in the following table to use the **MAC Access List** tab.

Name	Description
Name	Specify a name to create a MAC access list.
Mac	Specify the MAC address to add to the MAC access list, in (xx:xx:xx:xx:xx:xx) format.
AccessType	Specify allow or deny. By default, the access type is allow.

Create DHCPv6 Guard Policy

About This Task

Use this procedure to create the DHCPv6 Guard policy to block DHCPv6 reply and advertisement messages that originate from unauthorized DHCPv6 servers and relay agents.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **FHS**.
3. Select the **DHCPv6 Guard Policy** tab.
4. Select **Insert**.
5. Configure the parameters for the DHCPv6 Guard policy.
6. Select **Insert**.
7. (Optional) Select **Refresh** to update the results.

DHCPv6 Guard Policy Field Descriptions

Use the data in the following table to use the **DHCPv6 Guard Policy** tab.

Name	Description
PolicyName	Specifies the policy name to create or modify DHCPv6 Guard policy.
ServerAccessListName	<p>Enables verification of the sender IPv6 address in the DHCPv6 reply or advertisement packets against attached IPv6 server access list.</p> <p>Note: If the access-list is not attached, the source IPv6 address is not validated. If the list is attached and it does not match with any entries in attached IPv6 access list, the switch drops the DHCPv6 packet. To change this behavior, add an entry in the IPv6 access list with prefix 0::0/0 with access type as allow, which changes the drop by default to allow by default.</p>
ReplyPrefixListName	<p>Enables verification of the advertised prefixes in DHCPv6 reply messages against the attached prefix list. If not configured, this check is bypassed.</p> <p>Note: If the access-list is not attached, the advertised address/prefix is not validated. If the list is attached and it does not match with any entries in attached IPv6 access list, the switch drops the DHCPv6 packet. To change this behavior, an entry in the IPv6 access list with prefix 0::0/0 with access type as allow, which changes the drop by default to allow by default.</p>
PrefLimitMin	<p>Enables verification if the advertised preference (in reference option) is greater than the specified limit. If not specified, this check does not occur.</p> <p>The value range is from 0 to 255.</p>
PrefixLimitMax	<p>Enables verification if the advertised preference (in preference option) is less than the specified limit. If not specified, this check does not occur.</p> <p>The value range is from 0 to 255.</p> <p>Note: If both the maximum and minimum limit is 0, this preference check is ignored.</p>

Create RA Guard Policy

About This Task

Use this procedure to create a RA Guard policy to block or reject unwanted or rogue RA messages that arrive at the network device platform.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **FHS**.
3. Select the **RA Guard Policy** tab.
4. Select **Insert**.
5. Configure the parameters for the RA Guard policy.
6. Select **Insert**.
7. (Optional) Select **Refresh** to update the results.

RA Guard Policy Field Descriptions

Use the data in the following table to use the **RA Guard Policy** tab.

Name	Description
PolicyName	Specifies the name of the RA Guard policy to be created or modified.
SrcAddrList	<p>Specify the IPv6 access list name to verify the sender IPv6 address in the RA packets against the attached IPv6 access list.</p> <p>Note: The source address in the RA packet is not validated if the access-list is not attached. If the list is attached and the IPv6 source address in RA packet does not match any IPv6-prefix in the list, then the RA packet is dropped. To change this behavior, add an entry in the IPv6 access list with prefix 0::0/0 with access type as allow. The default value changes from drop to allow.</p>

Name	Description
PrefixList	<p>Specify the IPv6 prefix list name to verify the advertised prefixes in the RA packet against the attached IPv6 prefix list.</p> <p>Note: Advertised prefixes are not validated if the access-list is not attached. If the list is attached and the advertised prefix in the RA packet does not match any IPv6-prefix in the list, then the RA packet is dropped. To change this behavior, add an entry in the IPv6 access list with prefix 0::0/0 with access type as allow. The default value changes from drop to allow.</p>
MacAddrList	<p>Specify the MAC list name to verify the sender source MAC address against the attached MAC access list.</p> <p>Note: The source MAC address in the RA packet is not validated if the access-list is not attached. If the list is attached and the source MAC address in the RA packet does not match any MAC address in the list, then the RA packet is dropped.</p>
ManagedConfigFlag	<p>Select the managed configuration flag to verify managed address configuration in the advertised RA packet.</p> <p>By default, none is selected and managed configuration flag validation is skipped.</p>
RouterPrefMax	<p>Select the router preference maximum to verify the if the advertised default router preference parameter value is lower than or equal to a specified limit.</p> <p>By default, none is selected and router preference validation is skipped.</p>

Name	Description
HopLimitMin	Specify the minimum hop limit to verify the advertised hop count limit. The value range is from 0 to 255 By default, minimum hop limit is 0.
HopLimitMax	Specify the maximum hop limit to verify the advertised hop count limit. The value range is from 0 to 255 By default, the maximum hop limit is 0 and If both HopLimitMin and HopLimitMax are set to 0, then the hop limit parameter in the RA packet is not validated.

Create Port to Policy Mapping

About This Task

Use this procedure to map a port to a RA Guard or DHCPv6 Guard policy, DHCPv6 Guard or RA Guard statistics.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **FHS**.
3. Select the **Port Policy Mapping** tab.
4. Select **Insert**.
5. Configure the parameters for the port policy mapping.
6. Select **Insert**.
7. (Optional) Select **Refresh** to update the results.

Port Policy Mapping Field Descriptions

Use the data in the following table to use the **Port Policy Mapping** tab.

Name	Description
IfIndex	Specify the port.
DHCPv6GuardPolicyName	Enter an already-created DHCPv6 Guard policy name to map it with the port.
RAGuardPolicyName	Enter an already-created RA Guard policy name to map it with the port.
Dhcpv6gDeviceRole	Select server or client configuration. The default is server.
RagDeviceRole	Select host or router configuration. The default is router.

Enable DHCP Snooping Globally

Use the following procedure to enable DHCP Snooping globally. If DHCP Snooping is globally disabled, the switch forwards DHCP reply packets (received on trusted or untrusted ports) to all ports.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Snooping**.
3. Select the **DHCP Snooping Globals** tab.
4. Select **Enabled**.
5. Select **Apply**.

DHCP Snooping Globals Field Descriptions

Use data in the following table to use the **DHCP Snooping Globals** tab.

Name	Description
Enabled	Enables DHCP Snooping globally. By default, DHCP Snooping is disabled.

Configure DHCP Snooping on VLANs

Use the following procedure to configure DHCP Snooping on a specific VLAN. If DHCP Snooping is globally disabled, the switch forwards DHCP reply packets (received on trusted or untrusted ports) to all ports.

If you enable DHCP Snooping globally, the agent determines whether to forward DHCP reply packets based on the DHCP Snooping mode of the VLAN and trusted state of the port.



Note

You cannot enable DHCP Snooping on Private VLANs (E-Tree) and SPBM B-VLANs.

Before You Begin

You must enable DHCP Snooping globally.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Snooping**.
3. Select the **DHCP Snooping-VLAN** tab.
4. In the row for the VLAN, double-click **DhcpSnoopingEnabled**, and select **true**.
5. Select **Apply**.

DHCP Snooping-VLAN Field Descriptions

Use the data in the following table to use the **DHCP Snooping-VLAN** tab.

Name	Description
VlanId	Specifies the VLAN ID.
DhcpSnoopingEnabled	Specifies if DHCP Snooping is enabled or disabled for the particular VLAN. By default, DHCP Snooping is disabled.
Origin	Specifies the origin of DHCP Snooping configuration on the VLAN. The supported values are: <ul style="list-style-type: none">• config - Set by the user.• radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.

Configure Trusted and Untrusted Ports

Use the following procedure to set the trust factor associated with a port for DHCP Snooping. By default, the trust factor is set to untrusted on all ports.



Note

For ports that are members of an MLT, DHCP Snooping must be configured using the MLT configuration mode.

Before You Begin

To enable DHCP Snooping on a port, you must enable DHCP Snooping globally.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Snooping**.
3. Select the **DHCP Snooping-port** tab.
4. In the row for the port, double-click **DhcpSnoopingIfTrusted**, and select **trusted** or **untrusted**.
5. Select **Apply**.

DHCP Snooping-port Field Descriptions

Use data in the following table to use the **DHCP Snooping-port** tab.

Name	Description
Port	Specifies the port on the switch.
DhcpSnoopingIfTrusted	Specifies if the switch ports are trusted for DHCP Snooping. By default, it is set as untrusted.

Create DHCP Binding Table Entries

Use the following procedure to add entries for devices with static IP addresses to the DHCP binding table.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **DHCP Snooping**.
3. Select the **DHCP Bindings** tab.
4. Select **Insert**.
5. In **VlanId**, type the VLAN ID.
6. In **MacAddress**, type the MAC address of the DHCP client.
7. In **AddressType**, select a value.
8. In **Address**, type the IP address of the DHCP client.
9. In **Interface**, select a port.
10. In **LeaseTime(sec)**, type the time in seconds.
11. Select **Insert**.
12. Select **Apply**.

DHCP Bindings Field Descriptions

Use data in the following table to use the **DHCP Bindings** tab.

Name	Description
VlanId	Specifies the VLAN to which the DHCP client belongs.
MacAddress	Specifies the MAC address of the DHCP client.
AddressType	Specifies the type of address. The default address type is IPv4.
Address	Specifies the IP address assigned to the DHCP client.
Interface	Specifies the interface to which the DHCP client connects.
LeaseTime(sec)	Specifies the lease time (in seconds) of the particular DHCP binding entry. The time range is 0 to 2147483646 seconds.
TimeToExpiry(sec)	Species the time of expiry (in seconds) of the DHCP binding entry.
EntryType	Specifies the type of the DHCP binding entry. <ul style="list-style-type: none"> • If the entry was created through DHCP snooping, the type is learned(1). • If the entry was created through a management operation, the type is static(2).

Create an SBT Entry

You can build a snooping binding table (SBT), which contains entries from only trusted devices or hosts. This SBT table is used to validate Neighbor Discovery (ND) packets.

About This Task

Use this procedure to create an SBT entry.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **FHS**.
3. Select the **Snoop Binding** tab.
4. Select **Insert**.
5. Configure the parameters for the snoop binding.
6. Select **Insert**.
7. (Optional) Select **Refresh** to update the results.

Snoop Binding Field Descriptions

Use the data in the following table to use the **Snoop Binding** tab. The system displays a subset of these fields if you select **Insert**.

Name	Description
VlanId	Specify the VLAN to which the snooped entry belongs.
Ipv6Address	Enter the IPv6 address assigned to the IPv6 host.
MacAddress	Enter the MAC address of the snooped entry.
InterfaceIndex	Specify the interface on which the entry is learnt.
EntryType	Indicates the type of entry - static (1) or dynamic (2).
EntrySource	Indicates the method entry was learnt from - static (1) or dhcp (2).
ValidTime	Indicates the valid time for the snooped entry.
TimeToExpiry	Indicates the time to expiry of the snooped entry.

Enable IP Source Guard on a Port for IPv4 Addresses

Before You Begin

Ensure that the following conditions are all satisfied before you enable IPSPG on a port. Otherwise, the system displays error messages.

- DHCP Snooping is enabled globally.
- The port on which you want to enable IPSPG is a member of a VLAN that is configured with both DHCP Snooping and Dynamic ARP Inspection.
- The port is an untrusted port enabled with both DHCP Snooping and Dynamic ARP Inspection.
- The port has enough resources allocated to support the maximum number of 10 IP addresses allowed for IPSPG.

About This Task

Enable IP Source Guard (IPSPG) to add a higher level of security to a desired port by preventing IP spoofing. When you enable IPSPG on the interface, filters are installed for IPv4 addresses that are already learned on that interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Source Guard**.
3. Select the **IP Source Guard-port** tab.
4. In the row for the applicable port, edit **Mode** to select **ip**.
5. Select **Apply** to save your changes.
6. Select **Refresh** to update the **IP Source Guard-port** tab.

IP Source Guard-port Field Descriptions

Use the data in the following table to use the **IP Source Guard-port** tab.

Name	Description
Port	Identifies the port on which to enable IPSPG.
Mode	Displays whether IPSPG is enabled on the port. The default is disabled.
Origin	Specifies the origin of Source Guard configuration on the port. The supported values are: <ul style="list-style-type: none">• config - Set by the user.• radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.

View IPv4 Address Bindings

View the IPv4 address bindings that IPSP allows.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Source Guard**.
3. Select the **IP Source Guard-addresses** tab.

IP Source Guard-addresses Field Descriptions

Use the data in the following table to use the **IP Source Guard-addresses** tab.

Field	Description
Port	Indicates the port on which IPSP is configured.
Type	Indicates the address type.
Address	Indicates the IPv4 address that is allowed by IPSP on the port.
Source	Indicates the source of the IPv4 address, which is DHCP Snooping.

Configure IP Source Guard on a Port for IPv6 Addresses

Before You Begin

Ensure that the following conditions are all satisfied, before you enable IPSP on a port. Otherwise, the system displays error messages.

- DHCP Snooping is enabled globally.
- The port is a member of a VLAN that is configured with both DHCP Snooping and IPv6 Neighbor Discovery inspection.
- The port is an untrusted port enabled with both DHCP Snooping and IPv6 Neighbor Discovery inspection.
- The port has enough resources allocated to support the maximum number of 10 IP addresses allowed for IPSP.

About This Task

Enable IPSP to add a higher level of security to a desired port, by preventing IP spoofing. When you enable IPSP on an interface, filters are automatically installed for the IPv6 addresses that are already learned on that interface.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Source Guard** tab.
4. Double-click **InterfaceState**.

5. Select a value: **true** or **false**.
6. Double-click **MaxAddr**.
7. Enter the maximum number of IPv6 addresses that are allowed to transmit data on the port.
8. (Optional) To clear the overflow counters, double-click **ClearOverflowCount**, and then select **true**.
9. Select **Apply** to save your changes.
10. Select **Refresh** to update the **Source Guard** tab.

Source Guard Field Descriptions

Use the data in the following table to use the **Source Guard** tab.

Name	Description
IfIndex	Specifies a value that uniquely identifies the port.
InterfaceState	Specifies the state of the interface. The default value is false.
MaxAddr	<p>Specifies the maximum number of IPv6 addresses allowed to transmit data through the port. The default value is 4.</p> <p>Note: To reset the value to default, IPSPG must first be disabled on the interface.</p>
OverflowCount	Specifies the number of IPv6 addresses for which filters are not added on the IPSPG port, due to a lack of filter resources. The default value is 0.
ClearOverflowCount	Specifies whether the overflow counter must be cleared. By default, the value is false.

View IPv6 Address Bindings

View the IPv6 address bindings that IPSPG allows.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6**.
3. Select the **Source Guard Binding** tab.

Source Guard Binding Field Descriptions

Use the data in the following table to use the **Source Guard Binding** tab.

Field	Description
IfIndex	Specifies a value that uniquely identifies the port.
IPv6Addr	Specifies the binding entry for the IPv6 address.

Layer 2 security example scenarios

The following sections describe configuration examples to configure Layer 2 security features for IPv4 and IPv6 deployments.

FHS Deployment Scenario

In the following example, the Layer 2 switch “SW-1” is connected to another Layer 2 switch “SW-2”, two hosts and a DHCP server. Switch “SW-2” is connected to two other hosts and a router. Out of the two hosts connected to SW-2, one is a malicious host, which can generate bogus RA packets to advertise route prefix, and can also generate bogus DHCP reply packets to configure wrong IPv6 address or wrong default gateway. By doing this, it tries denial-of-service or adversary-in-the-middle attacks. These attacks must be prevented as it affects all the nodes present in the Layer 2 network and FHS can be effective in preventing these attacks.

These attacks can spread over the entire Layer 2 network and thus can affect the hosts connected to SW-2 as well as the hosts connected to SW-1. If you enable FHS only on SW-2, then it could only save the nodes which are directly connected to it. To prevent the good node connected to SW-1 from these attacks, the SW-1 switch also should be FHS enabled.

The following figure shows the FHS deployment scenario topology.




Figure 175: FHS deployment topology

By default, all the ports are trusted, until you configure DHCPv6 Guard or RA Guard policies.

See the following procedures to configure FHS RA Guard and DHCPv6 Guard for the preceding topology.

Creating FHS IPv6 ACL

About This Task

Filter IPv6 traffic by creating IPv6 Access Control Lists (ACLs) and applying them to the interfaces similar to the way that you create and apply IPv4 named ACLs.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Create an IP ACL name (ipv6_acl_1) to match the source IPv6 address of the router connected to the interface.


```
ipv6 fhs ipv6-access-list ipv6_acl_1
fe80:0:0:0:cef9:54ff:feb4:9481/128 mode allow
```

3. Create an IP ACL name (ipv6_acl_1) to match the source IPv6 address of the DHCPv6-server connected to the interface.

```
ipv6 fhs ipv6-access-list ipv6_acl_1  
fe80:0:0:0:cef9:54ff:feb4:9481/128 mode allow
```

What to Do Next

Create a First Hop Security MAC ACL.

Creating an FHS MAC ACL

About This Task

Filter the IPv6 traffic by creating a MAC access list with the ACL mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a MAC ACL name (rtr_smac) to match the source MAC of the router connected to the interface 1/2.

```
ipv6 fhs mac-access-list mac_acl_1 00:11:22:33:44:66 mode allow
```

Create a DHCPv6 Guard Policy for the Router

About This Task

Create a DHCPv6 Guard policy to provide Layer 2 security to DHCPv6 clients by protecting them against rogue DHCPv6 servers.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enter DHCP Guard mode with the DHCPv6 Guard policy name (dhcpv6g_pol_1). The DHCPv6 Guard policy for the interface is connected to a router.

```
ipv6 fhs dhcp-guard policy dhcpv6g_pol_1
```

3. Configure the source IPv6 access list to allow only a DHCPv6 server replies that originate from the IPv6 address fe80:0:0:0:cef9:54ff:feb4:9481/128 and check the preceding IPv6 ACL configuration for ipv6_acl_1 list.

```
match server access-list ipv6_acl_1
```

4. Verify the prefixes sent in the DHCPv6 server reply message so that the ipv6_acl_2 IPv6 ACL configuration allows only the prefix 1000::1/64.

```
match reply prefix-list ipv6_acl_1
```

Create an RA Guard Policy for the Router

About This Task

Create a **rag_pol_1** RA Guard policy for the router and configure the source IPv6 access list to allow only the RA packets that originate from the source IPv6 address **fe80:0:0:0:cef9:54ff:feb4:9481/128**. This configuration verifies the prefixes sent in the RA packets.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enter the RA Guard mode and configure RA Guard policy (rag_pol_1) for the interface connected to a router.
`ipv6 fhs ra-guard policy rag_pol_1`
3. Configure the source IPv6 access list to allow only RA packets originating from the source IPv6 address fe80:0:0:0:cef9:54ff:feb4:9481/128.
`match ipv6 ra-srcaddr-list ipv6_acl_1`
4. Verify the prefixes sent in the RA packets so that the rtr_pip IPv6 ACL configuration allows only the prefix 60::0/64.
`match reply ra-prefix-list ipv6_acl_1`

Attach FHS policies to the Interfaces

About This Task

Attach the FHS policies to the interfaces.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure DHCPv6 Guard and RA Guard policies on the interface (1/2) that connects to the router.
`interface ethernet 1/2`

`ipv6 dhcp-guard attach-policy dhcpv6g_pol_1`

`ipv6 ra-guard attach-policy rag_pol_1`

IPv6 DHCP Snooping and ND Inspection Configuration Example

This section shows examples of IPv6 DHCP snooping and ND inspection configuration.

Enable DHCPv6 Guard, ND inspection, and First Hop Security.

```
ipv6 fhs dhcp-guard enable
ipv6 fhs nd-inspection enable
```

```
ipv6 fhs enable
```

Create VLAN 1000 and add port members.

```
vlan create 1000 type port-mstprstp 0
vlan members add 1000 1/1-1/10
```

Enable DHCPv6 snooping and ND inspection on VLAN 1000.

```
interface vlan 1000
ipv6 fhs snooping dhcp enable
ipv6 fhs nd-inspection enable
exit
```

Add static SBT entry.

```
ipv6 fhs snooping static-binding ipv6-address 2001:DB8:0:0:0001:02ff:fe03:0405 vlan 1000
mac-address 00:01:02:03:04:05 port 1/5
```

Set the DHCPv6 Guard device-role on port 1/1 of the device on which DHCPv6 Guard is configured.

```
interface gigabitEthernet 1/1
ipv6 fhs dhcp-guard device-role server
exit
```

Enable ND inspection on ports 1/2 through 1/10.

```
interface gigabitEthernet 1/2-1/10
ipv6 fhs nd-inspection enable
exit
```

View the status.

```
show ipv6 fhs port-policy
show ipv6 fhs status
show ipv6 fhs status vlan
show ipv6 fhs snooping binding
```

Configure IP Source Guard

The following section describes a simple configuration example to configure IP Source Guard (IPSG) on a port.

When you enable IPSG on a port, filters are installed for the IPv4 or IPv6 addresses that are already learned on that port.

Procedure

Enable DHCP Snooping globally on the switch and verify the configuration.

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable DHCP Snooping globally.

```
ip dhcp-snooping enable
```

3. Verify the configuration.

```
show ip dhcp-snooping
```

Enable DHCP Snooping and Dynamic ARP Inspection on the VLAN that the port is a member of.

4. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```

5. Enable DHCP Snooping on the VLAN.

```
ip dhcp-snooping enable
```

6. Verify the configuration.

```
show ip dhcp-snooping vlan <1-4059>
```

7. Enable Dynamic ARP Inspection on the VLAN.

```
ip arp-inspection enable
```

8. Verify the configuration.

```
show ip arp-inspection vlan <1-4059>
```

9. Verify that the port on which you want to configure IPSPG is a DHCP Snooping and a Dynamic ARP Inspection untrusted port.

```
show ip dhcp-snooping interface gigabitEthernet [{slot/port[/sub-port]}
[-slot/port[/sub-port]][,...][slot/all][all]]
```

```
show ip arp-inspection interface gigabitEthernet [{slot/port[/sub-
port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

Configure IPSPG on a port and verify the configuration.

10. Perform one of the following steps to configure IPSPG on a port, for IPv4 or IPv6 addresses.

- Enable and verify IPSPG on a port for IPv4 addresses:

- a. `ip source verify enable`
- b. `show ip source verify interface gigabitEthernet [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]`

- Enable and verify IPSPG on a port for IPv6 addresses:

- a. `ipv6 source-guard enable`
- b. `ipv6 source-guard [max-allowed-addr <2-10>]`



Note

The default value is 4. To reset the value to default, IPSPG must first be disabled on the interface.

- c. `show ipv6 source-guard interface gigabitEthernet [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]`

Examples

The following example describes how to enable IPSC on port 1/5 which is a member of VLAN 10, for IPv4 or IPv6 addresses.

```
Switch:1>en
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

Enable DHCP Snooping globally and verify the configuration.

```
Switch:1(config)#ip dhcp-snooping enable
Switch:1(config)#show ip dhcp-snooping

=====
                        Dhcp Snooping General Info
=====
Dhcp Snooping                : Enabled
-----
```

Enable DHCP Snooping and Dynamic ARP Inspection on a VLAN that the port is a member of VLAN 10.

```
Switch:1(config)#interface vlan 10
Switch:1(config-if)#show ip dhcp-snooping vlan 10

=====
                        Dhcp Snooping Vlan Info
=====
VLAN      VRF
ID        NAME      ENABLE  ORIGIN
-----
10        GlobalRouter  true    RADIUS
-----

All 1 out of 1 Total Num of Dhcp Snooping entries displayed

Switch:1(config-if)#ip arp-inspection enable
Switch:1(config-if)#show ip arp-inspection vlan 10

=====
                        Arp Inspection Vlan Info
=====
VLAN      VRF
ID        NAME      ENABLE  ORIGIN
-----
10        GlobalRouter  true    CONFIG
-----

All 1 out of 1 Total Num of Arp Inspection entries displayed
```

Verify that the port is DHCP Snooping and Dynamic ARP Inspection untrusted.

```
Switch:1(config-if)#show ip dhcp-snooping interface gigabitEthernet 1/5

=====
                        Dhcp Snooping Interface Info
=====
PORT      PORT      TRUNK
NUM       CLASS    ID
-----
1/5       UNTRUSTED  none
```

```
-----
All 1 out of 1 Total Num of Dhcp Snooping entries displayed
Switch:1(config-if)#show ip arp-inspection interface gigabitEthernet 1/5
```

```
=====
                        Arp Inspection Port Info
=====
```

PORT NUM	PORT CLASS	TRUNK ID
1/5	UNTRUSTED	none

```
-----
```

```
All 1 out of 1 Total Num of Arp Inspection entries displayed
```

Enable IPSG on port 1/5 for IPv4 addresses, and verify the configuration. This port is a member of VLAN 10.

```
Switch:1(config-if)#ip source verify enable
Switch:1(config-if)#show ip source verify interface gigabitEthernet 1/5
```

```
=====
                        Source Guard Port Info
=====
```

PORT NUM	ENABLE	IPSC ORIGIN
1/5	true	RADIUS

```
-----
```


```
All 1 out of 1 Total Num of Ip Source Guard entries displayed
```

Enable IPSG on port 1/1 for IPv6 addresses, and verify the configuration. This port is a member of VLAN 10.

```
Switch:1(config-if)#ipv6 source-guard enable
Switch:1(config-if)#ipv6 source-guard max-allowed-addr 10

Switch:1(config-if)#show ipv6 source-guard interface gigabitEthernet 1/1
Slot/Port Source Guard Number of IPv6 Address
           Mode          address allowed  overflow count
=====
```

1/1	Enabled	10	0
-----	---------	----	---



Licensing

[Licensing Types](#) on page 2413

[License Installation using CLI](#) on page 2420

[License Installation using EDM](#) on page 2423

This chapter describes the different types of licenses supported on switches that run Fabric Engine.

The following license types are available on switches that run Fabric Engine:

- Perpetual licenses (permanent):
 - [Base](#)—base features
 - [Premier](#)—advanced features
 - [MACsec](#)—MACsec license
- Subscription and Time-Based licenses:
 - Pilot or Extreme Platform ONE—enables Premier features while managed by Extreme Platform ONE, ExtremeCloud IQ, or ExtremeCloud IQ Site Engine
 - Factory Default Premier Trial—grants Premier features for 30 days; extendable up to three times.



Note

There is no trial license for MACsec.

Licensing Types

The following sections provide detailed information about the various licensing types.

Perpetual Licenses

Table 178: Perpetual licenses product support

Feature	Product	Release introduced
Perpetual licenses	4220 Series	Not Applicable
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Perpetual licenses provide permanent access to premier or MACsec features on supported the switches.

Base License

The Base License is included with the switch and enables the basic networking capabilities of the device. It does not include access to Premier features or MACsec. You must purchase additional licenses separately to enable Premier features and MACSec on the switch.

Premier License

The Premier License enables features not available in the Base License. You must purchase the Premier License.

The following table provides information on the features enabled by the license and the applicable order codes.

Feature support can differ between models within the same product family. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Table 179: Premier License features and part numbers

Platform License	Enabled Features	Part Number / Order Code
5320 Series Premier 5420 Series Premier 5520 Series Premier 5720 Series Premier	<ul style="list-style-type: none"> • DvR Controller (5520 Series and 5720 Series,) • Fabric Connect Layer 3 Virtual Services Networks (VSNs) • Extreme Integrated Application Hosting, including Fabric IPsec Gateway (5720 Series) • Greater than two BGP peers • OpenAPI • Greater than four OSPF active interfaces • Precision Time Protocol version 2 • Segmented VRF if using Layer 3 VSN 	5000-PRMR-LIC-P
7520 Series Premier 7720 Series Premier 7830 Series Premier	<ul style="list-style-type: none"> • DvR Controller (7520 Series and 7720 Series) • Fabric Connect Layer 3 Virtual Services Networks (VSNs) • Extreme Integrated Application Hosting, including Fabric IPsec Gateway (7520 Series and 7720 Series) • Greater than 16 BGP peers • OpenAPI • Greater than four OSPF active interfaces • Precision Time Protocol version 2 (7520 Series and 7720 Series) • Segmented VRF if using Layer 3 VSN (7520 Series and 7720 Series) • VXLAN Gateway (7520 Series and 7720 Series) 	7000-PRMR-LIC-P

MACsec License

The MACsec License enables MACsec on the switch. For more information about feature support, see *Fabric Engine and VOSS Feature Matrix*.



Note

You must purchase the MACsec License.

The following table provides information on the MACsec License and the applicable order codes.

Table 180: MACsec License order codes

Platform License	MACsec License	Part Number / Order Code
5320 Series Premier 5420 Series Premier 5520 Series Premier 5720 Series Premier	IEEE 802.1AE MACsec	5000-MACSEC-LIC-P
7520-48YE-8CE Premier 7830 Series Premier	IEEE 802.1AE MACsec	7000-MACSEC-LIC-P

Install and Uninstall Licenses on Universal Switches

This topic explains how to obtain, activate, and transfer perpetual licenses on the switch.

Obtain and Install a License

For universal hardware switches, you can activate a license on the switch using either of the following methods:

- Manual activation by obtaining the license activation file from the [Extreme Networks Support portal](#).
- Automated activation using ExtremeCloud IQ. For information about using ExtremeCloud IQ for bulk automated license activation, go to <https://www.extremenetworks.com/support/documentation/extremecloud-iq>.

To obtain and install a license, follow these steps:



Note

You should have received by email a license voucher after your purchase of a switch. You will need the voucher ID number on this email to generate a license.

For manual license generation and installation:

1. Upgrade the switch so that it is running the most recent GA version of the software.

- To download the latest GA software, visit the [Extreme Networks Support portal](#).
2. Follow manual activation instructions for license generation available on the [Extreme Networks Support portal](#): **Products > ExtremeSwitching > Universal > Fabric Engine (VOSS) (for your switch) > Activation Instructions**.
 3. Download the license file onto the switch.

The license file name must meet the following requirements:

- Maximum of 63 alphanumeric characters
 - No spaces or special characters allowed
 - Underscore (_) is allowed
 - The file extension is `.lic`.
4. Install the license, using the command `load-license license_filename`.

The software verifies the license and activates the features that correspond to the license category.

For more information, see either of the following:

- [License Installation using CLI](#) on page 2420
- [License Installation using EDM](#) on page 2423

Transfer a License

You can permanently remove a license, which allows you to transfer the license to another switch. This should only be done when preparing to return a defective switch for a replacement switch (RMA).

For more information, see either of the following:

- [Revoke a License](#) on page 2421 (using CLI)
- [Revoke a License](#) on page 2424 (using EDM)

To transfer a license from a defective switch to a replacement switch (RMA), follow these steps.

1. Go to the [Extreme Networks Support portal](#).
2. Select **Assets > Licenses Home**.
3. Select **License Transfer**.
4. Enter the serial number of the defective unit, replacement unit serial number, and the RMA/case number.
5. Generate the new license, following the process in step 2 in [Obtain and Install a License](#) on page 2416.
6. Download the license file onto the switch.
7. Install the license, using the command `load-license license_filename`.

Subscription Licenses

Table 181: Subscription licenses product support

Feature	Product	Release introduced
Subscription licenses applied via cloud-management applications Note: Requires ExtremeCloud IQ Site Engine 25.2.10 or later ExtremeCloud IQ and Extreme Platform ONE Networking are automatically kept at the latest version.	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.2
	5420 Series	Fabric Engine 9.2
	5520 Series	Fabric Engine 9.2
	5720 Series	Fabric Engine 9.2
	7520 Series	Fabric Engine 9.2
	7720 Series	Fabric Engine 9.2
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	Not Supported
	VSP 7400 Series	Not Supported

This section provides details about subscription licenses applied to the switch using a cloud-management application, such as Extreme Networks Platform ONE, ExtremeCloud IQ, or ExtremeCloud IQ Site Engine.

For more information about cloud management support for subscription licenses, see the ExtremeCloud IQ, ExtremeCloud IQ Site Engine, and Extreme Networks Platform ONE documentation at <https://www.extremenetworks.com/support/documentation/>.

Pilot or Extreme Platform ONE License

The Pilot or Extreme Platform ONE License enables use of Premier features, but it does not include MACsec.



Note

MACsec, if supported, requires a perpetual MACsec License.

The following considerations apply to Pilot or Extreme Platform ONE Licenses:

- Premier License and Pilot or Extreme Platform ONE License can coexist on the same switch.
- If the switch has a Premier License, the Premier features remain enabled if the Pilot or Extreme Platform ONE License is revoked or if it expires.
- If you power on a switch with an expired license, the switch has 60 minutes to obtain updated license information from the cloud management application. If the switch fails to connect after 60 minutes, it revokes the Pilot or Extreme Platform ONE License.
- You can reassign a validated subscription license to a different switch in the following cases:
 - You removed or unmanaged a switch from the cloud management application.

- The switch is disconnected from the cloud management application and the 30-day factory default license expired.

Factory Default Premier Trial License

Table 182: Factory Default Premier Trial License

Feature	Product	Release introduced
Ability to extend the Factory Default Premier Trial License	4220 Series	Not Applicable
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.4
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.4
	VSP 7400 Series	VOSS 8.4



Note

The Trial License does not apply to 4220 Series.

New hardware switches include a Factory Default Premier Trial License to use all features, excluding MACsec. You can run the **extend-time-period** command up to three times to extend the trial license. You can configure all features, except MACsec, without restrictions, and save the configuration.



Note

There is no Trial License for MACsec.

You cannot configure new features after the trial period, but the switch continues to run with the existing configured features. If you reboot the switch after the trial period, and a valid software license is not present, licensed features in the configuration are not loaded.

You must install or apply a valid license to enable licensed features.

The initial trial period is 30 days.



Note

The evaluation period is based on the switch System Up Time.

If a Premier License was installed and then revoked, the trial period cannot be extended.

To extend the factory default trial license, see [Extend the Factory Default Premier Trial License](#) on page 2422.

License Installation using CLI

Install and manage a license file for the switch by using the Command Line Interface (CLI).



Note

This section applies to multiple platforms. The command syntax and example outputs may not be identical on all hardware platforms.

Obtain and Install a License File

Licenses are tied to the Base MAC address.



Note

Extreme Networks supports only a single host (system MAC address) for each license file. You cannot use the same license file on multiple hosts.

After you generate the license through the Extreme Networks Support Portal at <https://extremeportal.force.com/ExtrLicenseLanding>, perform the following steps to install the license on the switch to enable licensed features.

Before You Begin

- File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP) support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.
- You must enable File Transfer Protocol (FTP) or Trivial File Transfer Protocol (TFTP) in the boot configuration flags depending on which protocol you use to download the license file to the device.

Procedure

1. From a remote station or PC, use FTP or TFTP to download the license file to the device and store the license file in the /intflash directory.
2. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

3. Load the license:

```
load-license WORD<1-64>
```

Examples

Log in to the device and load the license. The following example shows a successful operation.

```
Switch:1(config)#load-license license_Switch_example
Switch:1(config)#2021-09-01T13:48:27.275Z Switch_example CP1 - 0x000006ef - 00000000
GlobalRouter SW INFO License key successfully loaded from <license_Switch_example>.
```

Variable Definitions

The following table defines parameters for the **copy** command.

Variable	Value
<code><a.b.c.d></code>	Specifies the IPv4 and IPv6 address of the TFTP server from which to copy the license file.
<code><file></code>	Specifies the name of the license file when copied to the flash.
<code><srcfile></code>	Specifies the name of the license file on the TFTP server.

The following table defines parameters for the **load-license** command.

Variable	Value
<code>WORD<1-64></code>	Specifies the name of the license file in the <code>/intflash</code> directory.

Revoke a License

Revoke a perpetual license on the switch if you need to transfer the license to a different switch.

You can permanently remove a license, which allows you to transfer the license to another switch. This should only be done when preparing to return a defective switch for a replacement switch (RMA).

About This Task

The **no license** command invalidates the feature license and generates a revocation certificate, which is the first step to releasing the license entitlement back to the license entitlement manager (LEM). The revocation certificate is contained in a file where the first part of the file name is the switch serial number and the file extension is .rvk.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Revoke a license:

```
no license <10G WORD<1-64> | macsec | premier>
```



Note

Beginning with Fabric Engine 9.0, 5320 Series switches no longer require a license to operate SFP+ ports at full 10 Gbps. Use the `10G WORD<1-64>` parameter to revoke a previously installed license.

Example

Revoke the Premier license.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#no license premier
Successfully revoked license of type <PRD-5000-PRMR>, revocation certificate available at
"/intflash/2011G-00165-PRD-5000-PRMR.rvk"
```

Extend the Factory Default Premier Trial License

Use the following procedure to extend the Factory Default Premier Trial License on your switch.

You can run the **extend-time-period** command up to three times to extend the evaluation license in 30-day increments for an additional 90 days.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Extend the trial software license on your device:
`extend-time-period`



Note

You must reboot your switch after each license extension.

Example

Extend the trial license period for 30-days:

```
Switch:1>enable
Switch:1#extend-time-period
Are you sure you want to reset the box to apply changes? (y/n) y
```

Show a License File

Display the existing software licenses on your device. If the switch uses a Factory Default Premier Trial License, the output shows the time remaining in the trial period.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the existing software licenses on your device:
`show license`

Example



Note

Example outputs might not be identical on all hardware platforms.

The following output shows a switch with a factory default premier license:

```
Switch:1>enable
Switch:1#show license

BASE License: Available
PREMIER License: Trial
Perpetual MACSEC License: Not Applicable
Subscription License Level: Trial, expires in 29 days
```

The following output shows a switch with a subscription license:

```
Switch:1>show license

BASE License: Available
PREMIER License: Subscription
Perpetual MACSEC License: Not Available
Subscription License Level: STANDARD, expires in 25 days
```

License Installation using EDM

Install and manage a license file for the switch by using Enterprise Device Manager (EDM).



Note

This section applies to multiple platforms. The fields may not be identical on all hardware platforms.

Install a License File

Install a perpetual license file on the switch to enable licensed features.

Before You Begin

- You must store the license file on a file server.

About This Task

IPv4 and IPv6 addresses are supported.

For additional information about the tabs referenced in this procedure, see the following sections:

- **Edit > File System > Copy File** tab: [Copy File Field Descriptions](#) on page 222
- **Edit > Chassis > System** tab: [System Field Descriptions](#) on page 2425

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **File System**.
3. Select the **Copy File** tab.
4. In the **Source** box, type the IP address of the file server where the license file is located and the name of the license file.

5. In the **Destination** box, type the flash device and the name of the license file.
The license file name must have a file extension of .xml.
6. Select **start**.
7. Select **Apply**.
The license file is copied to the flash of the device. The system displays the status of the file copy in the Result field.
8. In the navigation pane, expand **Configuration > Edit**.
9. Select **Chassis**.
10. Select the **System** tab.
11. In **ActionGroup1**, select **loadLicense**.
12. In **LicenseFileName** box, type the name of the license file.
13. Select **Apply**.

**Important**

If the loading fails, the switch cannot unlock the licensed features and reverts to base functionality.

14. On the **System** tab, in **ActionGroup1**, select **saveRuntimeConfig**.
15. Select **Apply**.

Revoke a License

Revoke a perpetual license on the switch if you need to transfer the license to a different switch.

You can permanently remove a license, which allows you to transfer the license to another switch. This should only be done when preparing to return a defective switch for a replacement switch (RMA).

About This Task

The revocation parameters invalidate the feature license and generate a revocation certificate, which is the first step to releasing the license entitlement back to the license entitlement manager (LEM). The revocation certificate is contained in a file where the first part of the file name is the switch serial number and the file extension is .rvk.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System** tab.
4. In **ActionGroup1**, select one of the following:
 - **revokeLicenseMacsec**
 - **revokeLicensePremier**
5. Select **Apply**.
6. In **ActionGroup1**, select **saveRuntimeConfig**.
7. Select **Apply**.

System *Field Descriptions*

Use the data in the following table to use the **System** tab.

Name	Description
sysDescr	Shows the system assigned name and the software version.
sysUpTime	Shows the elapsed time since the system last started.
sysContact	Configures the contact information.
sysName	Configures the name of this device.
sysLocation	Configures the physical location of this device.
VirtualIpAddr	Configures the virtual IP address that the CPU advertises and stores in the switch configuration file.
VirtualNetMask	Configures the net mask of the virtual management IP address.
VirtualIpv6Addr	Specifies the virtual IPv6 address.
VirtualIpv6PrefixLength	Specifies the length of the virtual IPv6 address prefix (in bits).
LastChange	Displays the time since the last configuration change.
LastVlanChange	Displays the time since the last VLAN change.
LastStatisticsReset	Displays the time since the statistics counters were last reset.
LastRunTimeConfigSave	Displays the last run-time configuration saved.
DefaultRuntimeConfigFileName	Displays the default Run-time configuration file directory name.
ConfigFileName	Specifies the name of a new configuration file.

Name	Description
ActionGroup1	<p>Performs one of the following actions:</p> <ul style="list-style-type: none"> • resetCounters— Resets all statistic counters. • saveRuntimeConfig— Saves the current run-time configuration. • loadLicense— Loads a software license file to enable features. • revokeLicensePremier— Revokes the Premier license and generates a revocation certificate in XML format. • revokeLicenseMacsec—Revokes the MACsec license and generates a revocation certificate in XML format. • revokeLicense10G4P— Revokes the 4-port 10 Gbps Port License on 5320 Series and generates a revocation certificate in XML format. • revokeLicense10G8P— Revokes the 8-port 10 Gbps Port License on 5320 Series and generates a revocation certificate in XML format.
LicenseFileName	Specifies the name of the license file in the <code>/intflash</code> directory.
ActionGroup2	<p>Specifies the following action:</p> <ul style="list-style-type: none"> • resetIstStatCounters—Resets the IST statistic counters
ActionGroup3	<p>Can be the following action:</p> <ul style="list-style-type: none"> • flushIpRouteTbl—flushes IP routes from the routing table
ActionGroup4	<p>Can be the following action:</p> <ul style="list-style-type: none"> • softReset—resets the device without running power-on tests • resetConsole—resets the switch console • softResetCoreDump—reset with coredump
Result	Displays a message after you select Apply .
LocatorLED	Configures the system Locator LED on or off. The default is off.

View License File Information

About This Task

View information about the license file for the switch.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **License** tab.

License *Field Descriptions*

Use the data in the following table to use the **License** tab.

Name	Description
LicenseType	Indicates the level type of the current license.



Link Aggregation Control Protocol

[Link Aggregation Overview](#) on page 2428

[MultiLink Trunking with LACP](#) on page 2429

[LACP Configuration Considerations](#) on page 2432

[Important Information and Restrictions](#) on page 2433

[LACP configuration using CLI](#) on page 2434

[LACP configuration using EDM](#) on page 2447

The following sections provide the concepts and procedures you need to configure the Link Aggregation Control Protocol (LACP) to dynamically aggregate links as they become available to a trunk group.

Link Aggregation Overview

Link aggregation provides link level redundancy and increases load sharing. Use Link aggregation to bundle the ports into a port group, which is represented as one logical interface to the Media Access Control (MAC) layer.

The switch supports the following types of link aggregation:

- MultiLink Trunking (MLT)—a statically configured link bundling method. MLT is not standards based, but it interoperates with static link methods of other vendors.
- IEEE 802.3ad based link aggregation, through the Link Aggregation Control Protocol (LACP), dynamically aggregates links as they become available to a trunk group. Link Aggregation Control Protocol dynamically detects whether links can be aggregated into a link aggregation group (LAG) and does so after links become available. Link Aggregation Control Protocol also provides link integrity checking at Layer 2 for all links within the LAG.

Both MLT and IEEE 802.3ad based link aggregation are point-to-point functions.

The switch software offers LACP functionality layered with MLT. This document uses the term MLT with LACP to refer to this functionality.

Split MultiLink Trunking (SMLT)

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency by providing for the addition of failure redundancy with subsecond failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. Use SMLT to connect a device that supports some form of link aggregation, be it a switch or a server, to two distinct separate SMLT endpoints or switches. These SMLT devices form

a virtualized Switch Cluster through the SPBM cloud and are referred to as a Virtual Inter-Switch Trunk (vIST) Core Switch pair.

For more information, see the following sections:

- [Virtual Inter-Switch Trunk \(vIST\)](#) on page 2657
- [Simplified Virtual-IST](#) on page 2659
- [Split MultiLink Trunking](#) on page 2662

LACP with SMLT

You can use LACP on SMLT configurations. The switch provides modifications to the LACP in SMLT configurations. LACP-capable devices can connect to an SMLT aggregation pair.



Note

Virtual IST is not supported on LACP-enabled MLTs.

VLACP with SMLT

You can also configure Virtual LACP (VLACP) with an SMLT configuration. VLACP is a modification that provides end-to-end failure detection. VLACP is not a link aggregation protocol.

VLACP implements link status control protocol at the port level. This mechanism periodically checks the end-to-end status of a point-to-point or end-to-end connection. You can run VLACP on single ports or on ports that are part of an MLT.



Note

Do not configure VLACP on LACP-enabled ports. VLACP does not operate properly with LACP.

MultiLink Trunking with LACP

MultiLink Trunking (MLT) with Link Aggregation Control Protocol (LACP) manages ports and port memberships to form a link aggregation group (LAG). Use Link Aggregation Control Protocol to gather one or more links to form a LAG, which a Media Access Control (MAC) client treats as a single link. Link Aggregation Control Protocol can dynamically add or remove LAG ports, depending on availability and state.

IEEE 802.3ad overview

The IEEE 802.3ad standard comprises service interfaces, the LACP, the Marker Protocol, link aggregation selection logic, a parser or multiplexer, frame distribution, and frame collection functions.

The following illustration shows the major functions of IEEE 802.3ad defined as multiple link aggregation.




Figure 176: Link aggregation sublayer (according to IEEE 802.3ad)

The link aggregation sublayer comprises the following functions:

- frame distribution

This block takes frames submitted by the MAC client and sends them for transmission on the appropriate port based on a frame distribution algorithm employed by the Frame Distributor.

Frame distribution also includes an optional Marker Generator/Receiver used for the Marker Protocol. The switch only implements the Marker Receiver function. For more information about the frame distribution algorithm, see [MLT Traffic Distribution Algorithm](#) on page 2661.

- frame collection

This block passes frames received from the various ports to the MAC client. Frame collection also includes a Marker Responder used for the Marker Protocol.

- aggregator parser or multiplexers

During transmission operations, these blocks pass frame transmission requests from the Distributor, Marker Generator, and Marker Responder to the appropriate port.

During receive operations, these blocks distinguish among Marker Request, Marker Response, MAC Client Protocol Data Units (PDU), and pass the blocks to the appropriate entity (Marker Responder, Marker Receiver, and Collector, respectively).

- aggregator

The combination of frame distribution and collection, and aggregator parser or multiplexers.

- aggregation control

This block configures and controls link aggregation. It incorporates LACP for the automatic communication of aggregation capabilities between systems and automatic configuration of link aggregation.

- control parser/multiplexers

During transmission operations, these blocks pass frame transmission requests from the aggregator and Control entities to the appropriate port.

During receive operations, these blocks distinguish Link Aggregation Control Protocol Data Units (LACPDUs) from other frames. The blocks pass, passing the LACPDUs to the appropriate sublayer entity and all other frames to the aggregator.

802.3ad link aggregation principles

Use link aggregation to group ports together to form a link group to another device. Link groups increase aggregate throughput between devices and provide link redundancy.

Link aggregation employs the following principles and concepts:

- A MAC client communicates with a set of ports through an aggregator, which presents a standard IEEE 802.3 service interface to the MAC client. The aggregator binds to one or more ports within a system.
- The aggregator distributes frame transmissions from the MAC client to various ports, collects received frames from the ports, and transparently passes the frames to the MAC client.
- A system can contain multiple aggregators serving multiple MAC clients. A port binds to a single aggregator at a time. A MAC client is served by a single aggregator at a time.
- The Link Aggregation Control function binds ports to aggregators within a system. The control function aggregates links, binds the system ports to an appropriate aggregator, and monitors conditions to determine if a change in aggregation is needed. Network managers can manually provide link aggregation control by manipulating the link aggregation state variables (for example, keys). You can also use LACP to automatically determine, configure, bind, and monitor link aggregation.
- LACP uses peer exchanges across links to continually determine the aggregation capability of the links and provide the maximum level of aggregation capability between a pair of systems.
- Frame ordering is maintained for certain sequences of frame exchanges between MAC Clients. The distributor ensures that all frames of a conversation pass to a single port. The collector passes frames to the MAC client in the order they are received from the port. The collector can select frames received from the aggregated ports.

Because the frames are not ordered on a single link, this guarantees that frame ordering is maintained for all conversations.

- Conversations move among ports within an aggregation for load balancing and for maintaining availability if a link fails.
- Each port is assigned a unique, globally administered MAC address.

After entities initiate frame exchanges within the link aggregation sublayer, the source address is the MAC address. An example of an entity that initiates frame exchanges is LACP and Marker Protocol exchanges.

- Each aggregator is assigned a unique, globally administered MAC address that is used from the perspective of the MAC client, both as a source address for transmitted frames and as the destination address for received frames. You can use one of the port MAC addresses in the associated LAG as the MAC address of the aggregator.

Input/output port redundancy

You can use the MLT link aggregation mechanism to protect I/O ports. MLT is compatible with 802.3ad static, and provides a load sharing and failover mechanism to protect against module, port, fiber, or complete link failures.

You can use MLT with Link Access Control Protocol (LACP) disabled or use LACP enabled by itself.

LACP Configuration Considerations

You can configure priorities, keys, modes, and timers for the LACP.

LACP Priority

You can configure LACP priority at the system and port level as follows:

- Port priority—determines which ports are aggregated into a LAG that has more than eight ports configured to it.
- System priority—generates the switch ID after communicating with other systems. As a best practice, use the default value. If you need to change it, first disable the LACP, and then enable it again after you change the value.

LACP Keys

You must use the LACP keys to determine which ports are eligible for link aggregation. The LACP keys are defined by the ports after you configure the multilink trunk. You can aggregate the ports key that match the MLT key into that multilink trunk.

- Keys do not have to match between two LACP peers.

LACP Timers

You can customize failover times by changing the LACP timer attributes (fast periodic time, slow periodic time, and aggregate wait time). Values are set by default to match the IEEE 802.3ad values. If you change the values, they must match on the ports participating in aggregation between two devices.

Changes to LACP timer values at the global level are reflected on all ports. However, you can change the LACP timer values for each port level. After you change an LACP timer globally, this value is set on all ports. The global timer value overwrites the local port value irrespective of the LACP state. You must configure port values that differ from the global values.

The switch software uses the following LACP timers:

- fast periodic timer—200 to 20 000 milliseconds (ms); default 1000 ms
- slow periodic timer—10 000 to 30 000 ms; default 30000 ms
- aggregation-wait timer—200 to 2000; default 2000

You cannot aggregate a link if it does not receive an LACPDU for a period of timeout \times slow periodic time = 3×30 seconds = 90 seconds. If you use the fast periodic time, the timeout period is 3×1000 ms = 3 seconds. You must make timer changes to all ports participating in link aggregation and to the ports on the partnering node.

Configuration changes to the LACP timers are not reflected immediately. Link Aggregation Control Protocol timers do not reset until the next time you restart LACP globally or on a port. This ensures consistency with peer switches.

After you enable LACP on a port, the timer values are set at the port level. You must toggle the LACP status after timer values change. This does not impact existing ports unless you toggle the LACP status on the port.

LACP modes

LACP uses two active and passive modes.

- Active mode—ports initiate the aggregation process. Active mode ports aggregate with other active mode ports or passive mode ports.
- Passive mode—ports participate in LACP but do not initiate the aggregation process. You must partner passive mode ports with active mode ports for aggregation to occur.

LACP and Private VLANs

When using LACP, configure the private-vlan at the interface level.

Important Information and Restrictions

Before you use the switch, review the following important information and restrictions.

LACP with Simplified vIST/SPB NNI links

You should not use LACP on SPB network-to-network interface (NNI) MLT links or on the Simplified Virtual IST.

vIST VLAN IP addresses

Do not configure a Rendezvous Point (RP) or Bootstrap Router (BSR) on the vIST VLAN because you cannot ping them outside of the vIST VLAN subnet. When you enter the **ip pim enable** command on the vIST VLAN, the following message displays:

WARNING: Please do not use virtual IST VLAN IP address for BSR and RP related configurations, as unicast packets to virtual IST vlan IP address from outside of virtual IST vlan subnet will be dropped. Use Loopback or CLIP interface IP address for BSR and RP related configurations.

Simplified vIST and egress port-based filters

On Simplified vIST nodes, egress port-based filters may not work for IP multicast routed traffic because vIST internal filter rules (to prevent duplicate traffic) have higher precedence than user-created filters.



Note

- This issue is specific to IP multicast routed traffic only.
- Egress port-based filters work for Layer 2 multicast, broadcast and unicast traffic.

LACP configuration using CLI

This section describes how to configure and manage link aggregation using the Command Line Interface (CLI), including Link Aggregation Control Protocol (LACP), to increase the link speed and redundancy for higher availability.

MultiLink Trunking (MLT) with LACP manages switch ports and port memberships to form a link aggregation group (LAG). Configure LACP to allow dynamic bundling of physical ports to form a single logical channel.

You can describe the LACP in terms of link aggregation operations within a single system. You can configure a single piece of equipment so it contains more than one system (from the point of view of the link aggregation operation).

Before You Begin

- Changes to LACP made at the global level overrides and resets all port level settings.



Important

After you globally configure the LACP system priority, it applies to all LACP-enabled aggregators and ports. After you enable the LACP on an aggregator or port, it uses the global system priority value.

- After you make a timer change, restart the LACP (globally or on the port) so the changes are consistent across the link.

**Important**

Configuration changes to LACP timers are not reflected immediately. LACP timers are not reset until the next time LACP is restarted globally or on a port. This action ensures consistency with peer switches.

- The switch does not support standby ports for LACP aggregation groups.

Configure Global LACP Parameters

Configure LACP parameters globally. After you configure the LACP system priority globally, it applies to all LACP-enabled aggregators and ports. After you enable the LACP on an aggregator or a port, it uses the global system priority value.

A change to the global parameter configuration takes effect after you restart the LACP globally or on each port.

About This Task

**Important**

Changes made at the global level override and reset all port level settings.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Change the system priority:
`lacp system-priority <0-65535>`
3. Configure additional LACP parameters as required:
`lacp enable [aggr-wait-time <200-2000>] [fast-periodic-time <200-20000>] [slow-periodic-time <10000-30000>] [smlt-sys-id <0x00:0x00:0x00:0x00:0x00:0x00>] [system-priority <0-65535>] [timeout-scale <2-10>]`

If you do not configure the optional parameters, the system uses the default values.

Example

```
Switch:1(config)#lacp fast-periodic-time 2000
Switch:1(config)#lacp enable
```

Variable Definitions

The following table defines parameters for the **lACP** command.

Variable	Value
<i>aggr-wait-time</i> <200-2000>	Configures the aggregation wait time (in milliseconds) globally. The default value is 2000.
<i>enable</i>	Enables LACP globally. The default value is disabled.
<i>fast-periodic-time</i> <200-20000>	Configures the fast periodic time (in milliseconds) globally. The default value is 1000.
<i>slow-periodic-time</i> <10000-30000>	Configures the slow periodic time globally. The default value is 30000.
<i>smlt-sys-id</i> <0x00:0x00:0x00:0x00:0x00:0x00> Note: Exception: not supported on 4220 Series or 5320 Series.	Configures the LACP system ID globally. Enter a MAC address in the following format: 0x00:0x00:0x00:0x00:0x00:0x00.
<i>system-priority</i> <0-65535>	Configures the LACP system priority globally. The default value is 32768.
<i>timeout-scale</i> <2-10>	Configures the timeout scale globally. The default value is 3.

Configure LACP on a Port



Important

Changes made at the global level override and reset the port-level configuration.

Configure LACP on a port to enable or disable LACP on the selected ports.

You must use the LACP keys to determine which ports are eligible for link aggregation. The LACP keys are defined by the ports after you configure the multilink trunk. You can aggregate the ports key that match the MLT key into that multilink trunk.

Before You Begin

When enabling or disabling LACP on a port, as a best practice, disable the port first and re-enable the port after the configuration is complete.

About This Task

The minimum LACP configuration is as follows:

- Assign a given key to a set of ports. In the following procedure steps, you must assign the key before you enable LACP on the port.

- Assign the same key to an MLT with no members. The ports will automatically become MLT members.

Keys do not need to match between two LACP peers.

A port can operate in active or passive mode. You can configure a port to be an individual link or an aggregated link.

**Note**

When using LACP with private VLANs, configure the private VLAN at the interface level.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Assign an LACP key:

```
lacp key <1-512|defVal>
```

3. Change the LACP mode:

```
lacp mode <active|passive>
```

4. Change the port priority:

```
lacp priority <0-65535>
```

5. (Optional) Change the system priority:

```
lacp system-priority <0-65535>
```

6. Configure aggregation for the port:

```
lacp aggr-wait-time <200-2000>[aggregation enable]
```

If you do not configure the optional parameter, the system uses the default values.

- Configure parameters for the partner device at the opposite end of the link:



Note

All parameters beginning with **partner** are for debug purposes only. If you configure these commands locally and there is a mismatch with what is learned from the partner, there will be trace or log messages.

```
lacp partner-key <0-65535|defVal>

lacp partner-port <0-65535>

lacp partner-port-priority <0-65535>

lacp partner-state <0-255 | 0x0-0xff>

lacp partner-system-id 0x00:0x00:0x00:0x00:0x00:0x00

lacp partner-system-priority <0-65535>
```

- Configure additional LACP parameters as required:

```
lacp enable [fast-periodic-time <200-20000>] [slow-periodic-time
<10000-30000>] [timeout-time <long|short>] [timeout-scale <2-10>]
```

If you do not configure the optional parameters, the system uses the default values.

- Enable LACP aggregation:

```
lacp aggregation enable
```

Examples

Configure LACP on ports 1/2 and 1/3:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/2-1/3
Switch:1(config-if)#lacp key 100
Switch:1(config-if)#lacp aggregation enable
Switch:1(config-if)#lacp enable
```

Return the key to the default value:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/2-1/3
Switch:1(config-if)#no lacp aggregation enable
Switch:1(config-if)#default lacp key
Switch:1(config-if)#lacp aggregation enable
```

Variable Definitions

The following table defines parameters for the **lACP** command.

**Note**

All parameters beginning with **partner** are for debug purposes only. If you configure these commands locally and there is a mismatch with what is learned from the partner, there will be trace or log messages.

Variable	Value
<i>aggr-wait-time</i> <200-2000>	Configures the aggregation wait time (in milliseconds) for this port. The default is 2000.
<i>aggregation enable</i>	Enables aggregation on the port, which makes it an aggregated link.
<i>enable</i>	Enables LACP for this port. The default is disabled.
<i>fast-periodic-time</i> <200-20000>	Configures the fast periodic time (in milliseconds) for this port. The default is 1000 ms.
<i>key</i> <1-512 defVal>	Configures the aggregation key for this port. Enter the aggregation key value or defVal (1024 + IfIndex) To return a configured key to the default value, you must first disable LACP aggregation on the port.
<i>mode</i> { <i>active</i> <i>passive</i> }	Configures the LACP mode to be active or passive.
<i>partner-key</i> <0-65535>	Configures the partner administrative key.
<i>partner-port</i> <0-65535>	Configures the partner administrative port value.
<i>partner-port-priority</i> <0-65535>	Configures the partner administrative port priority value.
<i>partner-state</i> <0-255 0x0-0xff>	Configures the partner administrative state bitmask. Specify the partner administrative state bitmap in the range 0x0-0xff. The bit to state mapping is Exp, Def, Dis, Col, Syn, Agg, Time, and Act. For example, to set the two partner-state parameters <ul style="list-style-type: none"> • Act = true • Agg = true specify a value of 0x05 (bitmap = 00000101).
<i>partner-system-id</i> <0x00:0x00:0x00:0x00:0x00:0x00>	Configures the partner administrative system ID. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00.
<i>partner-system-priority</i> <0-65535>	Configures the partner administrative system priority value.
<i>priority</i> <0-65535>	Configures the port priority. The default value is 32768.
<i>slow-periodic-time</i> <10000-30000>	Configures the slow periodic time for this port. The default is 30000 ms.

Variable	Value
<code>system-priority <0-65535></code>	Configures the system priority for this port. The default is 32768.
<code>timeout-scale <2-10></code>	Configures a timeout scale for this port. The default value is 3. The LACP timeout is equal to the slow periodic time or fast periodic time multiplied by the timeout-scale, depending how you configure the timeout-time variable.
<code>timeout-time {long short}</code>	Configures the timeout to either long or short.

Configure LACP on an MLT

Configure an MLT with LACP to use the dynamic link aggregation function.

About This Task



Important

Attach ports to an aggregator only if their system priorities are the same; otherwise, they are considered to be operating in two different switches. You can attach ports to an aggregator only if their keys are the same.

When you add a VLAN to a dynamic MLT, only the active ports of the MLT are added as port members of the VLAN. Ports configured with the same aggregation key, but not active, are not added to the VLAN. If these inactive ports become active later, the system does not automatically add them to the VLAN port member list.

You must add all inactive ports to the VLAN. If you do not add the inactive ports to the VLAN, when they become active later, hashing can result in choosing a newly active port for traffic forwarding. Because the port is not a port member of the VLAN, traffic will be dropped. When you add the VLAN to the MLT, also add the inactive aggregation ports to the VLAN. You may need to disable LACP on the inactive ports before you can add them to the VLAN. Because the ports are inactive, disabling LACP does not cause a traffic interruption.

Similarly when you remove a VLAN from a dynamic MLT, all active ports of the MLT are removed from the VLAN port member list but the inactive members are not removed. You must remove the inactive aggregation members from the VLAN.

If you later configure a port for the same aggregation, you must add this port to all VLANs that are members of the MLT.

Procedure

1. Enter MLT Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface mlt <1-512>
```

2. Configure LACP on an MLT:

```
lacp enable [key <0-512>] [system-priority <0-65535>]
```

Example

Enable LACP and configure the LACP key on MLT 3:

```
Switch:1(config)# interface mlt 3
Switch:1(config-mlt)# lacp enable key 1281
```

Variable definitions

Use the data in the following table to use the **lacp** command.

Variable	Value
<i>enable</i>	Enables LACP on the MLT interface.
<i>key <0-512></i>	Configures the LACP aggregator key for a specific MLT. <ul style="list-style-type: none">• <i>0-512</i> is the LACP actor admin key.
<i>system-priority <0-65535></i>	Configures the LACP system priority for a specific MLT. <ul style="list-style-type: none">• <i>0-65535</i> is the system priority.

Configure LACP and Private VLANs

About This Task

Use the following procedure to configure Link Aggregation Control Protocol (LACP) on a private VLAN.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable trunking:

```
encapsulation dot1q
```

3. Set private VLAN port type:

```
private-vlan [isolated|promiscuous|trunk]
```

Example

```
Switch:1(config)#interface GigabitEthernet 1/6
Switch:1(config-if)#encapsulation dot1q
Switch:1(config-if)#private-vlan promiscuous
Switch:1(config-if)#exit

Switch:1(config)#interface GigabitEthernet 1/37
Switch:1(config-if)#encapsulation dot1q
Switch:1(config-if)#private-vlan promiscuous
```

View LACP Configuration Information

View LACP configuration information to determine the LACP parameters and to ensure your configuration is correct.

Procedure

1. View the global configuration:

```
show lacp
```

2. View LACP administrative information for the local device:

```
show lacp actor-admin interface [gigabitethernet]
```

OR

```
show lacp actor-admin interface gigabitethernet [vid <1-4059>] [{slot/
port[-slot/port] [,...] [slot/all] [all]}]
```

3. View LACP operational information for the actor device:

```
show lacp actor-oper interface [gigabitethernet]
```

OR

```
show lacp actor-oper interface gigabitethernet [vid <1-4059>] [{slot/
port[-slot/port] [,...] [slot/all] [all]}]
```

4. View LACP timer information:

```
show lacp extension interface [gigabitethernet]
```

5. View LACP interface configuration information

```
show lacp interface
```

OR

```
show lacp interface gigabitethernet [vid <1-4059>] [{slot/port[-slot/
port] [,...] [slot/all] [all]}]
```

OR

```
show lacp interface mlt [<64-6399>]
```

OR

```
show lacp interface mlt [id<1-512>]
```


6. View LACP administrative information for the partner device:

```
show lacp partner-admin interface [gigabitethernet]
```

OR

```
show lacp partner-admin interface gigabitethernet [vid <1-4059>]
[{{slot/port[-slot/port]}[,...][slot/all][all]}]
```

7. View LACP operational information for the partner device:

```
show lacp partner-oper interface [gigabitethernet]
```

OR

```
show lacp partner-oper interface gigabitethernet [vid <1-4059>]
[{{slot/port[-slot/port]}[,...][slot/all][all]}]
```

Example

```
Switch:1# show lacp
```

```
=====
                        Lacp Global Information
=====
      SystemId: 00:24:7f:a1:70:00
      SmltSystemId: 00:00:00:00:00:00
      Lacp: enable
      system-priority: 32768
        timeout-admin: 3
      fast-periodic-time-admin: 1000
      slow-periodic-time-admin: 30000
        aggr-wait-time-admin: 2000
        timeout-oper: 3
      fast-periodic-time-oper: 2000
      slow-periodic-time-oper: 30000
        aggr-wait-time-oper: 2000
```

In the following example output, aggr indicates the port has become part of an aggregation. indi indicates individual.

```
Switch:1(config-if)# show lacp actor-admin interface gigabitethernet 1/1
```

```
=====
                        Actor Admin
=====
INDEX SYS  SYS      KEY  PORT  PORT  STATE
  PRIO  ID
-----
1/1  32768  00:24:7f:9f:c0:00 1    0x114  32768 act short aggr
```

```
Switch:1(config-if)# show lacp partner-admin interface gigabitethernet 1/1
```

```
=====
                        Partner Admin
=====
INDEX SYS  SYS      KEY  PORT  PORT  STATE
  PRIO  ID
-----
1/1  0      00:00:00:00:00:00 0    0x0    0      pas          long indi
```

Variable definitions

Use the data in the following table to use the **show lacp** command.

Variable	Value
<code>actor-admin interface gigabitethernet [vid <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/ all][all]}</code>	<p>Shows LACP actor (or local) administrative information for all interfaces or the specified interface.</p> <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> is the port or port list.
<code>actor-oper interface gigabitethernet [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/ all][all]}</code>	<p>Shows LACP actor operational information for all interfaces or the specified interface.</p> <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> is the port or port list.

Variable	Value
<code>extension interface gigabitethernet [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/ all][all]}]</code>	Shows LACP timer information for all interfaces or the specified interface. <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> is the port or port list.
<code>interface [gigabitethernet [vid <1-4059> [{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/ all][all]}] [mlt <64-6399>]] [mlt id <1-128>]]</code>	Shows all LACP port configuration information for all interfaces or the interface you specify. <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> is the port or port list. <code><64-6399></code> is the interface index of the mlt. <code><1-128></code> is the MLT ID.

Variable	Value
<code>partner-admin</code> <code>interface[gigabitethernet] [vid</code> <code><1-4059> [{slot/port[/sub-port]</code> <code>[-slot/port[/sub-port]][,...][slot/</code> <code>all][all]]}</code>	<p>Shows LACP partner administrative information for all interfaces or the specified interface.</p> <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> is the port or port list.
<code>partner-oper interface</code> <code>[gigabitethernet] [vid</code> <code><1-4059> [{slot/port[/sub-port]</code> <code>[-slot/port[/sub-port]][,...][slot/</code> <code>all][all]]}</code>	<p>Shows LACP partner operational information for all interfaces or the specified interface.</p> <ul style="list-style-type: none"> Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code> is the port or port list.

Display LACP Statistics for Specific Ports

Display individual LACP statistics for specific ports to manage network performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View statistics for specific ports:

```
show interfaces GigabitEthernet statistics lacp [{slot/port[/sub-port]
[-slot/port[/sub-port]][,...]]}
```

Example

View LACP statistics:

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet statistics lacp

=====
Port Stats Lacp
=====
PORT TX      RX      TX      RX      TX      RX      RX      RX
NUM  LACPDU LACPDU MARKERPDU MARKERPDU MARKERRESPDU MARKERRESPDU UNKNOWN ILLEGAL
-----
1/39  0        0        0        0        0        0        0        0
1/40  0        0        0        0        0        0        0        0
2/1   0        0        0        0        0        0        0        0
2/2   0        0        0        0        0        0        0        0
```

Variable Definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics lacp** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}</code>	<div>Identifies the slot and port in one of the following formats:</div> <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) <div>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</div>

LACP configuration using EDM

MultiLink Trunking (MLT) with Link Aggregation Control Protocol (LACP) manages switch ports and port memberships to form a link aggregation group (LAG). Configure LACP to allow dynamic bundling of physical ports to form a single logical channel.



Note
The switch does not support standby ports for LACP aggregation groups.

Configure Global LACP Parameters

Use LACP parameters to manage switch ports and their port memberships to form link aggregation groups (LAG). Link Aggregation Control Protocol (LACP) can dynamically add or remove LAG ports, depending on their availability and states.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **MLT/LACP**.
3. Select the **LACP Global** tab.
4. To enable LACP globally, select the **Enable** check box.
5. Configure the remaining parameters as required.

**Important**

Configuration changes to the LACP timers are not effective immediately. Link Aggregation Control Protocol timers are not reset until the next time LACP is restarted globally or on a port. This ensures consistency with peer switches.

6. Select **Apply**.

LACP Global *Field Descriptions*

Use the data in the following table to use the **LACP Global** tab.

Name	Description
Enable	Enables or disables LACP globally.
SystemPriority	Configures the system priority for all LACP enabled aggregators and ports. The default value is 32768.
FastPeriodicTime	Configures the number of milliseconds between periodic transmissions that use short timeouts. Sets this value to all LACP-enabled ports. The range is 200–20000. The default value is 1000.
FastPeriodicTimeOper	Displays the operating value of the fast periodic timer on the port. The default value is 1000.
SlowPeriodicTime	Configures the number of milliseconds between periodic transmissions that use long timeouts. All LACP enabled ports get the same value from this setting. The range is 10000–30000. The default value is 30000.
SlowPeriodicTimeOper	Displays the operating value of the slow periodic timer on the port. The default value is 30000.
AggrWaitTime	Configures the number of milliseconds to delay aggregation to allow multiple links to aggregate simultaneously. The range is 200–2000. The default value is 2000.

Name	Description
AggrWaitTimeOper	Displays the operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Configures the value used to calculate timeout time from the periodic time. All LACP-enabled ports get the same value from this setting. The range is 2–10. The default value is 3.
TimeoutScaleOper	Displays the operating value of the timeout scale on the port. The default value is 3.
SysId	Specifies the LACP system ID. The default value is f8:73:a2:00:90:00.
SmltSysId Note: Exception: not supported on 4220 Series or 5320 Series.	Specifies the LACP system ID for Split MultiLink Trunking (SMLT).

Configure LACP Parameters

Configure LACP parameters to manage LACP information.

About This Task



Important

The switch does not support standby ports for LACP aggregation groups.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **MLT/LACP**.
3. Select the **LACP** tab.
4. Double-click a field to change the value.
5. Select **Apply**.

LACP Field Descriptions

Use the data in the following table to use the **LACP** tab.

Name	Description
Index	The unique identifier the local system allocates to this aggregator. This attribute identifies an aggregator instance among the subordinate managed objects of the containing object. This value is read-only.
MACAddress	The six octet read-only value carrying the individual MAC address assigned to the aggregator.
ActorSystemPriority	The two octet read-write value indicating the priority value associated with the actor system ID. The default value is 32768.
ActorSystemID	The six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator. From the perspective of the link aggregation mechanisms, only a single combination of actor system ID and system priority are considered. No distinction is made between the values of these parameters for an aggregator and the ports that are associated with it. The protocol is described in terms of the operation of aggregation within a single system. However, the managed objects provided for the both the aggregator and the port allow management of these parameters. The result permits a single piece of equipment to be configured by management to contain more than one system from the point of view of the operation of link aggregation, which is useful in the configuration of equipment that has limited aggregation capability.
AggregateOrIndividual	Indicates whether the aggregator represents an aggregate (true) or an individual link (false).
ActorAdminKey	Specifies the current administrative value of the key for the aggregator, which is a 16-bit read-write value. The administrative key value can differ from the operational key value. This key needs to match the LAG key. The default value is 0.
ActorOperKey	Displays the current read-only operational value of the key for the aggregator. The operational key value can differ from the administrative key value. The meaning of particular key values is of local significance.
PartnerSystemID	The six octet read-only MAC address value consisting of the unique identifier for the current protocol partner of this aggregator. A value of zero indicates that there is no known partner. If the aggregation is manually configured, the value is assigned by the local system.

Name	Description
PartnerSystemPriority	The two octet read-only value that indicates the priority value associated with the partner system ID. If the aggregation is manually configured, this system priority value is a value assigned by the local system.
PartnerOperKey	The current operational value of the key for the aggregator current protocol partner, which is a 16-bit read-only value. If the aggregation is manually configured, the value is assigned by the local system.

Configure LACP on a Port



Important

Changes made at the global level override and reset the port-level configuration.

Configure LACP on a port to enable LACP.

You must use the LACP keys to determine which ports are eligible for link aggregation. The LACP keys are defined by the ports after you configure the multilink trunk. You can aggregate the ports key that match the MLT key into that multilink trunk.

Before You Begin

When enabling or disabling LACP on a port, as a best practice, disable the port first and re-enable the port after the configuration is complete.

About This Task

The minimum LACP configuration is as follows:

- Assign a given key to a set of ports. In the following procedure steps, you must assign the key before you enable LACP on the port.
- Assign the same key to an MLT with no members. The ports will automatically become MLT members.

Keys do not need to match between two LACP peers.

A port can operate in active or passive mode. You can configure a port to be an individual link or an aggregated link.



Note

When using LACP with private VLANs, configure the private VLAN at the interface level.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.

4. Select the **LACP** tab.
5. To assign an LACP key, configure **ActorAdminKey**.
6. Select **Apply**.
7. To enable LACP on the port, select **AdminEnable**.
8. Configure the remaining parameters as required.
9. Select **Apply**.

LACP Field Descriptions

Use the data in the following table to use the **LACP** tab.

All parameters beginning with **Partner** are for debug purposes only. If you configure these commands locally and there is a mismatch with what is learned from the partner, there will be trace or log messages.

Name	Description
AdminEnable	Enables LACP status for the port. The default value is false.
OperEnable	Displays the operational status of LACP for the port. The default value is false.
FastPeriodicTime	Specifies the number of milliseconds between periodic transmissions using short timeouts for all LACP enabled ports. The default value is 1000.
FastPeriodicTimeOper	Displays the operating value of the fast periodic timer on the port. The default value is 1000.
SlowPeriodicTime	Specifies the number of milliseconds between periodic transmissions using long timeouts for all LACP enabled ports. The default value is 30000.
SlowPeriodicTimeOper	Displays the operating value of the slow periodic timer on the port. The default value is 30000.
AggrWaitTime	Specifies the number of milliseconds to delay aggregation so multiple links can aggregate simultaneously. The default value is 2000.
AggrWaitTimeOper	Displays the operating value of the aggregate wait timer on the port. The default value is 2000.
TimeoutScale	Assigns the value used to calculate timeout time from the periodic time. Configure this value for all LACP-enabled ports. The default value is 3.

Name	Description
TimeoutScaleOper	Displays the operating value of the timeout scale on the port. The default value is 3.
ActorSystemPriority	Specifies the two octet read-write value indicating the priority value associated with the actor system ID. The default value is 32768.
ActorSystemID	Displays the six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator. From the perspective of the link aggregation mechanisms, only a single combination of actor system ID and system priority are considered, and no distinction is made between the values of these parameters for an aggregator and the ports that are associated with it; that is, the protocol is described in terms of the operation of aggregation within a single system. However, the managed objects provided for the aggregator and the port both enable management of these parameters. The result of this is to permit a single piece of equipment to be configured by management to contain more than one system from the point of view of the operation of link aggregation. This can be of particular use in the configuration of equipment that has limited aggregation capability.
ActorAdminKey	Configures the aggregation key for this port. The administrative key value can differ from the operational key value. This key must match the LAG key. To return a configured key to the default value, you must first disable LACP aggregation on the port.
ActorOperKey	Displays the current operational value of the key for the aggregator. The operational key value can differ from the administrative key value. The meaning of particular key values is of local significance.
SelectedAggID	Displays the identifier value of the aggregator that this aggregation port has currently selected. Zero indicates that the aggregation port has not selected an aggregator, either because it is in the process of detaching from an aggregator or because there is no suitable aggregator available for it to select.

Name	Description
AttachedAggID	Displays the identifier value of the aggregator to which this aggregation port is currently attached. Zero indicates that the aggregation port is not currently attached to an aggregator.
ActorPort	Displays the port number locally assigned to the aggregation port. The port number is communicated in LACPDUs as the Actor_Port.
ActorPortPriority	Specifies the priority value assigned to this aggregation port. The default value is 32768.
ActorAdminState	Configures one or more of the following: <ul style="list-style-type: none"> • lacpActive - Configures the LACP mode to be active. • lacpShortTimeout - Configures the timeout to short. • aggregation - Enables aggregation on the port, which makes it an aggregated link. By default, only lacpActive is enabled. These values enable administrative control over the values of LACP_Activity, LACP_Timeout, and aggregation.
ActorOperState	Displays a string of eight bits, corresponding to the current operational values of Actor_State as transmitted by the actor in LACPDUs. The default value is lacpActive.
PartnerAdminSystemPriority	Specifies the current administrative value of the port number for the protocol partner. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperSystemPriority	Displays a two octet value indicating the operational value of priority associated with the partner system ID. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemPriority if there is no protocol partner. The default value is 0.

Name	Description
PartnerAdminSystemID	Specifies a six octet MAC address value that represents the administrative value of the aggregation port protocol partners system ID. The assigned value is used with the values of PartnerAdminSystemPriority, PartnerAdminKey, PartnerAdminPort, and PartnerAdminPortPriority to achieve manually configured aggregation. The default value is 00:00:00:00:00:00.
PartnerOperSystemID	Displays a six octet MAC address value that indicates representing the current value of the aggregation port protocol partner system ID. A value of zero indicates that there is no known protocol partner. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemID if there is no protocol partner. The default value is 00:00:00:00:00:00.
PartnerAdminKey	Specifies the current administrative value of the key for the protocol partner. The assigned value is used with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminPort, and PartnerAdminPortPriority to achieve manually configured aggregation. The default value is 0.
PartnerOperKey	Displays the current operational value of the key for the aggregator current protocol partner. If the aggregation is manually configured, this value is assigned by the local system. The default value is 0.
PartnerAdminPort	Specifies the current administrative value of the port number for the protocol partner. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperPort	Displays the operational port number assigned to this aggregation port by the aggregation port protocol partner. The value of this attribute can contain the manually configured value carried in AggPortPartnerAdminPort if there is no protocol partner. The default value is 0.

Name	Description
PartnerAdminPortPriority	Specifies the current administrative value of the port priority for the protocol partner. The assigned value is used with the values of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPort to achieve manually configured aggregation. The default value is 0.
PartnerOperPortPriority	Displays the priority value assigned to this aggregation port by the partner. The value of this attribute can contain the manually configured value carried in PartnerAdminPortPriority if there is no protocol partner. The default value is 0.
PartnerAdminState	Specifies a string of eight bits, corresponding to the current administrative value of Actor_State for the protocol partner, by selecting check boxes. The assigned value is used to achieve manually configured aggregation. The default value is none.
PartnerOperState	Displays a string of eight bits, corresponding to the current values of Actor_State in the most recently received LACPDU transmitted by the protocol partner. In the absence of an active protocol partner, this value can reflect the manually configured value PartnerAdminState. The default value is none.

Configure LACP on an Extreme Integrated Application Hosting Port



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to configure Link Aggregation Control Protocol (LACP) on an Extreme Integrated Application Hosting (IAH) port. You must use the LACP keys to determine which IAH ports are eligible for link aggregation. The LACP keys are defined by the IAH ports after you configure the multilink trunk (MLT). You can aggregate the IAH port key that matches the MLT key into that multilink trunk.

The minimum LACP configuration is as follows:

- Assign a given key to a set of IAH ports.

- Assign the same key to an MLT with no members. The IAH ports automatically become MLT members.

The keys do not have to match between two LACP peers.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **LACP** tab.
4. Select **AdminEnable**.
5. In the **ActorAdminKey** field, enter a value.
6. Configure the other parameters as required.
7. Select **Apply**.

LACP Field Descriptions

Use data in the following table to use the **LACP** tab.



Note

The **Partner** fields in the **LACP** tab are for debug purpose only. If you configure them locally and there is a mismatch with what is learned from the partner, trace or log messages are generated.

Name	Description
AdminEnable	Enables LACP status for the Extreme Integrated Application Hosting port. The default value is false.
OperEnable	Specifies the LACP operational status for the Extreme Integrated Application Hosting port. The default value is false.
FastPeriodicTime	Specifies the number of milliseconds between periodic transmissions using short timeouts for all LACP enabled Extreme Integrated Application Hosting ports. The default value is 1000.
FastPeriodicTimeOper	Specifies the operating value of the fast periodic timer on the Extreme Integrated Application Hosting port. The default value is 1000.
SlowPeriodicTime	Specifies the number of milliseconds between periodic transmissions using long timeouts for all LACP enabled Extreme Integrated Application Hosting ports. The default value is 30000.

Name	Description
SlowPeriodicTimeOper	Specifies the operating value of the slow periodic timer on the Extreme Integrated Application Hosting port. The default value is 30000.
AggrWaitTime	Specifies the number of milliseconds to delay aggregation to enable multiple links to aggregate simultaneously. The default value is 2000.
AggrWaitTimeOper	Specifies the operating value of the aggregate wait timer on the Extreme Integrated Application Hosting port. The default value is 2000.
TimeoutScale	Specifies the value used to calculate timeout duration from the periodic time. Set the same value for all LACP enabled Extreme Integrated Application Hosting ports. The default value is 3.
TimeoutScaleOper	Specifies the operating value of the timeout scale on the Extreme Integrated Application Hosting port. The default value is 3.
ActorSystemPriority	Specifies the two octet read-write value indicating the priority value associated with the actor system ID. The default value is 32768.

Name	Description
ActorSystemID	<p>Specifies the six octet read-write MAC address value used as a unique identifier for the system that contains this aggregator.</p> <p>From the perspective of the link aggregation mechanisms, only a single combination of actor system ID and system priority are considered, and no distinction is made between the values of these parameters for an aggregator and the Extreme Integrated Application Hosting ports that are associated with it; that is, the protocol is described in terms of the operation of aggregation within a single system. However, the managed objects provided for the aggregator and the Extreme Integrated Application Hosting port both allow management of these parameters. The result of this is to permit a single piece of equipment to be configured by management to contain more than one system from the point of view of the operation of link aggregation. This can be of particular use in the configuration of equipment that has limited aggregation capability.</p>
ActorAdminKey	<p>Specifies the current read-write administrative value of the key for the aggregator, which is a 16-bit value. The administrative key value can differ from the operational key value. This key must match the LAG key.</p>
ActorOperKey	<p>Specifies the current read-only operational value of the key for the aggregator. The operational key value can differ from the administrative key value. The meaning of particular key values is of local significance.</p>
SelectedAggID	<p>Specifies the identifier value of the aggregator that this aggregation port has currently selected. Zero indicates that the aggregation port has not selected an aggregator, either because it is in the process of detaching from an aggregator or because there is no suitable aggregator available for it to select. This value is read-only.</p>
AttachedAggID	<p>Specifies the identifier value of the aggregator to which this aggregation port is currently attached. Zero indicates that the aggregation port is not currently attached to an aggregator. This value is read-only.</p>

Name	Description
ActorPort	Specifies the port number locally assigned to the aggregation port. The port number is communicated in LACPDUs as the Actor_Port. This value is read-only.
ActorPortPriority	Specifies the priority value assigned to this aggregation port. This 16-bit value is read-write. The default value is 32768.
ActorAdminState	Specifies a string of eight bits, corresponding to the administrative values as transmitted by the actor in LACPDUs. The values are: <ul style="list-style-type: none"> the first bit corresponds to bit 0 of Actor_State (LACP_Activity) (the default value) the second bit corresponds to bit 1 (LACP_Timeout) the third bit corresponds to bit 2 (Aggregation) the fourth bit corresponds to bit 3 (Synchronization) the fifth bit corresponds to bit 4 (Collecting) the sixth bit corresponds to bit 5 (Distributing) the seventh bit corresponds to bit 6 (Defaulted) the eighth bit corresponds to bit 7 (Expired) These values enable administrative control over the values of LACP_Activity, LACP_Timeout, and aggregation. This attribute value is read-write.
ActorOperState	Specifies a string of eight bits, corresponding to the current operational values of Actor_State as transmitted by the actor in LACPDUs. This attribute value is read-only. The default value is lacpActive.
PartnerAdminSystemPriority	Specifies the current administrative value of the Extreme Integrated Application Hosting port number for the protocol partner. It is a 16-bit read-write value. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.

Name	Description
PartnerOperSystemPriority	Displays a two octet read-only value indicating the operational value of priority associated with the partner system ID. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemPriority if there is no protocol partner. The default value is 0.
PartnerAdminSystemID	Specifies a six octet read-write MAC address value that represents the administrative value of the aggregation port protocol partners system ID. The assigned value is used with the values of PartnerAdminSystemPriority, PartnerAdminKey, PartnerAdminPort, and PartnerAdminPortPriority to achieve manually configured aggregation. The default value is 00:00:00:00:00:00.
PartnerOperSystemID	Displays a six octet read-only MAC address value that indicates representing the current value of the aggregation port protocol partner system ID. A value of zero indicates that there is no known protocol partner. The value of this attribute can contain the manually configured value carried in PartnerAdminSystemID if there is no protocol partner. The default value is 00:00:00:00:00:00.
PartnerAdminKey	Specifies the current administrative value of the key for the protocol partner. It is a 16-bit read-write value. The assigned value is used with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminPort, and PartnerAdminPortPriority to achieve manually configured aggregation. The default value is 0.
PartnerOperKey	Specifies the current operational value of the key for the aggregator current protocol partner. It is a 16-bit read-only value. If the aggregation is manually configured, this value is assigned by the local system. The default value is 0.

Name	Description
PartnerAdminPort	Specifies the current administrative value of the port number for the protocol partner. It is a 16-bit read-write value. The assigned value is used, along with the value of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPortPriority, to achieve manually configured aggregation. The default value is 0.
PartnerOperPort	Specifies the operational Extreme Integrated Application Hosting port number assigned to this aggregation port by the aggregation port protocol partner. The value of this attribute can contain the manually configured value carried in AggPortPartnerAdminPort if there is no protocol partner. This 16-bit value is read-only. The default value is 0.
PartnerAdminPortPriority	Specifies the current administrative value of the port priority for the protocol partner. It is a 16-bit read-write value. The assigned value is used with the values of PartnerAdminSystemPriority, PartnerAdminSystemID, PartnerAdminKey, and PartnerAdminPort to achieve manually configured aggregation. The default value is 0.
PartnerOperPortPriority	Specifies the priority value assigned to this aggregation port by the partner. The value of this attribute can contain the manually configured value carried in PartnerAdminPortPriority if there is no protocol partner. This 16 bit value is read-only. The default value is 0.

Name	Description
PartnerAdminState	Specifies a string of eight bits, corresponding to the current administrative value of Actor_State for the protocol partner, by selecting check boxes. This attribute value is read-write. The assigned value is used to achieve manually configured aggregation. The default value is none.
PartnerOperState	Specifies a string of eight bits, corresponding to the current values of Actor_State in the most recently received LACPDU transmitted by the protocol partner. In the absence of an active protocol partner, this value can reflect the manually configured value PartnerAdminState. This attribute value is read-only. The default value is none.

View LACP Port Statistics

View LACP port statistics to monitor the performance of the port.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **LACP** tab.
5. To change the poll interval, in the toolbar select the **Poll Interval** box, and then select a new interval.

LACP Field Descriptions

Use the data in the following table to use the **LACP** tab.

Name	Description
LACPDUsRx	The number of valid LACPDU received on this aggregation port.
MarkerPDUsRx	The number of valid marker PDUs received on this aggregation port.
MarkerResponsePDUsRx	The number of valid marker response PDUs received on this aggregation port.

Name	Description
UnknownRx	The number of frames received that either: <ul style="list-style-type: none">• carry Slow Protocols Ethernet type values, but contain an unknown PDU.• are addressed to the Slow Protocols group MAC Address, but do not carry the Slow Protocols Ethernet Type.
IllegalRx	The number of frames received that carry the Slow Protocols Ethernet Type value (43B.4), but contain a badly formed PDU or an illegal value of Protocol Subtype (43B.4).
LACPDUsTx	The number of LACPDU's transmitted on this aggregation port.
MarkerPDUsTx	The number of marker PDUs transmitted on this aggregation port.
MarkerResponsePDUsTx	The number of marker response PDUs transmitted on this aggregation port.



Link Layer Discovery Protocol

[Link Layer Discovery Protocol \(802.1AB\) Fundamentals](#) on page 2466
[Link Layer Discovery Protocol-Media Endpoint Discovery](#) on page 2470
[Link Layer Discovery Protocol configuration using CLI](#) on page 2471
[LLDP-MED Configuration Using CLI](#) on page 2487
[Link Layer Discovery Protocol configuration using EDM](#) on page 2494
[LLDP-MED Configuration Using EDM](#) on page 2506

Table 183: Link Layer Discovery Protocol product support

Feature	Product	Release introduced
Industry Standard Discovery Protocol (ISDP) (CDP compatible)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Link Layer Discovery Protocol (LLDP)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 183: Link Layer Discovery Protocol product support (continued)

Feature	Product	Release introduced
Link Layer Discovery Protocol-Media Endpoint Discovery	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The following sections describe how to use Link Layer Discovery Protocol (LLDP) and Industry Standard Discovery Protocol (ISDP).

Link Layer Discovery Protocol (802.1AB) Fundamentals

With Link Layer Discovery Protocol (LLDP) you can obtain node and topology information to help detect and correct network and configuration errors.

LLDP

802.1AB is the IEEE standard called Station and Media Access Control Connectivity Discovery. This standard defines the Link Layer Discovery Protocol.

LLDP stations connected to a local area network (LAN) can advertise station capabilities to each other, allowing the discovery of physical topology information for network management.

LLDP-compatible stations can comprise any interconnection device, including PCs, IP Phones, switches, and routers.

Each LLDP station stores LLDP information in a standard Management Information Base (MIB), making it possible for a network management system (NMS) or application to access the information.

The functions of an LLDP station include:

- Advertising connectivity and management information about the local station to adjacent stations
- Receiving network management information from adjacent stations
- Enabling the discovery of certain configuration inconsistencies or malfunctions that can result in impaired communications at higher layers

For example, you can use LLDP to discover duplex mismatches between an IP Phone and the connected switch.

LLDP is compatible with IETF PROTO MIB (IETF RFC 2922).

The following figure shows an example of a LAN using LLDP.




Figure 177: LLDP in a LAN

Legend:

1. The switch and an LLDP-enabled router advertise chassis and port IDs and system descriptions to each other
2. The devices store the information about each other in local MIB databases, accessible with SNMP
3. A network management system retrieves the data stored by each device and builds a network topology map
4. Switch
5. Router
6. Management work station
7. IP Phone

LLDP Modes

LLDP is a one-way protocol.

An LLDP agent can transmit information about the capabilities and current status of the system associated with its MAC service access point (MSAP) identifier.

The LLDP agent also can receive information about the capabilities and current status of the system associated with a remote MSAP identifier.

However, LLDP agents cannot solicit information from each other.

You can configure the local LLDP agent to transmit and receive.

Connectivity and Management Information

The information parameters in each LLDP frame are in a LLDP Protocol Data Unit (PDU) as a sequence of short, variable length information elements known as TLVs (type, length, value).

Each LLDP PDU includes the following mandatory TLVs:

- Chassis ID
- Port ID
- Time To Live
- Port Description
- System Name
- System Description
- System Capabilities (indicates both the system supported capabilities and enabled capabilities, such as end station, bridge, or router)
- Management Address

The chassis ID and the port ID values are concatenated to form a logical MSAP identifier that the recipient uses to identify the sending LLDP agent and port.

A non-zero value in the Time to Live (TTL) field of the TTL TLV indicates to the receiving LLDP agent how long the LLDP PDU information from the MSAP identifier remains valid.

The receiving LLDP agent automatically discards all LLDP PDU information, if the sender fails to update it in a timely manner.

A zero value in TTL field of Time To Live TLV tells the receiving LLDP agent to discard the information associated with the LLDP PDU MSAP identifier.

LLDP for the Segmented Management Instance

LLDP and SONMP both advertise the same topology IP address for the Segmented Management Instance management interface. LLDP supports IPv4 and IPv6 advertisement.

If all three management interfaces are configured, the advertised default topology IP priority is management CLIP, then management VLAN, then management OOB. You can change the default topology IP using CLI or EDM. If multiple IPv4 addresses are

configured on an OOB or VLAN management interface, the advertised IP priority is static IP address, then DHCP IP address, then link-local IP address.

Transmitting LLDP PDUs

When a transmit cycle is initiated, the LLDP manager extracts the managed objects from the LLDP local system MIB and formats this information into TLVs. TLVs are inserted into the LLDP PDU.

LLDP PDUs are regularly transmitted at a user-configurable transmit interval (tx-interval) or when any of the variables in the LLDPDU is modified on the local system; for example, system name or management address.

Transmission delay (tx-delay) is the minimum delay between successive LLDP frame transmissions.

TLV system MIBs

The LLDP local system MIB stores the information to construct the various TLVs for transmission.

The LLDP remote systems MIB stores the information received from remote LLDP agents.

LLDP PDU and TLV error handling

The system discards LLDP PDUs and TLVs that contain detectable errors.

The system assumes that TLVs that contain no basic format errors, but that it does not recognize, are valid and stores them for retrieval by network management.

LLDP and MultiLink Trunking

You must apply TLVs on a per-port basis.

Because LLDP manages trunked ports individually, TLVs configured on one port in a trunk do not propagate automatically to other ports in the trunk.

And the system sends advertisements to each port in a trunk, not on a per-trunk basis.

LLDP and Fabric Attach

Fabric Attach uses LLDP to signal a desire to join the SPB network. When a switch is enabled as an FA Server, it receives IEEE 802.1AB LLDP messages from FA Client and FA Proxy devices requesting the creation of Switched UNI service identifiers (I-SIDs). All of the discovery handshakes and I-SID mapping requests are using LLDP TLV fields. Based on the LLDP standard, FA information is transmitted using organizational TLVs within LLDP PDUs.

FA also leverages LLDP to discover directly connected FA peers and to exchange information associated with FA between those peers.

Link Layer Discovery Protocol-Media Endpoint Discovery

Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) defined in ANSI/TIA-1057, is an extension to the LLDP standard protocol as defined in IEEE 802.1AB. LLDP-MED provides support to deploy Voice over Internet Protocol (VoIP) telephones into the LAN environment. LLDP-MED provides additional TLVs for basic configuration, network policy configuration, location identification, and inventory management.

Following are the types of LLDP-MED devices:

- Network connectivity devices: provide access to the IEEE 802 based LAN infrastructure for LLDP-MED endpoint devices. The LLDP-MED Network Connectivity device is a LAN access device based on:
 - LAN Switch or Router
 - IEEE 802.1 Bridge
 - IEEE 802.3 Repeater
 - IEEE 802.11 Wireless Access Point
 - Any device that supports the IEEE 802.1AB, LLDP-MED, and can relay IEEE 802 frames.
- Endpoint devices: located at the IEEE 802 LAN network edge, participating in the IP communication service using the LLDP-MED framework. The endpoint devices are divided into three classes:
 - Class 1 - LLDP-MED Generic Endpoint devices, for example, IP communication controllers.
 - Class 2 - LLDP-MED Media Endpoint devices, for example, media servers, conference bridges.
 - Class 3 - LLDP-MED Communication Endpoint devices, for example, IP telephones.

Organizational-specific TLVs for LLDP-MED

The organizational-specific TLVs for use by LLDP-MED network connectivity and endpoint devices are:

- Capabilities TLV — enables a network element to determine whether particular connected devices support LLDP-MED, and also discover the TLVs supported by specific network connectivity or endpoint devices.
- Network Policy Discovery TLV — enables both network connectivity and endpoint devices to advertise VLAN information, Layer 2, and Layer 3 priorities associated with a specific set of applications on a port. In addition, an LLDP-MED endpoint advertises this TLV for supported application types to enable the discovery of specific policy information and the diagnosis of network policy configuration mismatch issues.
- Location Identification Discovery TLV — allows network connectivity devices to advertise the appropriate location information for communication endpoint devices,

including emergency call service location, to use in the context of location-based applications.

- Extended Power-via-MDI Discovery TLV — enables advanced power management between an LLDP-MED network connectivity and endpoint devices. The Extended Power-via-MDI TLV enables the advertisement of fine grained power requirement details, endpoint power priority, and power status for network connectivity and endpoint devices.
- Inventory Management Discovery TLV — enables tracking and identification of inventory-related attributes for endpoint devices. For example, manufacturer, model name, and software version.

Link Layer Discovery Protocol configuration using CLI

This section describes how to configure Link Layer Discovery Protocol using the Command Line Interface (CLI).

IPv4 management IP addresses are supported by LLDP, including the management virtual IP address, and they are advertised in the Management address TLV.

Configure Global LLDP Transmission Parameters

Before You Begin

- In the GigabitEthernet Interface Configuration mode, specify the LLDP port status as transmit only or transmit and receive.

About This Task

Use this procedure to configure global LLDP transmission parameters on the switch. If required, you can also restore these parameters to their default values.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the LLDP transmission interval:
`lldp tx-interval <5-32768>`
3. Configure the LLDP transmission parameters:
`lldp tx-hold-multiplier <2-10>`

Example

Configure the LLDP transmission interval.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#lldp tx-interval 31
```

Variable Definitions

The following table defines parameters for the **lldp** command.

Variable	Value
<code>tx-interval<5-32768></code>	Specifies the global LLDP transmit interval in seconds, that is, the interval in which LLDP frames are transmitted. The default is 30 seconds.
<code>tx-hold-multiplier<2-10></code>	Configures the multiplier for the transmit interval used to compute the Time To Live (TTL) value in LLDP frames. The default is 4 seconds.

Configure LLDP Status on Ports

About This Task

Use this procedure to configure LLDP and configure the status to transmit and receive on a port, or ports, on your switch.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the status for transmit and receive on a port or ports:

```
lldp port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all]  
[all]} status <txAndRx>
```

Variable Definitions

The following table defines parameters for the **lldp port** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>status <txAndRx></code>	Enables LLDP PDU transmit and receive on the port(s).

Enable CDP Mode on a Port

To configure the switch as CDP-compatible, you must enable the Industry Standard Discovery Protocol (ISDP) on a port, or ports, on the switch.

About This Task

By default, if you enable CDP, the interface accepts only CDP packets. Similarly, if you disable CDP but enable LLDP, the interface accepts only LLDP packets. You can optionally configure the interface to send and receive both LLDP and CDP packets.

If you plan to configure a CDP-enabled port for features that rely on LLDP like Auto-sense or Fabric Attach (FA), ensure you enable both CDP and LLDP packet transmission on the port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all] [all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable CDP on the port:

```
lldp cdp enable
```

3. (Optional) Configure the port to send and receive both CDP and LLDP packets:

```
lldp cdp dual-mode-enable
```

Disable LLDP Tx-TLVs

About This Task

Use this procedure to disable optional management Link Layer Discovery Protocol (LLDP) TLVs for transmission.

LLDP TLVs for transmission are enabled by default.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable LLDP TLVs for transmission:

```
no lldp tx-tlv {local-mgmt-addr | port-desc | sys-cap | sys-desc | sys-  
name}
```

3. Verify the configuration:

```
show lldp tx-tlv port {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

Variable Definitions

The following table defines parameters for the **lldp tx-tlv** command.

Variable	Value
<i>local-mgmt-addr</i>	Specifies the local management address TLV. The default is enabled.
<i>port-desc</i>	Specifies the port description transmission TLV. The default is enabled.
<i>sys-cap</i>	Specifies the system capabilities transmission TLV. The default is enabled.

Variable	Value
<i>sys-desc</i>	Specifies the system description transmission TLV. The default is enabled.
<i>sys-name</i>	Specifies the system name transmission TLV. The default is enabled.

Enable Dot3 LLDP TLVs

About This Task

Use this procedure to enable transmission of dot3 MAC/PHY configuration/status Link Layer Discovery Protocol (LLDP) TLV.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable LLDP TLVs for transmission on a specific port:

```
lldp tx-tlv dot3 mac-phy-config-status
```

3. Verify the configuration:

```
show lldp tx-tlv dot3
```

Example

Enable LLDP TLVs for transmission on a specific port:

```
Switch:1>enable
Switch:1#config
Configuring from terminal or network [terminal]? t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/2
Switch:1(config-if)#lldp tx-tlv dot3 mac-phy-config-status
Switch:1#show lldp tx-tlv dot3
=====
                        LLDP port dot3 tlvs
=====
-----
Port      MAC Physical
          Config Status
-----
1/1       disabled
```

1/2	enabled
1/3	enabled
1/4	disabled
1/5	disabled

Variable Definitions

The following table defines parameters for the **lldp tx-tlv dot3** command.

Variable	Value
<i>mac-phy-config-status</i>	Specifies the status of the MAC Physical Config Status TLV. The default is disabled.

Disable LLDP MED TLVs for Transmission on Ports

About This Task

Use this procedure to configure organizational-specific TLVs for Link Layer Discovery Protocol (LLDP) Media Endpoint Devices (MED). LLDP MED TLVs are enabled by default.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable LLDP MED TLVs for transmission:

```
no lldp tx-tlv med {extendedPSE | inventory | location | med-
capabilities | network-policy}
```

3. Verify the configuration:

```
show lldp tx-tlv med
```

Example

Disable LLDP MED TLVs

```
Switch:1>enable
Switch:1#config
Configuring from terminal or network [terminal]? t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/2
Switch:1(config-if)#no lldp tx-tlv med inventory
Switch:1(config-if)#no lldp tx-tlv med location
Switch:1#show lldp tx-tlv med
```

LLDP port med tlvs					
Port	Med Capabilities	Network Policy	Location	ExtendedPower ViaMDI	Inventory
1/2	enabled	enabled	disabled	enabled	disabled

Variable Definitions

The following table defines parameters for the **lldp tx-tlv med** command.

Variable	Value
<i>extendedPSE</i>	Specifies the extended power-via-MDI Media Endpoint Device (MED) transmission TLV. This TLV is applicable for POE-capable switches only. The default is enabled on POE-capable switches only.
<i>inventory</i>	Specifies the system capabilities MED transmission TLV. The default is enabled.
<i>location</i>	Specifies the system description MED transmission TLV. The default is enabled.
<i>med-capabilities</i>	Specifies the system name MED transmission TLV. The default is enabled.
<i>network-policy</i>	Specifies the network policy MED transmission TLV. The default is enabled.

View Global LLDP Information

About This Task

Use this procedure to view global LLDP information, to know which LLDP settings and parameters are configured.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display LLDP local system data:
`show lldp local-sys-data [med]`
3. Display detailed LLDP neighbor information:
`show lldp neighbor detail [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]`

4. Display the summary of LLDP neighbors:

```
show lldp neighbor summary [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

5. Display the list of ports:

```
show lldp port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

6. Display the LLDP reception statistics:

```
show lldp rx-stats [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

7. Display the LLDP statistics:

```
show lldp stats
```

8. Display the LLDP transmission statistics:

```
show lldp tx-stats [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

Examples

View global LLDP information:

```
Switch:1#show lldp
802.1ab Configuration:
-----
          TxInterval: 30
      TxHoldMultiplier: 4
          ReinitDelay: 1
              TxDelay: 1
      NotificationInterval: 5
```

View the LLDP local system data on the switch:

```
Switch:1#show lldp local-sys-data
=====
                        LLDP Local System Data
=====

          ChassisId: MAC Address          b0:ad:aa:4c:54:00
          SysName   : LLDP agent
          SysDescr  : 5320-24P-8XE-FabricEngine (9.0.0.0)
          SysCap    : Br / Br
-----
Capabilities Legend: (Supported/Enabled)
B= Bridge,      D= DOCSIS,      O= Other,      R= Repeater,
S= Station,    T= Telephone, W= WLAN,      r= Router
```

View the LLDP neighbor information. You can also view this information for a specific port.

```
Switch:1#show lldp neighbor
=====
                        LLDP Neighbor
=====

Port: 1/28      Index    : 1              Time: 0 day(s), 01:16:25
                Protocol  : LLDP
                ChassisId: MAC Address      a4:25:1b:52:54:00
                PortId   : MAC Address      a4:25:1b:52:54:1b
                SysName  : BEB
```

```

SysCap    : Br / Br
PortDescr: Extreme Networks 7520-48XT-6C-FabricEngine - 10GBaseT Port 1/13
SysDescr  : 7520-48XT-6C-FabricEngine (9.0.0.0)
Address   : 192.0.2.47

```

Total Neighbors : 1

Capabilities Legend: (Supported/Enabled)

B= Bridge, D= DOCSIS, O= Other, R= Repeater,
S= Station, T= Telephone, W= WLAN, r= Router



Note

For information on the Mgmt-sdwan interface, see [Fabric Extend \(FE\) States](#) on page 32.

```
Switch:1#show lldp neighbor
```

LLDP Neighbor

```

Port: 1/20      Index    : 3
                  Protocol : LLDP
                  ChassisId: MAC Address      00:07:32:a6:8e:d4
                  PortId   : MAC Address      00:07:32:a6:8e:d4
                  SysName  : SDWAN-4
                  SysCap   : BWrS / Br
                  PortDescr: lan1
                  SysDescr : ExtremeCloud SDWAN S/N:A1122W1412A0
                  Address  : 172.16.0.1
                  IPv6 Address : 0:0:0:0:0:0:0:0

```

Total Neighbors : 1

View the LLDP neighbor details of all ports on the switch.

```
Switch1:1#show lldp neighbor detail
```

LLDP Detail

LLDP neighbor with index 3 discovered on port 1/1:

Chassis:

Index: 3

Protocol: LLDP

ID Subtype: MAC Address(4)

ID: 00:1b:4f:f9:cc:00

Name: Mgmt_M7_10.3.44.250

Description: Ethernet Routing Switch 5952GTS

HW:R0D.7

FW:7.2.0.5

SW:v7.9.2.017

Available Capabilities:

Bridge(4)

Router(16)

Enabled Capabilities:

Bridge(4)

Router(16)

Management addresses:

Port index: 1

Protocol: LLDP

Description: Port 41

MFS: 0

DOT3:

```

Autonegotiation: enabled
Advertised Autonegotiation: 27713
mau type: 30
Aggregation ID: 0
Dot3_at: not present
DOT1:
  p-vid: 0
  Vlans:
Remote Call Servers: not present
Remote File Servers: not present

```

```
-----
Total LLDP detailed entries displayed: 3
```

View LLDP neighbor details on a specific port on the switch.

```
Switch1:1#show lldp neighbor detail port 1/19
```

```
=====
                        LLDP Detail
=====
```

```
LLDP neighbor with index 1 discovered on port 1/19:
```

```

Chassis:
  Index: 1
  Protocol: LLDP
  ID Subtype: MAC Address(4)
  ID: e4:db:ae:71:bc:00
  Name: BN_5_7520
  Description: 7520-48YE-8CE-FabricEngine (9.3.0.0)
  Available Capabilities:
    Bridge(4)
    Router(16)
  Enabled Capabilities:
    Bridge(4)
    Router(16)
  Management addresses:
    IPv4: 10.3.34.38

```

```

Port index: 1
Protocol: LLDP
Description: Extreme Networks 7520-48YE-8CE-FabricEngine - 10GbSR Port 1/1
MFS: 0
DOT3:
  Autonegotiation: not present
  Advertised Autonegotiation: 0
  mau type: 65535
  Aggregation ID: 0
Dot3_at: not present
DOT1:
  p-vid: 0
  Vlans:
Fabric Connect: expanded
  Primary bvlan: 4051
  Secondary bvlan: 4052
  System ID: 0049:0002:0115
  Expansion:
Remote Call Servers: not present
Remote File Servers: not present

```

```
-----
Total LLDP detailed entries displayed: 1
```

View the LLDP neighbor summary of all ports on the switch. You can also view this on a specific port.



Note

The EDM tab (**Edit > Diagnostic > LLDP & MED > LLDP > Neighbor >**) displays only IPv4 addresses. Use CLI to see both IPv4 and IPv6.

```
Switch:1#show lldp neighbor summary
```

LLDP Neighbor Summary						
LOCAL	IP/IPv6		CHASSIS	REMOTE		
PORT	PROT	ADDR	ID	PORT	SYSNAME	SYSDDESCR
1/1	LLDP	192.168.1.1	bc:ad:ab:00:1c:00	1/23	Mgmt_L10_10.3.44.248	Ethernet Routing Switch 59100GTS
1/2	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/2	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
1/3	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/3	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
1/4	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/4	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
1/5	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/5	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
1/6	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/6	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
1/7	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/7	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
1/8	LLDP	169.254.252.4	f4:ce:48:9e:fc:00	1/8	5320-24T-8XE-FabricEngi~	5320-24T-8XE-FabricEngine (9.0.2.0)
Total Neighbors : 8						

View the LLDP administrative status of ports 1/1 to 1/4.

```
Switch:1#show lldp port 1/1-1/4
```

LLDP Admin Port Status				
Port	AdminStatus	ConfigNotificationEnable	CdpAdminState	Dual-Mode
1/1	txAndRx	disabled	disabled	disabled
1/2	txAndRx	disabled	disabled	disabled
1/3	txAndRx	disabled	disabled	disabled
1/4	txAndRx	disabled	disabled	disabled

View the LLDP reception statistics. You can also view this on a specific port.

```
Switch:1#show lldp rx-stats
```

LLDP Rx-Stats						
Port Num	Frames Discarded	Frames Errors	Frames Total	TLVs Discarded (Non FA)	TLVs Unsupported (Non FA)	AgeOuts
1/1	0	0	0	0	0	0
1/2	0	0	0	0	0	0
1/3	0	0	0	0	0	0
1/4	0	0	0	0	0	0
...						
...						

View the LLDP statistics:

```
Switch:1#show lldp stats
=====
                        LLDP Stats
=====
Inserts      Deletes      Drops      Ageouts
-----
4             0             0             0
-----
```

View the LLDP transmission statistics:

```
Switch:1#show lldp tx-stats
=====
                        LLDP Tx-Stats
=====
PORT NUM      FRAMES
-----
1/1            95
1/2            95
1/3            95
1/4            95
1/5            95
...
...
```

Display LLDP-MED local system data:

```
Switch:1#show lldp local-sys-data med
=====
                        LLDP Local System Data
=====

ChassisId: MAC Address      f4:ce:48:9f:50:00
SysName   : 5320-24P-8XE-FabricEngine
SysDescr  : 5320-24P-8XE-FabricEngine (9.0.0.0)
SysCap    : Br / Br

MED Capabilities:          CNLSI
MED Device Type:          Network Connectivity Device
MED Power Device Type:    PSE Device
HWRev: 02                  FWRev: -
SWRev: 9.0.0.0             SerialNumber: TB012132K-H0057
ManufName: Extreme Networks. ModelName: 5320-24P-8XE-FabricEngine
Asset ID: TB012132K-H0057
-----

Capabilities Legend: (Supported/Enabled)
B= Bridge,    D= DOCSIS,    O= Other,    R= Repeater,
S= Station,   T= Telephone, W= WLAN,      r= Router
-----

MED Capabilities Legend: (Supported/Enabled)
C= Capabilities, N= Network Policy; L= Location Identification;
I= Inventory; S= Extended Power via MDI - PSE; D= Extended Power via MDI - PD.
```


Variable Definitions

The following table defines parameters for the **show lldp** command.

Variable	Value
<i>local-sys-data</i>	Displays the LLDP local system data.
<i>neighbor detail</i>	Displays detailed LLDP neighbor information. You can also view this on a specific port.
<i>neighbor summary</i>	Displays a summary of LLDP neighbor(s). You can also view this on a specific port.
<i>port</i> [{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]]}	Displays the LLDP administrative status of a port or all ports on the switch.
<i>rx-stats</i> [port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]]}	<p>Displays the LLDP reception statistics on all ports on the switch, or on a specific port.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>stats</i>	Displays the LLDP statistics.
<i>tx-stats</i> [port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]]}	<p>Displays the LLDP transmission statistics on all ports on the switch or on a specific port.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <p>] t</p>

View LLDP Neighbor Information

About This Task

Use this procedure to display LLDP neighbor information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View LLDP neighbor information:

```
show lldp neighbor [detail] [dot3] [med] [port {slot/port[/sub-  
port] [-slot/port[/sub-port]] [,...][slot/all][all]] [summary] [vendor-  
specific]
```

Variable Definitions

The following table defines parameters for the **show lldp neighbor** command.

Variable	Value
detail	Displays LLDP neighbor-specific information.
dot3	Displays remote LLDP dot3 information.
med	Displays LLDP neighbors learned based on LLDP-MED TLV information.
port {slot/ port[/sub-port] [-slot/port[/sub-port]] [,...][slot/ all][all]}	<p>Displays LLDP neighbor information on the specified port. Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
summary [port {slot/ port[/sub-port] [-slot/port[/sub-port]] [,...][slot/ all][all]}	<p>Displays the summary of LLDP neighbors of a port or all ports on the switch. Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
vendor-specific	Displays the remote LLDP vendor-specific call server and file server information.

View Global LLDP Statistics

Use this procedure to view and verify global LLDP statistics.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View LLDP statistics:
show lldp stats
3. View LLDP reception statistics:
show lldp rx-stats
4. View LLDP transmission statistics:
show lldp tx-stats
5. (Optional) Clear global LLDP statistics:
clear lldp stats [summary]

Example

View LLDP statistics:

```
Switch:1>enable
Switch:1#show lldp stats
```

LLDP Stats			
Inserts	Deletes	Drops	Ageouts
0	0	0	0

View LLDP transmission statistics:

```
Switch:1#show lldp tx-stats
```

LLDP Tx-Stats	
PORT NUM	FRAMES
1/2	100

View LLDP reception statistics:

```
Switch:1#show lldp rx-stats
```

LLDP Rx-Stats						
Port Num	Frames Discarded	Frames Errors	Frames Total	TLVs Discarded	TLVs Unrecognized	AgeOuts
1/2	0	0	46	0	0	0

View Port-Based LLDP Statistics

Use this procedure to verify port-based LLDP statistics.

About This Task

LLDP operates at the interface level. Enabling FA on a port automatically enables LLDP transmission and reception on the port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on all ports in the MLT.



Note

When FA is enabled on ports in an MLT or LACP MLT, tagging is enabled and spanning tree is disabled on those ports.

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Verify successful LLDP transmission on a port:

```
show lldp tx-stats port {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```

3. Verify that a port receives LLDP PDUs successfully:

```
show lldp rx-stats port {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```

4. (Optional) Clear LLDP statistics on a port:

```
clear lldp stats {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}
```

Example

Verify LLDP transmission statistics on a port:

```
Switch:1>en
Switch:1#show lldp tx-stats port 1/2
=====
LLDP Tx-Stats
=====
PORT NUM          FRAMES
-----
1/2                100
```

Verify that the port is receiving LLDP PDUs:

```
Switch:1#show lldp rx-stats port 1/2
```

LLDP Rx-Stats						
Port Num	Frames Discarded	Frames Errors	Frames Total	TLVs Discarded (Non FA)	TLVs Unsupported (Non FA)	AgeOuts
1/2	0	0	46	0	0	0

LLDP-MED Configuration Using CLI

Configure LLDP-MED information for local and remote systems on specific ports. LLDP-MED is enabled by default and all its TLVs are enabled for transmission.

To configure LLDP-MED TLVs in the LLDP PDUs on an interface:

- Configure LLDP-MED.
- Configure LLDP-MED network policy and location information.
- The switch automatically configures LLDP-MED capabilities, power, and inventory information.

Configure LLDP-MED Network Policies on Ports

About This Task

Perform this procedure to configure network policies on specific ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure a network policy:

```
lldp med-network-policies {guest-voice | guest-voice-signaling |  
softphone-voice | streaming-video | video-conferencing | video-  
signaling | voice | voice-signaling} [dscp <0-63>] [priority <0-7>]  
[tagging {tagged|untagged}] [vlan-id <0-4059>]
```

Example

Configuring guest voice network policy on port 1/2:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#lldp med-network-policies guest-voice dscp 1 priority 5 tagging
tagged vlan-id 5
```

Variable Definitions

The following table defines parameters for the **lldp med-network-policies** command.

Variable	Value
<i>{guest-voice guest-voice-signaling softphone-voice streaming-video video-conferencing video-signaling voice voice-signaling}</i>	Specifies the type of network policy.
<i>dscp <0-63></i>	Specifies the Layer 3 DiffServ Code Point (DSCP) value, as defined in IETF RFC 2474 and RFC 2475. The default is 0.
<i>priority <0-7></i>	Specifies the priority level, as defined in IEEE 802.1D. The default is 0.
<i>tagging {tagged untagged}</i>	Specifies the type of VLAN tagging to apply on the selected ports. The default is untagged.
<i>vlan-id <0-4059></i>	Specifies the VLAN ID for the port, as defined in IEEE 802.1Q. If you configure priority tagged frames, the system recognizes only the 802.1D priority level and uses a value of 0 for the VLAN ID of the ingress port. The default is 0.

Configure LLDP-MED Civic Address Location Information

About This Task

Perform the following procedure to configure civic address location information of local LLDP-MED on specific ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the civic address location by configuring the country-code and at least one other location parameter:

```
lldp location-identification civic-address country-code WORD<2-2>  
(additional-code additional-information apartment block building city  
city-district county floor house-number house-number-suffix landmark  
leading-street-direction name place-type pobox postal-community-name  
postal-zip-code room-number state street street-suffix trailing-  
street-suffix) WORD<0-255>
```



Note

If you try to configure a civic-address with a large number of arguments, 26 or more, the command fails and a software message informs you to split the command into multiple smaller commands.

Example

Configuring civic address location on port 2/1:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 2/1  
Switch:1(config-if)#lldp location-identification civic-address country-code US city New  
York
```

Variable Definitions

The following table defines parameters for the **lldp location-identification civic-address** command.

Variable	Value
additional-code WORD<0-255>	Specifies the location information parameters.
additional-information WORD<0-255>	Example: South Wing
apartment WORD<0-255>	Example: Apt 42
block WORD<0-255>	Specifies a block, e.g. 3

Variable	Value
building WORD<0–255>	Example: Low Library
city WORD<0–255>	Specifies a city, e.g. Sunnyvale
city-district WORD<0–255>	Specifies a city district, e.g. Santa Clara
country-code WORD<2–2>	Specifies a country using a 2 character string, example US (United States), CA (Canada).
county WORD<0–255>	Specifies a county, e.g. Alameda
floor WORD<0–255>	Example: 8
house-number WORD<0–255>	Specifies a house number, e.g. 123
house-number-suffix WORD<0–255>	Specifies a house number suffix, e.g. A, 1/2
landmark WORD<0–255>	Specifies a landmark, e.g. Columbia University
leading-street-direction WORD<0–255>	Specifies a leading street direction, e.g. N
name WORD<0–255>	Example: Joe's Barbershop
place-type WORD<0–255>	Example: office
pobox WORD<0–255>	Example: 12345
postal-community-name WORD<0–255>	Example: Leonia
postal-zip-code WORD<0–255>	Specifies a postal or zip code, e.g. 95054
room-number WORD<0–255>	Example: 450F
state WORD<0–255>	Specifies a state, e.g. NJ, FL
street WORD<0–255>	Specifies a street, e.g. Great America Parkway
street-suffix WORD<0–255>	Specifies a street suffix, e.g. Ave, Blvd
trailing-street-suffix WORD<0–255>	Specifies a trailing street suffix, e.g. SW

Configure LLDP-MED Coordinate Based Location Information

About This Task

Perform the following procedure to configure coordinate based location information of local LLDP-MED on specific ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure coordinate based location:

```
lldp location-identification coordinate (altitude WORD<1-13> {floors |  
meters} datum {NAD83/MLLW | NAD83/NAVD88 | WGS84} latitude WORD<1-14>  
{NORTH | SOUTH} longitude WORD<1-14> {EAST | WEST})
```

Example

Configuring coordinate based location on port 1/2:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 1/2  
Switch:1(config-if)#lldp location-identification coordinate-base altitude 3 floors
```

Variable Definitions

The following table defines parameters for the **lldp location-identification coordinate** command.

Variable	Value
<i>altitude</i> WORD<1-13>	Specifies the value for altitude. The units of measurement are: <ul style="list-style-type: none"> • <i>floors</i> • <i>meters</i>
<i>datum</i>	Specifies the reference datum. The formats are: <ul style="list-style-type: none"> • NAD83/MLLW • NAD83/NAVD88 • WGS84
<i>latitude</i> WORD<1-14>	Specifies the latitude in degrees, and its relation to the equator from North or South.
<i>longitude</i> WORD<1-14>	Specifies the longitude in degrees, and its relation to the prime meridian from East or West.

Configure LLDP-MED Emergency Call Service Location

Perform the following procedure to configure emergency call service location of local LLDP-MED on specific ports.

About This Task

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure emergency call service location:

```
lldp location-identification ecs-elin WORD<10-25>
```

Example

Configuring emergency call service location on port 2/1-2/3:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitEthernet 2/1-2/3  
Switch:1(config-if)#lldp location-identification ecs-elin 123456789
```

Variable Definitions

The following table defines parameters for the **lldp location-identification ecs-elin** command.

Variable	Value
WORD<10-25>	Specifies the emergency line information number for emergency call service.

Display Local LLDP-MED Location Information

About This Task

Perform this procedure to display location information of the LLDP-MED configured locally.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display location information for local LLDP-MED:

```
show lldp [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...]  
[slot/all][all]] location-identification
```

Example

```
Switch:1>enable  
Switch:1#show lldp port 1/1-1/3 location-identification  
=====
```

LLDP-MED Location Information

```
=====
```

Port: 1/1

MED Location - Coordinate-based LCI:
Latitude: +12.3 (degrees) North
Longitude: +42 (degrees) East
Altitude: +45 (meters)
Datum: World Geodesic System (WGS84)
MED Location - Civic Address LCI:
Country code: RO
Country: Romania
City: Bucuresti
Block: 12
Street: Calea Floreasca
Floor: 3
MED Location - Emergency Call Service ELIN:
ECS ELIN: 121416182022

Port: 1/2

MED Location - Civic Address LCI:
Country code: RO
Country: Romania
City: Bucuresti
Block: 12
Street: Calea Floreasca
Floor: 3

Port: 1/3

MED Location - Coordinate-based LCI:
Latitude: +12.3 (degrees) North
Longitude: +42 (degrees) East
Altitude: +45 (meters)
Datum: World Geodesic System (WGS84)
MED Location - Emergency Call Service ELIN:
ECS ELIN: 121416182022

Display LLDP-MED Local Network Policies Configuration

About This Task

Perform this procedure to display LLDP-MED network policies locally configured on specific ports.

Procedure

1. Enter Privileged EXEC mode:
enable

2. Display LLDP-MED network policies configured:

```
show lldp [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...]  
[slot/all][all]] med-network-policies [guest-voice | guest-voice-  
signaling | softphone-voice | streaming-video | video-conferencing |  
video-signaling | voice | voice-signaling]
```

Example

```
Switch:1>enable  
Switch:1#show lldp med-network-policies
```

LLDP-MED Network Policies						
Port	Application Type	VlanID	Tagging	DSCP	Priority	Origin
1/2	Voice	4	Untagged	0	0	CONFIG
1/2	Guest Voice	0	Untagged	3	0	CONFIG

Link Layer Discovery Protocol configuration using EDM

This section describes how to configure LLDP on your switch using EDM.

Configure LLDP Global Information

Use this procedure to configure or view LLDP global information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **Globals** tab.
4. After you make the required configuration changes, select **Apply** to save changes.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Field	Description
lldpMessageTxInterval	Specifies the interval at which LLDP messages are transmitted. The default is 30 seconds.
lldpMessageTxHoldMultiplier	Specifies the multiplier used to calculate the time-to-live (TTL) value of an LLDP message. The default value is 4 seconds.
lldpReinitDelay	Specifies the delay in seconds between the time a port is disabled and the time it is re-initialized. The default is 1 second.

Field	Description
lldpTxDelay	Specifies the delay in seconds between successive LLDP transmissions. The default is 1 second. The value is as follows: $1 < \text{lldpTxDelay} < (0.25 \times \text{lldpMessageTxInterval})$
lldpNotificationInterval	Specifies the time interval between successive LLDP notifications. It controls the transmission of notifications. The default is 5 seconds.
Stats	
RemTablesLastChangeTime	Specifies the timestamp of LLDP missed notification events on a port, for example, due to transmission loss.
RemTablesInserts	Specifies the number of times the information advertised by a MAC Service Access Point (MSAP) is inserted into the respective tables.
RemTablesDeletes	Specifies the number of times the information advertised by an MSAP is deleted from the respective tables.
RemTablesDrops	Specifies the number of times the information advertised by an MSAP was not entered into the respective tables.
RemTablesAgeouts	Specifies the number of times the information advertised by an MSAP was deleted from the respective tables.

Configure LLDP on a Port

Use this procedure to view and change the LLDP port configuration.

About This Task

To modify a value, double-click on a table cell, and then select a value.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **Port** tab.
4. In **AdminStatus**, configure LLDP frame direction on the port.
5. In **NotificationEnable**, configure notification status.
6. In **TLVsTxEnable**, configure which TLVs LLDP transmits:
 - a. To enable a TLV, select one or more values, and then select **Ok**.
 - b. To enable all TLVs, select **Select All**, and then select **Ok**.
 - c. To disable all TLVs, select **Disable All**, and select **Ok**.
7. In **TLVsTxEnable(med)**, configure which TLVs LLDP transmits for MED devices:
 - a. To enable a TLV, select one or more values, and then select **Ok**.
 - b. To enable all TLVs, select **Select All**, and then select **Ok**.

- c. To disable all TLVs, select **Disable All**, and select **Ok**.
8. Select **Apply** to save configuration changes.
9. Select **Refresh** to verify the configuration.

Port Field Descriptions

Use the data in the following table to use the **Port** tab.

Name	Description
PortNum	Displays the port number.
AdminStatus	<p>Specifies the administrative status of the port. The options are as follows:</p> <ul style="list-style-type: none"> • txOnly: LLDP frames are only transmitted on this port. • rxOnly: LLDP frames are only received on this port. • txAndRx: LLDP frames are transmitted and received on this port. • disabled: LLDP frames are neither transmitted or received on this port. Any information received on this port from remote systems before this is disabled, ages out. <p>The default is disabled.</p>
NotificationEnable	<p>Specifies whether the port is enabled or disabled for notifications.</p> <ul style="list-style-type: none"> • true: indicates that the notifications are enabled. • false: indicates that the notifications are disabled. <p>The default is false.</p>
TLVsTxEnable	<p>Specifies the set of TLVs whose transmission using LLDP is always allowed by network management.</p> <p>The following list describes the TLV types:</p> <ul style="list-style-type: none"> • portDesc • sysName • sysDesc • sysCap <p>The default is an empty set of TLVs.</p>
CdpAdminState	<p>Configures the CDP administrative status of the port. Configure this parameter to true to enable the Industry Standard Discovery Protocol (ISDP) on the port. ISDP is CDP-compatible.</p> <p>The default is false.</p> <p>By default, if you enable CDP, the interface accepts only CDP packets. Similarly, if you disable CDP but enable LLDP, the interface accepts only LLDP packets. You can optionally configure the interface to send and receive both LLDP and CDP packets.</p>

Name	Description
DualMode	Configures the port to send and receive both CDP and LLDP packets. The default is false. If you plan to configure a CDP-enabled port for features that rely on LLDP like Auto-sense or Fabric Attach (FA), ensure you enable both CDP and LLDP packet transmission on the port.
CapSupported(med)	Specifies the LLDP-MED capabilities. The default is enabled.
TLVsTxEnable(med)	Specifies the set of TLVs whose transmission using LLDP is always allowed by network management. The following list describes the TLV types: <ul style="list-style-type: none"> • capabilities • networkPolicy • location • extendedPSE • inventory The default is an empty set of TLVs.

NEW Enable CDP Mode on a Port

To configure the switch as CDP-compatible, you must enable the Industry Standard Discovery Protocol (ISDP) on a port, or ports, on the switch.

About This Task

By default, if you enable CDP, the interface accepts only CDP packets. Similarly, if you disable CDP but enable LLDP, the interface accepts only LLDP packets. You can optionally configure the interface to send and receive both LLDP and CDP packets.

If you plan to configure a CDP-enabled port for features that rely on LLDP like Auto-sense or Fabric Attach (FA), ensure you enable both CDP and LLDP packet transmission on the port.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **Port** tab.
4. In **CdpAdminState**, select **true** to enable CDP,.
5. (Optional) In **DualMode**, select **true** to configure the port to send and receive both CDP and LLDP packets.
6. Select **Apply**.

View LLDP Transmission Statistics

Use this procedure to view the LLDP transmission statistics. You can also view the statistics graphically.

About This Task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on all ports in that MLT.



Note

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **TX Stats** tab.
4. (Optional) To view the transmission statistics graphically for a port, see [Graph LLDP Transmission Statistics](#) on page 2498.

TX Stats Field Descriptions

Use the data in the following table to use the **TX Stats** tab.

Name	Description
PortNum	Specifies the port number.
FramesTotal	Specifies the total number of LLDP frames transmitted.

Graph LLDP Transmission Statistics

Use this procedure to view the LLDP transmission (TX) statistics. You can also view the statistics graphically.

About This Task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on all ports in that MLT.

**Note**

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. However, if the mappings are learned on another port on the MLT, the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **TX Stats** tab.
The system displays the transmission statistics.
4. To view the transmission statistics graphically for a port:
 - a. Select a row, and select **Graph**.
The system displays the **TX Stats-Graph,<port-number>** tab.
 - b. To view the graph, select one of the parameters, and then select the appropriate icon on the upper-left corner of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
5. To clear the existing counters, and fix a reference point in time to restart the counters, select **Clear Counters**.
6. To export the statistical data to a file, select **Export**.
7. To configure a poll interval, select an appropriate value from **Poll Interval**.

TX Stats Field Descriptions

Use the data in the following table to use the **TX Stats** tab.

Name	Description
FramesTotal	Specifies the total number of LLDP frames transmitted.

View LLDP Reception Statistics

Use this procedure to view the LLDP reception statistics. You can also view these statistics graphically.

About This Task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on all ports in that MLT.



Note

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. If however, the mappings are learned on another port on the MLT, then the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **RX Stats** tab.
4. (Optional) To view the reception statistics graphically for a port:
 - a. Select a row, and then select **Graph**.
 - b. Select one or more of the parameters, and then select the appropriate icon on the top left of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
 - c. Select **Refresh Counters** to clear the existing counters, and fix a reference point in time to restart the counters.
 - d. Select **Export** to export the statistical data to a file.
 - e. To fix a poll interval, select an appropriate value from **Poll Interval**.

RX Stats Field Descriptions

Use the data in the following table to use the **RX Stats** tab.

Name	Description
PortNum	Specifies the port number.
FramesDiscardedTotal	Specifies the number of LLDP frames received on the port, but discarded, for any reason. This counter provides an indication of possible LLDP header formatting problems in the sending system, or LLDP PDU validation problems in the receiving system.

Name	Description
FramesErrors	Specifies the number of invalid LLDP frames received on the port.
FramesTotal	Specifies the total number of LLDP frames received on the port.
TLVsDiscardedTotal	Specifies the number of LLDP TLVs discarded on the port, for any reason.
TLVsUnrecognizedTotal	Specifies the number of LLDP TLVs on the port, that are unrecognized on that port. An unrecognized TLV is referred to as the TLV whose type value is in the range of reserved TLV types (000 1001–111 1110). An unrecognized TLV could be, for example, a basic management TLV from a later LLDP version.
AgeoutsTotal	Specifies the number of LLDP age-outs that occur on a specific port. An age-out is the number of times the complete set of information advertised by a particular MSAP is deleted, because the information timeliness interval has expired.

Graph LLDP Reception Statistics

Use this procedure to graphically view the LLDP reception statistics.

About This Task

LLDP operates at the port interface level. Enabling FA on a port automatically enables LLDP transmission and reception on that port. It also enables traffic tagging and disables spanning tree on that port.

Enabling FA on an MLT enables LLDP transmission and reception on all ports in that MLT.



Note

When a port is removed from an MLT, LLDP transmission on that port stops and spanning tree is enabled. Any I-SID-to-VLAN mappings on that port are removed, if not already learned on any other port in the MLT. This also causes the Switched UNI I-SID to be deleted from the MLT. However, if the mappings are learned on another port on the MLT, the Switched UNI I-SID continues to exist for that MLT.

For ports in an LACP MLT, when FA is enabled, tagging is enabled on all ports in the LACP MLT. The consistency check for FA is based on key membership. If all ports with the same key do not support FA, FA is not successfully enabled on those ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.

3. Select the **RX Stats** tab.
4. To view the reception statistics graphically for a port:
 - a. Select a row, and then select **Graph**.
The system displays the **RX Stats-Graph,<port-number>** tab.
 - b. Select one of the parameters, and then select the appropriate icon in the upper-left corner of the menu bar to draw a line chart, area chart, bar chart or a pie chart.
5. To clear the existing counters, and fix a reference point in time to restart the counters, select **Clear Counters**.
6. To export the statistical data to a file, select **Export**.
7. To configure a poll interval, select an appropriate value from **Poll Interval**.

RX Stats Field Descriptions

Use the data in the following table to use the **RX Stats** tab.

Name	Description
FramesDiscardedTotal	Specifies the number of LLDP frames received on the port, but discarded, for any reason. This counter provides an indication of possible LLDP header formatting problems in the sending system, or LLDP PDU validation problems in the receiving system.
FramesErrors	Specifies the number of invalid LLDP frames received on the port.
FramesTotal	Specifies the total number of LLDP frames received on the port.
TLVsDiscardedTotal	Specifies the number of LLDP TLVs discarded on the port, for any reason.
TLVsUnrecognizedTotal	Specifies the number of LLDP TLVs on the port, that are unrecognized on that port. An unrecognized TLV is referred to as the TLV whose type value is in the range of reserved TLV types (000 1001–111 1110). An unrecognized TLV could be, for example, a basic management TLV from a later LLDP version.

View LLDP Local System Information

Use this procedure to view the LLDP local system information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.

3. Select the **Local System** tab.

Local System *Field Descriptions*

Use the data in the following table to use the **Local System** tab.

Name	Description
ChassisIdSubType	Indicates the encoding used to identify the local system chassis. <ul style="list-style-type: none">• chassisComponent• interfaceAlias• portComponent• macAddress• networkAddress• interfaceName• local
ChassisId	Indicates the chassis ID of the local system.
SysName	Indicates local system name.
SysDesc	Indicates local system description.
SysCapSupported	Indicates the system capabilities supported on the local system.
SysCapEnabled	Indicates the system capabilities that are enabled on the local system.

View LLDP Local Port Information

Use this procedure to view the LLDP local port information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **Local Port** tab.

Local port Field Descriptions

Use the data in the following table to use the **Local Port** tab.

Name	Description
PortNum	Specifies the port number.
PortIdSubType	Specifies the type of port identifier. <ul style="list-style-type: none">• interfaceAlias• portComponent• macAddress• networkAddress• interfaceName• agentCircuitId• local
PortId	Specifies the identifier associated with the port, on the local system.
PortDesc	Specifies the description of the port, on the local system.

View LLDP CDP Remote Information

Use this procedure to view the remote device information received through the Industry Standard Discovery Protocol (ISDP) (CDP-compatible) messages on the local interface.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **CDP Remote** tab.

CDP Remote Field Descriptions

Use the data in the following table to use the **CDP Remote** tab.

Name	Description
PortNum	Indicates the port number on which the remote system information is received.
CallServers	Indicates the remote call server IP addresses that are received on this port.
FileServers	Indicates remote file server IP addresses that are received on this port.

View LLDP Neighbor Information

Use this procedure to view the LLDP neighbor information.

About This Task

LLDP IPv6 neighbors do not display in EDM. LLDP IPv6 is only supported in CLI. To display LLDP IPv6 neighbors, use the **show lldp neighbor summary** command.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **LLDP**.
3. Select the **Neighbor** tab.

Neighbor Field Descriptions

Use the data in the following table to use the **Neighbor** tab.

Name	Description
TimeMark	Indicates the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Identifies the port on which the remote system information is received.
Index	Indicates a particular connection instance that is unique to the remote system.
ProtocolType	Indicates whether the entry protocol is CDP or LLDP.
SysName	Indicates the name of the remote system.
IpAddress	Indicates the neighbor's IP address.
PortIdSubType	Indicates the type of encoding used to identify the remote port.
PortId	Indicates the remote port ID.
PortDesc	Indicates the remote port description.
ChassisIdSubtype	Indicates the type of encoding used to identify the remote system chassis. <ul style="list-style-type: none"> • chassisComponent • interfaceAlias • portComponent • macAddress • networkAddress • interfaceName • local
ChassisId	Indicates the chassis ID of the remote system.
SysCapSupported	Identifies the system capabilities supported on the remote system.
SysCapEnabled	Identifies the system capabilities enabled on the remote system.
SysDesc	Indicates the description of the remote system.

LLDP-MED Configuration Using EDM

Configure LLDP-MED information for local and remote systems on specific ports. LLDP-MED is enabled by default and all its TLVs are enabled for transmission.

To configure LLDP-MED TLVs in the LLDP PDUs on an interface:

- Configure LLDP-MED.
- Configure LLDP-MED location information.
- The switch automatically configures LLDP-MED capabilities, power, and inventory information.

Configure LLDP-MED Network Policies on Ports

About This Task

Perform this procedure to configure LLDP-MED network policies on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Local Policy** tab.
4. Select **Insert**.
5. In **PortNum**, select the ellipsis (...).
6. In **Port Editor: PortMembers** dialog box, select the desired ports.
7. Select **OK**.
8. In **PolicyAppType**, select the application type.
9. (Optional) In **PolicyVlanId**, type the VLAN ID for the port.
10. (Optional) In **PolicyPriority**, type the priority level.
11. (Optional) In **PolicyDscp**, type DSCP value.
12. (Optional) Select **Policy Tagged** to enable VLAN tagging on the port.
13. Select **Insert**.

Local Policy Field Descriptions

Use the data in the following table to use the **Local Policy** tab.

Name	Description
PortNum	Specifies the port.
PolicyAppType	Specifies the application type.
PolicyVlanId	Specifies the VLAN ID for the port, as defined in IEEE 802.1Q-2003. The value 0 is used if the device is using priority tagged frames, which means only the 802.1D priority level is significant, and the default VLAN ID of the ingress port is used instead.

Name	Description
PolicyPriority	Specifies the Layer 2 priority used for the specified application type, as defined in IEEE 802.1D-2004. The default is 0.
PolicyDscp	Specifies the value of the Differentiated Service Code Point (DSCP) associated with a specific port on the local LLDP-MED, as defined in IETF RFC 2474 and RFC 2475. The default is 0.
PolicyTagged	Specifies whether the application uses a tagged or untagged VLAN, as defined by IEEE 802.1Q-2003. <ul style="list-style-type: none">• true — uses tagged VLAN• false — uses untagged VLAN or does not support a port-based VLAN

Add LLDP-MED Local Location Information

About This Task

Perform this procedure to add location information of local LLDP-MED configured on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Local Location** tab.
4. In the **LocationInfo** column, double-click on the cell, and type the information.
5. Select **Apply**.

Local Location Field Descriptions

Use the data in the following table to use the **Local Location** tab.

Name	Description
PortNum	Specifies the port number.
LocationSubtype	Specifies the location subtype of the local LLDP-MED: <ul style="list-style-type: none">• coordinateBased• civicAddress• elin
LocationInfo	Specifies the location information of local LLDP-MED. The parsing of this information depends on the location subtype.

View LLDP-MED Local PoE PSE Information

About This Task

Perform this procedure to view PoE Power Sourcing Entity (PSE) information for local LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Local PoE PSE** tab.

Local PoE PSE Field Descriptions

Use the data in the following table to use the **Local PoE PSE** tab.

Name	Description
PortNum	Specifies the port.
PSEPortPowerAvailable	Specifies the value of the power available (in units of 0.1 watts) from the PSE through the specific port.
PSEPortPDPriority	Specifies the Power Device (PD) power priority for the PSE port (see RFC 3621): <ul style="list-style-type: none">• unknown — priority is not configured or known by the PD• critical• high• low

View LLDP-MED Neighbor Capabilities Information

About This Task

Perform this procedure to view capabilities information for remote LLDP-MED on specific ports based on the information advertised by the remote device and received on each port in the capabilities TLV.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor Capabilities** tab.

Neighbor Capabilities *Field Descriptions*

Use the data in the following table to use the **Neighbor Capabilities** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.
Index	Specifies the connection instance unique to the remote LLDP-MED.
CapSupported	Specifies the LLDP-MED capabilities supported on the remote system.
CapCurrent	Specifies the LLDP-MED capabilities that are enabled on the remote system.
DeviceClass	Specifies the remote LLDP-MED device class.

View LLDP-MED Neighbor Policy Information**About This Task**

Perform this procedure to view policy information of remote LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor Policy** tab.

Neighbor Policy *Field Descriptions*

Use the data in the following table to use the **Neighbor Policy** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.
Index	Specifies the particular connection instance that is unique to the remote LLDP-MED.
PolicyAppType	Specifies the policy application type.

Name	Description
PolicyVlanId	Specifies the VLAN ID for the port, as defined in IEEE 802.1Q-2003. The value 0 is used if the device is using priority tagged frames, which means only the 802.1D priority level is significant, and the default VLAN ID of the ingress port is used instead.
PolicyPriority	Specifies the Layer 2 priority used for the specified application type, as defined in IEEE 802.1D-2004. The default value is 0.
PolicyDscp	Specifies the value of the Differentiated Service Code Point (DSCP) associated with a specific port on the remote LLDP-MED, as defined in IETF RFC 2474 and RFC 2475. The default value is 0.
PolicyUnknown	Specifies the network policy for the remote LLDP-MED is currently unknown.
PolicyTagged	Specifies whether the application uses a tagged or untagged VLAN, as defined by IEEE 802.1Q-2003. <ul style="list-style-type: none"> • true — uses tagged VLAN • false — uses untagged VLAN or does not support a port based VLAN

View LLDP-MED Neighbor Location Information

About This Task

Perform this procedure to view location information of remote LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor Location** tab.

Neighbor Location *Field Descriptions*

Use the data in the following table to use the **Neighbor Location** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.
Index	Specifies the particular connection instance that is unique to the remote LLDP-MED.
LocationSubType	Specifies the subtype of the remote LLDP-MED location: <ul style="list-style-type: none">• coordinateBased• civicAddress• elin
LocationInfo	Specifies the location information of the remote LLDP-MED.

View LLDP-MED Neighbor PoE Information**About This Task**

Perform this procedure to view PoE information of remote LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor PoE** tab.

Neighbor PoE *Field Descriptions*

Use the data in the following table to use the **Neighbor PoE** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.

Name	Description
Index	Specifies the particular connection instance that is unique to the remote LLDP-MED.
PoEDeviceType	Specifies the type of PoE LLDP-MED. <ul style="list-style-type: none">• PseDevice: specifies the device as a Power Sourcing Entity (PSE).• pdDevice: specifies the device as a Power Device (PD).• none: specifies that the device does not support PoE.

View LLDP-MED Neighbor PoE PSE Information

About This Task

Perform this procedure to view PoE Power Sourcing Entity (PSE) information for remote LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor PoE PSE** tab.

Neighbor PoE PSE Field Descriptions

Use the data in the following table to use the **Neighbor PoE PSE** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.
Index	Specifies the particular connection instance that is unique to the remote LLDP-MED.
PSEPowerAvailable	Specifies the power available from the PSE connected remotely to the specific port.

Name	Description
PSEPowerSource	Specifies the type of PSE Power Source for the remote LLDP-MED. <ul style="list-style-type: none">• unknown• primary• backup
PSEPowerPriority	Specifies the power priority associated with the PSE LLDP-MED, for more information, see RFC 3621. <ul style="list-style-type: none">• unknown• critical• high• low

View LLDP-MED Neighbor PoE PD Information

About This Task

Perform this procedure to view PoE Powered Device (PD) information for remote LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor PoE PD** tab.

Neighbor PoE PD Field Descriptions

Use the data in the following table to use the **Neighbor PoE PD** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.
Index	Specifies the particular connection instance that is unique to the remote LLDP-MED.
PDPowerReq	Specifies the value of the power required by a PD LLDP-MED connected remotely to the port.

Name	Description
PDPowerSource	Specifies the power source being utilized by the PD LLDP-MED. <ul style="list-style-type: none"> from PSE: specifies that the device advertises its power source as received from a PSE. local: specifies that the device advertises its power source as local. local and PSE: specifies that the device advertises its power source as using both local and PSE power.
PDPowerPriority	Specifies the priority of the PD LLDP-MED connected remotely to the port, for more information, see RFC 3621. <ul style="list-style-type: none"> unknown — priority is not configured for the PD LLDP-MED critical high low

View LLDP-MED Neighbor Inventory Information

About This Task

Perform this procedure to view inventory attributes for LLDP-MED on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > LLDP & MED**.
2. Select **Port MED**.
3. Select the **Neighbor Inventory** tab.

Neighbor Inventory Field Descriptions

Use the data in the following table to use the **Neighbor Inventory** tab.

Name	Description
TimeMark	Specifies the time filter. For more information about TimeFilter, see the TimeFilter textual convention in IETF RFC 2021.
LocalPortNum	Specifies the port on which the remote LLDP-MED information is received.
Index	Specifies the particular connection instance that is unique to the remote LLDP-MED.
HardwareRev	Specifies the current hardware revision of the LLDP-MED.

Name	Description
FirmwareRev	Specifies the current firmware revision of the LLDP-MED.
SoftwareRev	Specifies the current software revision of the LLDP-MED.
SerialNum	Specifies the current serial number of the LLDP-MED.
MfgName	Specifies the manufacturer of the LLDP-MED.
ModelName	Specifies the model name of the LLDP-MED.
AssetID	Specifies the asset tracking identifier for the LLDP-MED.



Link State Change Control

[Link Debounce](#) on page 2517

[Link State Change Control Using CLI](#) on page 2518

[Link State Change Control Using EDM](#) on page 2521

Rapid fluctuation in a port link state is called link flapping.

Link flapping is detrimental to network stability because it can trigger recalculation in spanning tree and the routing table.

If the number of port down events exceeds a configured limit during a specified interval, the system forces the port out of service.

You can configure link flap detection to control link state changes on a physical port. You can set thresholds for the number and frequency of changes allowed.

You can configure the system to take one of the following actions if changes exceed the thresholds:

- send a trap
- bring down the port

If changes exceed the link state change thresholds, the system generates a log entry.

Link Debounce

Table 184: Link Debounce for WAN Links

Feature	Product	Release introduced
Link Debounce	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.4.2
	VSP 7400 Series	VOSS 8.4.2

In a WAN environment, when a carrier-side link failure occurs, switchover on the carrier side can take a few hundred milliseconds. During that time, a lag in the sending and receiving of packets can occur. Use Link Debounce to hold the connection path until the switchover is complete. You can configure Link Debounce on each port.

Link Debounce protects the upper layers from unnecessary state changes by delaying the change of a port link state when the following situations occur:

- There are frequent flaps in a short interval at the physical layer in the case of Fiber WAN services.
- There is a delay in switching from the working path to the protected path in the case of Carrier Wave WAN services.

Link Debounce works only on Layer 1 protocol applications. Layer 2 / Layer 3 protocols make decisions based on how they receive packets. For example, STP makes the decision according to the lack of traffic and port up condition; OSPF and IS-IS can still fail adjacencies.

**Note**

You cannot configure Link Debounce on Integrated Application Hosting (IAH) ports (also known as Insight ports).

ExtremeCloud SD-WAN

Auto-sense automatically configures Link Debounce on the switch port that connects to SD-WAN Appliance. This configuration enables the switch that connects to the appliance LAN1 port to keep using its FE VXLAN tunnels over MPLS transport, even if SD-WAN Appliance is down, Layer 3 WAN Internet ports are lost, and the appliance is in Bypass mode.

If you do not configure a timer value and the port connects to SD-WAN Appliance, Auto-sense configures a value of 8000 milliseconds. Auto-sense does not overwrite a configured timer value.

Auto-sense Ports

The interface-level Link Debounce configuration does not apply to ports with Auto-sense enabled. Use Auto-sense commands to configure Link Debounce on Auto-sense UNI ports. For more information, see [Auto-sense Configuration using CLI](#) on page 34 or [Auto-sense Global Configuration using EDM](#) on page 57.

Link State Change Control Using CLI

This section provides procedures to detect and to control link flapping using the command line interface (CLI).

Control Link State Changes

Configure link flap detection to control state changes on a physical port.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the interval for link state changes:
`link-flap-detect interval <2-600>`
3. Configure the number of changes allowed during the interval:
`link-flap-detect frequency <1-9999>`
4. Enable automatic port disabling:
`link-flap-detect auto-port-down`
5. Enable sending a trap:
`link-flap-detect send-trap`

Example

Enable automatic disabling of the port:

```
Switch:1(config)#link-flap-detect auto-port-down
```

Configure the link-flap-detect interval:

```
Switch:1(config)#link-flap-detect interval 20
```

Enable sending traps:

```
Switch:1(config)#link-flap-detect send-trap
```

Variable Definitions

Use the data in the following table to use the **link-flap-detect** command.

Variable	Value
<code><auto-port-down></code>	Automatically disables the port if state changes exceed the link-flap threshold. By default, auto-port-down is enabled.
<code>frequency <1-9999></code>	Configures the number of changes that are permitted during the time specified by the <code>interval</code> command. The default is 20.
<code>interval <2-600></code>	Configures the link-flap-detect interval in seconds. The default value is 60.
<code>send-trap</code>	Activates traps transmission. The default setting is activated.

Display Link State Changes

Displays link flap detection state changes on a physical port.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display link state changes:
`show link-flap-detect`

Example

```
Switch:1>enable
Switch:1#show link-flap-detect

Auto Port Down : enable
Send Trap      : enable
Interval       : 60
Frequency      : 20
```

Configure Link Debounce

About This Task

Configure Link Debounce for ports to hold the connection path until the carrier side switchover is complete.



Note

Use this procedure to configure ports with Auto-sense disabled. To configure Link Debounce on Auto-sense ports, see [Configure Link Debounce for Auto-sense Ports](#) on page 56

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable Link Debounce:

```
link-debounce <0-300000>
```

3. Verify the configuration:

```
show interfaces gigabitEthernet link-debounce [{slot/port[/sub-port] [-
slot/port[/sub-port]][,...][slot/all][all]}
```

Examples

Configure the Link Debounce timer on port 1/2 to 300000 milliseconds.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#link-debounce 300000
```

Display the Link Debounce configuration for ports 1/1 through 1/4.

```
Switch:1>enable
Switch:1#show interfaces gigabitEthernet link-debounce 1/1-1/4
=====
                        Port link debounce
=====
-----
```

PORT NUM	LINK-DEBOUNCE STATUS	LINK-DEBOUNCE TIMER MILLISECONDS	DEBOUNCE COUNT	DEBOUNCE SUCCEEDED
1/1	Enable	1500	0	0
1/2	Enable	300000	0	0
1/3	Disable	--	0	0
1/4	Enable	4000	0	0

Variable Definitions

The following table defines parameters for the **link-debounce** command.

Variable	Value
<0-300000>	Link debounce configures time for ports in milliseconds. The default status is disabled for all ports when not initially configured. If you run the default link-debounce command, the default configuration is enabled with a value of 1,000 milliseconds. To return to the initial disabled state, you must run the no link-debounce command or set the Link Debounce timer to 0.

Link State Change Control Using EDM

This section provides procedures to detect and to control link flapping using Enterprise Device Manager (EDM).

Control Link State Changes

About This Task

Configure link flap detection to control link state changes on a physical port.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **General**.
3. Select the **Link Flap** tab.
4. Configure the parameters as required.
5. Select **Apply**.

Link Flap Field Descriptions

Use the data in the following table to use the **Link Flap** tab.

Name	Description
AutoPortDownEnable	Enables or disables Link Flap Detect. If you enable Link Flap Detect, the system monitors the number of times a port goes down during a designated interval. If the number of drops exceeds a specified limit, the system forces the port out-of-service. The default is enabled.
SendTrap	Specifies that a trap is sent if the port is forced out-of-service.

Name	Description
Frequency	Specifies the number of times the port can go down. The default is 20.
Interval	Specifies the interval (in seconds) between port failures. The default is 60.

Configure Link Debounce

About This Task

Configure Link Debounce for ports to hold the connection path until the carrier side switchover is complete.



Note

Use this procedure to configure ports with Auto-sense disabled. To configure Link Debounce on Auto-sense ports, see [Configure Link Debounce for Auto-sense Ports](#) on page 65.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select **Interface** tab.
5. In **LinkDebounce**, type the value in milliseconds to configure the timer.
6. Select **Apply**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Index	Displays the index of the port, written in the slot/port[/sub-port] format.
Name	Configures the name of the port.
Descr	Displays the description of the port. A textual string containing information about the interface.
Type	Displays the type of connector plugged in the port.
Mtu	Displays the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram which can be sent or received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.

Name	Description
PhysAddress	Displays the physical address of the port. The address of the interface at the protocol layer immediately below the network layer in the protocol stack. For interfaces which do not have such an address, for example, a serial line, this object should contain an octet string of zero length.
VendorDescr	Displays the vendor of the connector plugged in the port.
DisplayFormat	Identifies the slot and port numbers (slot/port). If the port is channelized, the format also includes the sub-port in the format slot/port/sub-port
AdminStatus	Configures the port as enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
OperStatus	Displays the current status of the port. The status includes enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
ShutdownReason	Indicates the reason for a port state change.
LastChange	Displays the timestamp of the last change.
LinkTrap	Enable or disable link trapping.
AutoNegotiate	Enables or disables Auto-Negotiation for this port. The default Auto-Negotiation behavior depends on the switch model and transceiver type.


Name	Description
AutoNegAd	<p>Specifies the port speed and duplex abilities to advertise during link negotiation.</p> <p>Supported speeds and duplex modes vary, depending on your hardware.</p> <p>The abilities specified in this object are only used when auto-negotiation is enabled on the port. If all bits in this object are disabled, and auto-negotiation is enabled on the port, then the physical link process on the port is disabled (if hardware supports this ability).</p> <p>Any change to this configuration restarts the auto-negotiation process, which has the same effect as physically unplugging and reattaching the cable attached to the port.</p> <p>If you select default, all capabilities supported by the hardware are advertised.</p>
AdminDuplex	Configures the administrative duplex setting for the port.
OperDuplex	Indicates the operational duplex setting for the port.
AdminSpeed	Configures the administrative speed for the port.
OperSpeed	Indicates the operational speed for the port.
QoSLevel	Selects the Quality of Service (QoS) level for this port. The default is level1.
DiffServ	Enables the Differentiated Service feature for this port. The default is disabled.
Layer3Trust	Configures if the system should trust Layer 3 packets coming from access links or core links only. The default is core.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled (selected) or disabled (cleared) on the port. The default is disabled (clear).
MLtId	Shows the MLT ID associated with this port. The default is 0.
Locked	Shows if the port is locked. The default is unlocked.
UnknownMacDiscard	Discards packets that have an unknown source MAC address, and prevents other ports from sending packets with that same MAC address as the destination MAC address. The default is disabled.

Name	Description
DirectBroadcastEnable	<p>Specifies that an Isolated Routing Port (IRP) can forward directed broadcast traffic. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcast on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling this function protects a host from possible denial of service (DoS) attacks.</p> <p>With the feature enabled, the Control Processor (CP) does not receive a copy of the directed broadcast. As a result, the system does not respond to a subnet broadcast ping sent from a remote subnet.</p> <p>The default is disabled.</p>
OperRouting	Shows the routing status of the port.
HighSecureEnable	Enables or disables the high secure feature for this port.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
FlexUniEnable	Enables Flex UNI on the port. The default is disabled.
EgressRateLimitState	Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the port. The default is disabled.
EgressRateLimit	<p>Specifies the egress rate limit in Kbps. You cannot configure the egress shaper rate to exceed the port capability.</p> <p>If you configure this value to 0, shaping is disabled on the port.</p>
TxFlowControl	<p>Configures if the port sends pause frames. By default, an interface does not send pause frames.</p> <p>You must also enable the flow control feature globally before an interface can send pause frames.</p>
TxFlowControlOperState	Shows the operational state of flow control.
BpduGuardTimerCount	Shows the time, starting at 0, since the port became disabled. When the BpduGuardTimerCount reaches the BpduGuardTimeout value, the port is enabled. Displays in 1/100 seconds.

Name	Description
BpduGuardTimeout	Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.
BpduGuardAdminEnabled	Enables BPDU Guard on the port. The default is disabled.
ForwardErrorCorrection	Configures one of the following options for Forward Error Correction (FEC) on the port: <ul style="list-style-type: none"> • CL 91 • CL 108 • CL 74 • disable • auto The disable option disables this configuration on the port.
ForwardErrorCorrectionApplicability	Displays whether FEC is applicable on the interface.
OperAutoNegotiate	Shows the operational state of Auto-Negotiation.
OperForwardErrorCorrection	Shows the negotiated operational FEC clause. If the value is off, the port supports FEC and is up but not configured for FEC. If the value is notApplicable, the port does not support FEC. If the value is unknown, the port supports FEC but is down.
Action	Performs one of the following actions on the port <ul style="list-style-type: none"> • none - none of the following actions • flushMacFdb - flush the MAC forwarding table • flushArp - flush the ARP table • flushIp - flush the IP route table • flushAll - flush all tables • triggerRipUpdate — manually triggers a RIP update The default is none.
Result	Displays the result of the selected action. The default is none.
AutoSense	Enables or disables Auto-sense on the specific port. The default value is disabled for existing configurations but enabled for new Zero Touch Fabric Configuration deployments.

Name	Description
AutoSenseKeepAutoConfig	Retains the Auto-sense configuration if you disable Auto-sense on the port. The dynamic configuration becomes a manual configuration and is visible in the show running-config output.
CustomAutoNegAdOrigin	Specifies the origin of Custom Auto Negotiation Advertisements (CANA) configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
BpduGuardOrigin	Specifies the origin of BPDU Guard configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
AutoSenseState	Displays the Auto-sense port state.
LinkDebounce	Specifies the extended debounce timer on the port. The value 0 milliseconds disables debounce time. The default value is 1000. Note: This field does not apply to Auto-sense ports. To configure Link Debounce on Auto-sense ports, see Configure Link Debounce for Auto-sense Ports on page 65.
AutoSenseDataIsid	Specifies the Auto-sense data I-SID per port.
AutoSenseWaitInterval	Specifies the Auto-sense Wait Interval for the port. The default is 35.
NodeAliasEnable	Enables or disables Node Alias on this specific port. The default is enabled on Auto-sense ports.

Name	Description
NodeAliasClear	Clears Node Alias on this specific port.
IpfixEnable	Enables or disables IPFIX on the port. The default is true, which means enabled. If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.



Link-state tracking (LST)

[Link-state tracking \(LST\) Overview](#) on page 2529

[LST Configuration Using CLI](#) on page 2531

[LST Configuration Using EDM](#) on page 2533

Table 185: Link-state tracking (LST) product support

Feature	Product	Release introduced
Link-state tracking (LST)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Link-state tracking (LST) Overview

Link-state tracking (LST) binds the link state of multiple interfaces, creating LST groups with upstream (to-be-followed) and downstream (to-follow) interfaces. LST monitors the state of upstream interfaces and automatically transfers the upstream state to the downstream interfaces. If all the upstream interfaces in a LST group are down, the downstream interfaces are administratively configured as down after approximately five seconds. If any upstream interface in a LST group is up, the downstream interfaces are not affected. The role of the LST group is to keep the downstream interfaces in the same state as the upstream interface.

An interface can be an aggregation of ports, multi-link trunks (MLT) or link aggregation groups (LAG). Interfaces can only belong to one LST group. You can configure LST using CLI or EDM. LST receives updates from Port Manager, MLT, and VLACP regarding the upstream state of ports and trunks in the group.

LST can detect a link failure of upstream interfaces and shutdown downstream interfaces, eliminating loss of traffic and allowing the source to reroute traffic. When

a LST group disables a downstream interface, the interface can only be enabled by the LST group. You can recover the downstream interfaces that LST disabled by removing the interfaces from the LST group or by disabling the LST group. You can administratively enable or disable LST group downstream interfaces with shutdown commands. A LST group cannot enable ports that you administratively disabled.

**Note**

If you administratively enable an interface which LST disabled, only the administrative status of the interface changes. The interface remains disabled until the LST group enables the interface, or until you remove the interface from the LST group.

MLT interactions

For MLTs, a last-link-down event triggers a down operational state and a first-link-up event triggers an up operational state. You cannot delete an MLT that is a member of a LST group.

LAG and LACP interactions

For LAGs, a static LAG trunk ID associated with a LACP administrative key can be a member of a LST group. You can add LAGs to a LST group by specifying the LAG trunk ID. You cannot break the association between trunk ID, LACP key, and ports while the LAG trunk is in a LST group. For LACP, you cannot add ports with link-aggregation enabled to a LST group, or enable link-aggregation on interfaces already in a LST group.

VLACP interactions

You can add LST group interfaces configured with VLACP. For upstream interfaces with VLACP enabled, when the physical link is up with a VLACP partner, the operational state is up. Otherwise the operational state is down. If VLACP is enabled, the value of the VLACP have partner field and the link status correspond to the operational state of upstream interfaces. For upstream interfaces with VLACP disabled, the up and down operational states correspond directly with the physical link.

SLPP Guard interactions

You can add LST group interfaces configured with SLPP Guard. When LST disables a port that is already disabled by SLPP Guard, the interface is unblocked by SLPP Guard and the blocking timer clears.

BPDU Guard and MACsec interactions

You can add LST group interfaces configured with BPDU Guard or MACsec. BPDU Guard and MACsec can enable or disable ports administratively. An interface is enabled if both LST and BPDU Guard or MACsec consider the port enabled. If BPDU Guard or MACsec disables the port, the port remains down and does not link up.

LST Configuration Using CLI

Configure Link-state tracking (LST) groups composed of upstream and downstream interfaces. Valid LST group members are switch ports, multi-link trunks (MLT), or link aggregation groups (LAG).

Configure LST

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the LST group upstream and downstream interface members:

```
link-state group <1-48> <upstream | downstream>
```

```
interface gigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
```

```
[,...][slot/all][all]}
```
3. Configure the LST group upstream and downstream mlt members:

```
link-state group <1-48> <upstream | downstream> mlt <1-512>
```
4. Enable the group:

```
link-state group <1-48> enable
```
5. Display the status of the groups:

```
show link-state group <1-48> [detail]
```

Example

The following example shows an LST group created with ports 1–4 and mlt 1 as upstream members and ports 11–14 and mlt 2 as downstream members. The LST group is enabled. Because port 1 is up, the LST group receives an up operational state and the downstream interfaces are enabled. Upstream ports 2, 4 and mlt 1 have VLACP admin enabled and are listed in the VLACP Upstream State section.

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# link-state group 1 upstream interface gigabitEthernet 1/1-1/4
Switch:1(config)# link-state group 1 upstream mlt 1
Switch:1(config)# link-state group 1 downstream interface gigabitEthernet 1/11-1/14
Switch:1(config)# link-state group 1 downstream mlt 2
Switch:1(config)# link-state group 1 enable
Switch:1(config)# show link-state group 1
```

```
=====
Link State Tracking General Info
=====
Group          1
Status         Enabled
Operational Status  UP
=====
```

```
Switch:1(config)# show link-state detail
```

```
=====
Link State Tracking Detailed Info
=====
Group:         1
Status:        Enabled
```

```

Operational Status:    UP
VLACP Upstream State:  Active on ports: 1/2, 1/4
                      Active on Trunks: 1

Upstream Ports:
1/1  (UP)  1/2  (DW)  1/3  (DW)  1/4  (DW)
Upstream Trunks:
1      (UP)
Downstream Ports:
1/11 (UP)  1/12 (DW)  1/13 (DW)  1/14 (UP)
Downstream Trunks:
2      (DW)

```

```

-----
Group                : 2
Status               : Disabled
Operational Status   : N/A
VLACP Upstream State : N/A
Upstream Ports       : Not Configured
Upstream Trunks      : Not Configured
Downstream Ports     : Not Configured
Downstream Trunks    : Not Configured
-----

```

Variable definitions

Use the data in the following table to use the **link-state** command.

Variable	Value
<i>enable</i>	Activates the action specified for the LST group or specified interfaces.
<i>group <1-48></i>	Specifies a LST group ID.
<i>interface gigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</i>	<p>Adds ports to the specified upstream or downstream LST group.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>mlt <1-512></i>	Adds an MLT to the specified upstream or downstream LSTgroup.
<i><upstream downstream></i>	Specifies the upstream or downstream interfaces in the LST group.

Use the data in the following table to use the **show link-state** command.

Variable	Value
detail	Displays the specified LST group status as enabled or disabled and if the operational status is up or down. Detail displays LST group additional information.

LST Configuration Using EDM

Configure Link-state tracking (LST) groups composed of upstream and downstream interfaces. Valid LST group members are switch ports, multi-link trunks (MLT), or link aggregation groups (LAG).

Configure LST

About This Task

Use this procedure to configure LST groups.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Link State Tracking**
3. Select the **Link State Tracking** tab.
4. Select a **GroupId** row to configure.
5. To add upstream ports to the LST group, double-click **UpstreamPortList**, select the upstream ports to add, and then select **OK**.
6. To add downstream ports to the LST group, double-click **DownstreamPortList**, select the upstream ports to add, and then select **OK**.
7. To add an upstream MLT to the LST group, double-click **UpstreamMltList**, and then type an upstream MLT ID.
8. To add a downstream MLT to the LST group, double-click **DownstreamMltList**, and then type an upstream MLT ID.
9. To enable an LST group, double-click **Enabled**, and then select **true**.
10. To activate the LST group configuration, select **Apply**.
11. View the **OperState** of the LST groups to verify the current operating state.

Link State Tracking Field Descriptions

Use the data in the following table to use the **Link State Tracking** tab.

Name	Description
GroupId	Specifies the LST group number between 1 and 48.
Enabled	Specifies if the LST group is enabled. Values are true or false. Default is false.
UpstreamPortList	Specifies upstream interface ports in the LST group.

Name	Description
DownstreamPortList	Specifies downstream interface ports in the LST group.
UpstreamMltList	Specifies an upstream multi-link trunk in the LST group.
DownstreamMltList	Specifies a downstream multi-link trunk in the LST group.
OperState	Displays the current operating status of the LST group. Values are up, down, or notConfigured.



Logs and Traps

[Logs and Traps Fundamentals](#) on page 2536
[Log Configuration Using CLI](#) on page 2548
[Log Configuration Using EDM](#) on page 2565
[SNMP Trap Configuration Using CLI](#) on page 2570
[SNMP Trap Configuration Using EDM](#) on page 2574
[Traps Reference](#) on page 2579

Table 186: Logs and traps product support

Feature	Product	Release introduced
Logging to a file and syslog (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Logging to a file and syslog (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 186: Logs and traps product support (continued)

Feature	Product	Release introduced
Simple Mail Transfer Protocol (SMTP) for log notification	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
System Logging compliance with RFC 5424 and RFC 3339	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
TLS client for secure syslog	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Logs and Traps Fundamentals

This section details SNMP traps and log files, available as part of the switch System Messaging Platform.

Overview of Traps and Logs

System Log Messaging

On a UNIX-based management platform, you can use system log (syslog) messaging to manage event messages. The switch syslog software communicates with a server software component named syslogd on the management workstation.

The UNIX daemon syslogd is a software component that receives and locally logs, displays, prints, and forwards messages that originate from sources internal and external to the workstation. For example, syslogd on a UNIX workstation concurrently handles messages received from applications that run on the workstation, as well as messages received from the switch that runs in a network accessible to the workstation.

The remote UNIX management workstation performs the following actions:

- Receives system log messages from the switch.
- Examines the severity code in each message.
- Uses the severity code to determine appropriate system handling for each message.

Log Consolidation

The switch generates a system log file and can forward that file to a syslog server for remote viewing, storage, and analyzing.

The system log captures messages for the following components:

- Central Processing Unit (CPU)
- Command line interface (CLI) log
- Extensible Authentication Protocol (EAP)
- filter
- hardware (HW)
- Internet Protocol (IP)
- MultiLink Trunking (MLT)
- policy
- Quality of Service (QoS)
- Remote Authentication Dial-in User Service (RADIUS)
- Remote Monitoring (RMON)
- Simple Network Management Protocol (SNMP) log
- software (SW)
- Virtual Local Area Network (VLAN)
- web

The switch can send information in the system log file, including CLI command log and the SNMP operation log, to a syslog server.

View logs for CLILOG module to track all CLI commands executed and for fault management purposes. The CLI commands are logged to the system log file as CLILOG module.

View logs for SNMPLOG module to track SNMP logs. The SNMP operation log is logged to the system log file as SNMPLOG module.

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs CLI Log and SNMP Log information regardless of the logging level you configure. This is not the case for other INFO messages.

System Log Client over IPv6 Transport

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration except in the following case. When you configure the system log table in EDM, under the **System Log Table** tab, you must select either IPv4 or IPv6.

Log Files with Enhanced Secure Mode

Enhanced secure mode allows the system to provide role-based access levels to log file commands. If you enable enhanced secure mode, the system encrypts the entire log file.

Log files are generated to `/inflash/shared`.

The current log file is protected against wiping for Telnet, SSH, FTP, SFTP, TFTP, and SCP applications for the following commands:

- Telnet and SSH:
 - mv
 - rename
 - delete
 - copy
 - cp
- FTP:
 - delete
 - mput
 - put
- TFTP:
 - put
- SCP:
 - destination file only

Log Commands Accessible for Various Users

The following table summarizes log file command access based on role-based access levels.

Access level role	Commands
Administrator and Privilege (except for the active log file)	The following commands: <ul style="list-style-type: none"> • edit • rename • copy • remove • delete
Administrator, Privilege, and Auditor	All show commands for log files.
All users (Administrator, Privilege, Auditor, Security, Operator)	All show commands for log configurations.

With enhanced secure mode enabled, authorized users can use Telnet, SSH, FTP, SFTP, TFTP, SCP to transfer files to a remote server with the content encrypted. Users with Administrator and Privilege access level roles can remove old log files but they cannot remove the active log file.

SNMP Traps

The SNMP trap is an industry-standard method used to manage events. You can set SNMP traps for specific types of log message (for example, warning or fatal), from specific applications, and send them to a trap server for further processing. For example, you can configure the switch to send SNMP traps to a server after a port is unplugged or if a power supply fails.

This section only describes SNMP commands related to traps. For more information about how to configure SNMP community strings and related topics, see [Simple Network Management Protocol \(SNMP\)](#) on page 3467.

Secure Syslog

Syslog is a standard used to send event log messages to devices within a network. The switch sends event messages to a logging server called syslog server. The syslog server stores the log messages and displays them for event reporting. Syslog messages are used for monitoring system activities and troubleshooting.

The secure syslog feature adds security and authenticated access to the plain text event log messages that are communicated between a remote syslog server and a syslog client. The secure syslog feature helps prevent unauthorized access to confidential data transmitted on an unsecured communication channel between a remote syslog server and client.

To implement the security, this feature employs port forwarding using the Transport Layer Security (TLS) to provide encrypted communication between a syslog server and client.

After starting the syslog server, to ensure authentication, you must set up a remote port forwarding connection to connect the switch with a remote TLS Server.

TLS Client for Secure Syslog

The syslog server is installed on a host that serves as a TLS Server. The switch plays the role of a TLS client for secure syslog.

A TLS handshake is initiated between the syslog server and the switch. The syslog server transmits a certificate which has a subject common name and an optional subject alternative name (SAN). The subject common name is always present in the certificate but the SAN is optional. The server-cert-name must match the SAN name, if present in the certificate. If the SAN name is not present, the server-cert-name must match the subject common name. Otherwise, TLS negotiation fails and the connection to the server is closed. If the server-cert-name part is not configured, this check is not done.

After the TLS handshake is successful, the log messages sent from the switch to the syslog server are encrypted. The syslog server decrypts these messages using a private key. The server then stores the messages or forwards them to other servers.

This feature supports the Rsyslog, syslog-ng, and stunnel, which are Linux based open-source syslog servers for TLS tunneling. Ensure you use the correct version of these applications or the switch displays a `no cipher match` error and the TLS communication fails.

Automatic Reconnection

If connectivity fails between the switch and the syslog server, the syslog host remains enabled and the switch attempts to reconnect with the syslog server every two minutes. This automatic reconnection occurs only if the following conditions are true:

- The syslog feature is globally enabled.
- The specific syslog host is enabled.
- The secure mode is TLS.
- The TLS session between the switch and syslog server is not active.



Note

If the route to the syslog server is removed from the switch and the IP address is not reachable, the switch does not try to reconnect after two minutes if there are no syslog messages in the queue. A TCP connection can take up to two hours to timeout after a connectivity issue. You can use the **logging write test** command to send a test message to the server and force the syslog application to reconnect. You can also configure the TCP keepalive interval using either CLI ([Configure TCP Keepalive and TCP Timestamp](#) on page 3424) or EDM ([Configure TCP Keepalive and TCP Timestamp](#) on page 3430).

Log messages on the console indicate connection issues between the switch and syslog server or with the TLS session. You can also use trace to diagnose connection issues. Use trace level 28 for syslog or trace level 115 for TLS.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) provides facilities to manage and monitor network resources. SNMP consists of:

- Agents—An agent is software that runs on a device that maintains information about device configuration and current state in a database.
- Managers—An SNMP manager is an application that contacts an SNMP agent to query or modify the agent database.
- The SNMP protocol—SNMP is the application-layer protocol SNMP agents and managers use to send and receive data.
- Management Information Bases (MIB)—The MIB is a text file that specifies the managed objects by an object identifier (OID).



Important

The switch does not reply to SNMP requests sent to the Virtual Router Redundancy Protocol (VRRP) virtual interface address; it does, however, reply to SNMP requests sent to the physical IP address.

An SNMP manager and agent communicate through the SNMP protocol. A manager sends queries and an agent responds; however, an agent initiates traps. Several types of packets transmit between SNMP managers and agents:

- Get request—This message requests the values of one or more objects.
- Get next request—This message requests the value of the next object.
- Set request—This message requests to modify the value of one or more objects.
- Get response—An SNMP agent sends this message in response to a get request, get next request, or set request message.
- Trap—SNMP trap is a notification triggered by events at the agent.

Log Message Format

The log messages for the switch have a standardized format. All system messages are tagged with the following information, except that alarm type and alarm status apply to alarm messages only:

Table 187: Log message tagging

Information type	Description
CPU slot number	Indicates the CP slot where the command is logged.
timestamp	Records the date and time at which the event occurred. The format is MM/DD/YY hh:mm:ss.uuu, where uuu is milliseconds. Example: [11/01/10 11:41:21.376].
hostname	The name of the host from which the message is generated.
event code	Precisely identifies the event reported.

Table 187: Log message tagging (continued)

Information type	Description
alarm code	Specifies the alarm code.
alarm type	Identifies the alarm type (Dynamic or Persistent) for alarm messages.
alarm status	Identifies the alarm status (set or clear) for alarm messages.
VRF name	Identifies the Virtual Routing and Forwarding (VRF) instance, if applicable.
module name	Identifies the software module or hardware from which the log is generated.
severity level	Identifies the severity of the message.
sequence number	Identifies a specific CLI command.
context	Specifies the type of the session used to connect to the switch. If the session is a remote session, the remote IP address is identified.
user name	Specifies the user name used to login to the switch.
CLI command	Specifies the commands typed during the CLI session. The system logs anything type during the CLI session as soon as the user presses the Enter key.

The following messages are examples of an informational message for CLILOG:

```

CP1 [07/18/14 13:23:11.253] 0x002c0600 00000000 GlobalRouter CLILOG INFO 13
TELNET:192.0.2.200 rwa show log file name-of-file log.40300001.1806

CP1 [07/18/14 13:24:19.739] 0x002c0600 00000000 GlobalRouter CLILOG INFO 15
TELNET:192.0.2.200 rwa term more en

CP1 [07/18/14 13:24:22.577] 0x002c0600 00000000 GlobalRouter CLILOG INFO 16
TELNET:192.0.2.200 rwa show log

CP1 [01/12/17 15:13:59.056] 0x002c0600 00000000 GlobalRouter CLILOG INFO 5
TELNET:198.51.100.108 rwa syslog host 4

CP1 [01/12/17 15:13:35.520] 0x002c0600 00000000 GlobalRouter CLILOG INFO 4
TELNET:198.51.100.108 rwa syslog host enable

CP1 [01/12/17 15:13:14.576] 0x002c0600 00000000 GlobalRouter CLILOG INFO 3
TELNET:198.51.100.108 rwa show syslog

CP1 [01/12/17 15:12:44.640] 0x002c0600 00000000 GlobalRouter CLILOG INFO 2
TELNET:198.51.100.108 rwa show logging file tail

```

The following messages are examples of an informational message for SNMPLOG:

```

CP1 [05/07/14 10:24:05.468] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 1 ver=v2c
public rcVlanPortMembers.2 =

CP1 [05/07/14 10:29:58.133] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 2 ver=v2c
public rcVlanPortMembers.2 =

```

```
CPI [05/07/14 10:30:20.466] 0x002c4600 00000000 GlobalRouter SNMPLOG INFO 3 ver=v2c
public rcVlanPortMembers.1 =
```

The system encrypts AP information before writing it to the log file.

The encrypted information in a log file is for debugging purposes. Only a Customer Service engineer can decrypt the encrypted information in a log file. CLI commands display the logs without the encrypted information. Do not edit the log file.

The following table describes the system message severity levels.

Table 188: Severity levels

Severity level	Definition
EMERGENCY (0)	A panic condition that occurs when the system becomes unusable. A severity level of emergency is usually a condition where multiple applications or servers are affected. You must correct a severity level of emergency immediately.
ALERT (1)	Any condition requiring immediate attention and correction. You must correct a severity level of alert immediately, but this level usually indicates failure of a secondary system, such as an Internet Service Provider connection.
CRITICAL (2)	Any critical conditions, such as a hard drive error.
ERROR (3)	A nonfatal condition occurred. You can be required to take appropriate action. For example, the system generates an error message if it is unable to lock onto the semaphore required to initialize the IP addresses used to transfer the log file to a remote host.
WARNING (4)	A nonfatal condition occurred. No immediate action is needed. An indication that an error can occur if action is not taken within a given amount of time.
NOTIFICATION (5)	Significant event of a normal nature. An indication that unusual, but not error, conditions have occurred. No immediate action is required.
INFO (6)	Information only. No action is required.
DEBUG (7)	Message containing information useful for debugging.
FATAL (0)	A fatal condition occurred. The system cannot recover without restarting. For example, a fatal message is generated after the configuration database is corrupted.

Based on the severity code in each message, the platform dispatches each message to one or more of the following destinations:

- workstation display
- local log file
- one or more remote hosts

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration except in the following case. When you configure the system log table in EDM, you must select either IPv4 or IPv6.

Internally, the switch has four severity levels for log messages: INFO, WARNING, ERROR, and FATAL. The system log supports eight different severity levels:

- Debug
- Info
- Notice
- Warning
- Critical
- Error
- Alert
- Emergency

The following table shows the default mapping of internal severity levels to Syslog severity levels.

Table 189: Default and system log severity level mapping

UNIX system error codes	System log severity level	Internal severity level
0	Emergency	Fatal
1	Alert	
2	Critical	
3	Error	Error
4	Warning	Warning
5	Notice	
6	Info	Info
7	Debug	

Log Files

The log file captures hardware and software log messages, and alarm messages. The switch logs to internal flash.

The system saves internal log messages in a circular list in memory, which overwrite older log messages as the log fills. Unlike the log messages in a log file, the internal log messages in memory do not contain encrypted information, which can limit the information available during troubleshooting. Free up the disk space on the flash if the system generates the `disk space 75% full` alarm. After the disk space utilization returns below 75%, the system clears the alarm, and then starts logging to a file again.

Log file naming conventions

The following list provides the naming conventions for the log file:

- The log file is named as log.xxxxxxxx.sss format. The prefix of the log file name is log. The six characters after the log file prefix contain the last three bytes of the chassis base MAC address. The next two characters are 01. The last three characters (sss) denote the sequence number of the log file.
- The sequence number of the log file is incremented for each new log file created after the existing log file reaches the maximum configured size.
- At initial system start up when no log file exists, a new log file with the sequence number 000 is created. After a restart, the system finds the newest log file from internal flash based on file timestamps. If the newest log file is on the flash that is used for logging, the system continues to use the newest log file. And once the maximum configured size is reached, system continues to create a new log file with incremental sequence number on the internal flash for logging.

Log File Transfer

The system logs contain important information for debugging and maintaining the switch. After the current log file reaches the configured maximum size, the system creates a new log file for logging. The system transfers old log files to a remote host. You can configure up to 10 remote hosts, which creates long-term backup storage of your system log files.

Of the 10 configured remote hosts, 1 is the primary host and the other 9 are redundant. Upon initiating a transfer, system messaging attempts to use host 1 first. If host 1 is not reachable, system messaging tries hosts 2 to 10.

If log file transfer is unsuccessful, the system keeps the old log files on internal flash. The system attempts to transfer old log files after the new log file reaches the configured maximum size. The system also attempts to transfer old log files periodically (once in one hundred log writes) if the disk space on the flash is more than 75% full.

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration.

Specify the following information to configure the transfer criteria:

- The maximum size of the log file.



Note

The maximum number of log files is 1,000. The maximum size of one log file is 1,024 KB. You cannot configure the log file parameters on 4220 Series, 5320 Series, or 5420 Series. On these platforms, the number of log files is restricted to 145 and the size of a log file is restricted to 126 KB.

- The IP address of the remote host.
- The name prefix of the log file to store on the remote host.

The system appends a suffix of .xxxxxxx.sss to the file name. The first six characters of the suffix contain the last three bytes of the chassis base MAC address. The next

two characters are 01. The last three characters (sss) denote the sequence number of the log file. For example, if you configure the name prefix as mylog, a possible file name is mylog.90000001.001.

- To use Secure Copy (SCP) for the transfer, configure the user name.

After you configure the first remote host entry for an SCP server, the switch generates an SSH key pair for authentication with all SCP servers. You must manually copy the public key to the remote server. Use the **show logging transferfile {1-10}** command to determine the location of the public key.



Note

With Enhanced Secure Mode enabled, authorized users can use SFTP to transfer files to a remote server with the content encrypted.

- To use File Transfer Protocol (FTP) for the transfer, use the **boot config host** commands to configure the user name and password.

For more information about configuration, see [Configure the Remote Host for Log Transfer](#) on page 2556.

Restrictions

Be aware of the following restrictions to transfer log files to a remote host:

- The remote host IP address must be reachable.
- If you transfer a log file from a host to the system, (for example, to display it with a show command), rename the log file. Failure to rename the log file can cause the system to use the recently transferred file as the current log, if the sequence number in the extension is higher than the current log file. For example, if bf860005.002 is the current log file and you transfer bf860005.007 to the system, the system logs future messages to the bf860005.007 file. You can avoid this if you rename the log file to something other than the format used by system messaging.
- If the TFTP server is a UNIX-based machine, files written to the server must already exist. For example, you must create dummy files with the same names as your system logs. This action is commonly performed by using the touch command (for example, **touch bf860005.001**).
- The internal flash must be less than 75% full for the system to log a file. If the internal flash is more than 75% full, logging to a file stops to prevent exhausting disk space.

Log File Transfer Using a Wild Card File Name

File transfers using SFTP require file permissions. Use the command **attribute WORD<1-99> [+/-] R** to change the permissions of a file.

To change permissions for all log files, use the wild card file name **log.***. Using the command in the wild card form **attribute log.* [+/-]R** changes permissions for log files with names that begin with the characters "log.".



Important

You cannot use a wild card pattern other than **log.*** for this command.

Email Notification

The switch can send email notification for failed components or other critical log-event conditions. The switch can also send periodic status notifications.

Enable and configure a Simple Mail Transfer Protocol (SMTP) client on the switch for one SMTP server by specifying the server hostname or IPv4 address. To use a hostname, you must also configure a Domain Name System (DNS) client on the switch.

You must configure at least one email recipient and can create a maximum of five email recipients.

The switch can periodically send general condition status notifications. Status email messages include information about the following items:

- General switch
- Chassis
- Card
- Temperature
- Power supplies
- Fans
- LEDs
- System errors
- Port lock
- Message control
- Operational configuration changes
- Current Uboot
- Port interfaces
- Port statistics

The switch maintains a default list of event IDs for which it generates an email notification. You can add or remove event IDs. Use the **show smtp event-id** command to display the list of active event IDs. To display information about a particular event, use the **show logging file event code** command.

The following example shows an email that the switch sends for log events.

```
Subject: Logs from LabSwitch - 50712100008
From: <LabSwitch@default.com>
To: <test1@default.com>
CP1 [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1 [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1 [08/04/15 21:48:04.527:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter
SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1 [08/04/15 21:50:03.511:UTC] 0x00088524 00000000 GlobalRouter SW INFO Boot sequence
successful
```

If you enable the SMTP client but the switch cannot reach the SMTP server, the switch generates an alarm. The switch holds log and status information in a queue until the

connection with the SMTP server is restored. The message queue holds a maximum of 2,000 messages. If the queue fills, the switch drops new messages.

The following text is an example of the alarm that the switch generates when it cannot connect to the SMTP server.

```
CP1 [06/10/15 19:27:07.901:EST] 0x00398600 0e600000 DYNAMIC SET GlobalRouter SMTP
WARNING SMTP: Unable to establish connection with server: mailhost.usae.company.com,
port:25
```

If the switch cannot establish a connection to the SMTP server, verify that the server IP address or hostname, and the TCP port are correct. If you specify the server hostname, confirm that the IP address for the DNS server is correct. Check for network issues such as unplugged cables.

If the SMTP server rejects the email message, the switch generates a log message.

Log Configuration Using CLI

Use log files and messages to perform diagnostic and fault management functions.

For information about how to configure or extend system message control, see [Run-Time Process Management Using CLI](#) on page 371.

Configure a UNIX System Log and Syslog Host

Configure the syslog to control a facility in UNIX machines that logs SNMP messages and assigns each message a severity level based on importance.

About This Task

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the system log:
`syslog enable`
3. Configure the maximum number of syslog hosts:
`syslog max-hosts <1-10>`
4. Create the syslog host:
`syslog host <1-10>`
5. Configure the IP address for the syslog host:
`syslog host <1-10> address WORD <0-46>`

6. Enable the syslog host:

```
syslog host <1-10> enable
```

Configure optional syslog host parameters by using the variables in the following variable definition tables.

7. View the configuration to ensure it is correct:

```
show syslog [host <1-10>]
```

Examples

```
Switch:1(config)#syslog enable
Switch:1(config)#syslog host 7 address 192.0.2.1
Switch:1(config)#syslog host 7 enable
Switch:1(config)#show syslog host 7

      Id : 7
      IpAddr : 192.0.2.1
      UdpPort : 514
      Facility : local7
      Severity : info|warning|error|fatal
      MapInfoSeverity : info
      MapWarningSeverity : warning
      MapErrorSeverity : error
      MapMfgSeverity : notice
      MapFatalSeverity : emergency
      Enable : true
      SecureForwardingMode: none
      Tcp Port : 1025
Switch:1(config)#show syslog

Enable      : true
Max Hosts   : 5
OperState   : active
header      : default
Total number of configured hosts : 3
Total number of enabled hosts : 1
Configured host : 7  8  9
Enabled host : 7

TLS-minimum-version      : tlsv11
Ciphers-Tls               : TLS_DHE_RSA_WITH_AES_256_CBC_SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA
                           TLS_RSA_WITH_AES_256_CBC_SHA TLS_RSA_WITH_AES_128_CBC_SHA
```

Variable Definitions

The following table defines parameters for the **syslog** command.

Variable	Value
<i>enable</i>	Enables the sending of syslog messages on the device. The default is enabled.
<i>max-hosts</i> <1-10>	Specifies the maximum number of syslog hosts supported, from 1-10. The default is 5.

The following table defines parameters for the **syslog host** command.

Variable	Value
<i>1-10</i>	Creates and configures a host instance.
<i>address WORD <0-46></i>	Configures a host location for the syslog host. <i>WORD <0-46></i> is the IPv4 or IPv6 address of the UNIX system syslog host in the format A.B.C.D or x:x:x:x:x:x. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration.
<i>enable</i>	Enables the syslog host. The default is disabled.
<i>facility {local0 local1 local2 local3 local4 local5 local6 local7}</i>	Specifies the UNIX facility in messages to the syslog host. <i>{local0 local1 local2 local3 local4 local5 local6 local7}</i> is the UNIX system syslog host facility. The default is <i>local7</i> .
<i>maperror {emergency alert critical error warning notice info debug}</i>	Specifies the syslog severity to use for error messages. The default is <i>error</i> .
<i>mapfatal {emergency alert critical error warning notice info debug}</i>	Specifies the syslog severity to use for fatal messages. The default is <i>emergency</i> .
<i>mapinfo {emergency alert critical error warning notice info debug}</i>	Specifies the syslog severity level to use for information messages. The default is <i>info</i> .
<i>mapwarning {emergency alert critical error warning notice info debug}</i>	Specifies the syslog severity to use for warning messages. The default is <i>warning</i> .
<i>severity <info warning error fatal> [<info warning error fatal>] [<info warning error fatal>] [<info warning error fatal>]</i>	Specifies the severity levels for which to send syslog messages. You can specify up to four severity levels in the same command string. The default is <i>info</i> .
<i>udp-port <514-530></i>	Specifies the User Datagram Protocol port number on which to send syslog messages to the syslog host. This value is the UNIX system syslog host port number from 514–530. The default is 514.

Configure Secure Forwarding to the Syslog Server

Configuring secure forwarding includes setting the mode for the particular syslog host and setting the TCP port through which the logs are sent to the syslog server.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create the syslog host:

```
syslog host <1-10>
```

3. Configure an IP address for the syslog host:

```
syslog host <1-10> address WORD<0-46>
```

4. Enable the syslog host:

```
syslog host <1-10> enable
```

5. Enable syslog globally:

```
syslog enable
```

6. Set the mode for secure forwarding on the host:

```
syslog host <1-10> secure-forwarding mode <none | tls [server-cert-name WORD<1-64>]>
```

7. Set the TCP port:

```
syslog host <1-10> secure-forwarding tcp-port <1025-49151>
```

8. Display the secure forwarding configured values:

```
show syslog host <1-10>
```

What to Do Next

After configuring secure forwarding on the switch, set the syslog server to be able to see the log messages on the interactive syslog viewer.

For TLS secure syslog, on the rsyslog server, configure the server to use TLS method and install the root certificate on the server in the switch.

Variable Definitions

The following table defines parameters for the **syslog host** command.

Variable	Value
host <1-10>	Specifies the ID for the syslog host. The range is 1-10.
address WORD<0-46>	Configures a host location for the syslog host. WORD <0-46> is the IPv4 or IPv6 address of the UNIX system syslog host in the format A.B.C.D or x:x:x:x:x:x:x. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration using CLI.
enable	Enables the syslog host. The default is disabled.
secure-forwarding	Adds protected syslog using remote port forwarding for host.

The following table defines parameters for the **syslog host secure-forwarding** command.

Variable	Value
<code>host <1-10></code>	Creates and configures a host instance.
<code>mode <none tls [server-cert-name WORD<1-64>]></code>	Specifies the mode of secure forwarding of syslog on the host. The default mode is none, that is, tls mode is disabled by default. Note: Certificate validation is done only if the server-cert-name is configured.
<code>tcp-port <1025-49151></code>	Set tcp-port for secure forwarding of syslog for host. The default tcp-port is 1025. To set the TCP port to default value, use command default syslog host <1-10> secure-forwarding tcp-port . Important: The tcp-port 6000 cannot be used, as it is used as an internal port for Internal Spanning Tree (IST).

Configure Certificate Authority Trustpoint for Syslog

Before You Begin

Configure and install the digital certificate. For digital certificate information, see [Digital Certificate Configuration](#) on page 3399.

About This Task

Use this procedure to configure the certificate authority (CA) for the syslog server. This command adds a root or intermediate CA certificate to the syslog certificate store.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Configure the CA trustpoint:
syslog certificate ca-common-name WORD<1-45>
3. Save the configuration.
4. (Optional) Display the installed syslog server CA trustpoint:
show syslog certificate
5. (Optional) Display CA details:
show certificate ca [WORD<1-45>]

The UsedFor field displays SYSLOG.

Variable Definitions

The following table defines parameters for the **syslog certificate ca-common-name** and **show certificate ca** commands.

Variable	Value
<i>WORD</i> <1-45>	Specifies the Certificate Authority common name.

Configure the Minimum Version of TLS for Syslog

About This Task

Use this procedure to configure the minimum version of TLS protocol supported by the syslog client.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Configure the TLS protocol version for the syslog client:

```
syslog tls-min-verion {tlsv11 | tlsv12}
```

Variable Definitions

The following table defines parameters for the **syslog tls-min-verion** command.

Variable	Value
<i>tlsv11</i> <i>tlsv12</i>	<p>Specifies the minimum version of the TLS protocol supported by the syslog client.</p> <ul style="list-style-type: none">• <i>tlsv11</i> - configures version 1.1• <i>tlsv12</i> - configures version 1.2 <p>If the switch operates in Enhanced Secure Mode, the default TLS minimum version for Syslog is <i>tlsv12</i>. Otherwise, the default is <i>tlsv11</i>.</p>

Configure Logging

Configure logging to determine the types of messages to log and where to store the messages. For information about how to transfer a log file to a remote host, see [Configure the Remote Host for Log Transfer](#) on page 2556.

About This Task



Note

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs CLI Log and SNMP Log information regardless of the logging level you configure. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Define which messages to log:

```
logging level <0-4>
```
3. Write the log file from memory to a file:

```
logging write WORD<1-1536>
```
4. Show logging on the screen:

```
logging screen
```

Example

```
Switch:1(config)#logging level 0
Switch:1(config)#logging write log2
Switch:1(config)#logging screen
```

Variable Definitions

The following table defines parameters for the **logging** command.

Variable	Value
<code>level <0-4></code>	Shows and configures the logging level. The level is one of the following values: <ul style="list-style-type: none"> • 0: Information — all messages are recorded • 1: Warning — only warning and more serious messages are recorded • 2: Error — only error and more serious messages are recorded • 3: Manufacturing — this parameter is not available for customer use • 4: Fatal — only fatal messages are recorded
<code>screen</code>	Enables the log display on the screen.
<code>write WORD<1-1536></code>	Writes the log file with the designated string. <code>WORD<1-1536></code> is the string or command that you append to the log file. If the string contains spaces, you must enclose the string in quotation marks (").

Enable System Logging

System logs are a valuable diagnostic tool. You can send log messages to flash files for later retrieval.

About This Task

Configure logging to a flash file at all times as a best practice.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable system logging:

```
boot config flags logging
```

Variable Definitions

The following table defines parameters for the **boot config** command.

Variable	Value
<i>flags logging</i>	Enables or disables logging to a flash file. The log file is named using the format log.xxxxxxx.sss. The first six characters after the prefix of the file name log contain the last three bytes of the chassis base MAC address. The next two characters specify the slot number. The last three characters denote the sequence number of the log file.
<i>logfile</i> <64-500> <500-16384> <10-90>	<p>Configures the following log file parameters:</p> <ul style="list-style-type: none"> • <64-500> specifies the minimum free memory space on /intflash from 64-500 KB. Although you can currently use this parameter, it does not alter the configuration. • <500-16384> specifies the maximum size of the log file from 500-16384 KB. • <10-90> specifies the maximum percentage, ranging from 10-90 percent, of space on /intflash the logfile can use. Although you can currently use this parameter, it does not alter the configuration. <p>Note: The maximum number of log files is 1,000. The maximum size of one log file is 1,024 KB. You cannot configure the log file parameters on 4220 Series, 5320 Series, or 5420 Series. On these platforms, the number of log files is restricted to 145 and the size of a log file is restricted to 126 KB.</p>

Configure the Remote Host for Log Transfer

Configure the remote host for log transfer. The system transfers the current log file to a remote host after the log file size reaches the maximum size or the local storage capacity is reached.

Before You Begin

- The IP address you configure for the remote host must be reachable at the time of configuration.

Procedure

- Enter Global Configuration mode:

```
enable

configure terminal
```
- (Optional) For FTP, configure the host user name and password:

```
boot config host user WORD<0-16>

boot config host password WORD<0-16>
```
- Configure the remote host address for log transfer:
 - To use TFTP or FTP: `logging transferFile {1-10} address {A.B.C.D} [filename-prefix WORD<0-200>]`
 - To use SCP: `logging transferFile {1-10} scp username WORD<0-50> [address {A.B.C.D}] [filename-prefix WORD<0-200>]`



Note

As a best practice for SCP transfer, first configure parameters other than the remote server address so the switch can generate the SSH key pair, and then configure the remote server address after you upload the public key to the remote server.

- Verify the configuration:

```
show logging transferfile {1-10}
```

Example

Configure the remote host for log transfer using SCP:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#logging transferfile 1 scp username scp_user filename-prefix scp
Generating ssh key-pair ...
SSH key pair generated. Public key at location /intflash/shared/my_key.pub.
#Upload the public key to the destination host.
Switch:1(config)#logging transferfile 1 address 192.0.2.10
```

Verify the configuration:

```
Switch:1>show logging transferfile 1

RemoteIPAddress : 192.0.2.10
FileName : scp.02d80001.131
```

```

FileNamePrefix : scp
Username       : scp_user
TransferType   : SCP
PublicKey      : /intflash/shared/my_key.pub

```

Variable Definitions

The following table defines parameters for the **logging transferFile** command.

Variable	Value
<i>1-10</i>	Specifies the host ID. You can configure a maximum of 10 remote hosts.
<i>address {A.B.C.D}</i>	Specifies the IP address of the host to which to transfer the log file. The remote host must be reachable or the configuration fails.
<i>filename-prefix WORD<0-200></i>	Specifies the prefix for the filename on the remote host. If you do not configure a value, the current log filename is the default.
<i>scp username WORD<0-50></i>	Specifies the SCP user name for the remote server.

View Logs

View log files by file name, category, or severity to identify possible problems.

About This Task

View CLI command and SNMP trap logs, which are logged as normal log messages and logged to the system log file.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Show log information:
show logging file [alarm] [CPU WORD<0-100>] [detail] [event-code WORD<0-10>] [module WORD<0-100>] [name-of-file WORD<1-99>] [save-to-file WORD<1-99>] [severity WORD<0-25>] [tail] [vrf WORD<0-32>]

Example

Display log file information:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#show logging file
CP1  [02/06/15 22:38:20.678:UTC] 0x00270428 00000000 GlobalRouter SW INFO Lifecy
cle: Start
CP1  [02/06/15 22:38:21.770:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s sockserv started, pid:4794
CP1  [02/06/15 22:38:21.771:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom95 started, pid:4795
CP1  [02/06/15 22:38:21.771:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces
s oom90 started, pid:4796
CP1  [02/06/15 22:38:21.772:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Proces

```

```

s imgsync.x started, pid:4797
CP1 [02/06/15 22:38:22.231:UTC] 0x0026452f 00000000 GlobalRouter SW INFO No patch set.
CP1 [02/06/15 22:38:22.773:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process logServer started, pid:4840
CP1 [02/06/15 22:38:22.774:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process trcServer started, pid:4841
CP1 [02/06/15 22:38:22.774:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process oobServer started, pid:4842
CP1 [02/06/15 22:38:22.775:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process cbcp-main.x started, pid:4843
CP1 [02/06/15 22:38:22.776:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process rssServer started, pid:4844
CP1 [02/06/15 22:38:22.777:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process dbgServer started, pid:4845
CP1 [02/06/15 22:38:22.777:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process dbgShell started, pid:4846
CP1 [02/06/15 22:38:22.778:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process coreManager.x started, pid:4847
CP1 [02/06/15 22:38:22.779:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process ssio started, pid:4848
CP1 [02/06/15 22:38:22.780:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process hckServer started, pid:4849
CP1 [02/06/15 22:38:22.780:UTC] 0x0027042b 00000000 GlobalRouter SW INFO Process remCmdAgent.x started, pid:4850
CP1 [02/06/15 22:38:24.717:UTC] 0x000006cc 00000000 GlobalRouter SW INFO rcStart: FIPS Power Up Self Test SUCCESSFUL - 0
CP1 [02/06/15 22:38:24.718:UTC] 0x000006c2 00000000 GlobalRouter SW INFO rcStart: Security Stack Init SUCCESSFUL - 0
CP1 [02/06/15 22:38:24.718:UTC] 0x000006c3 00000000 GlobalRouter SW INFO rcStart: IPSEC Init SUCCESSFUL
CP1 [02/06/15 22:38:24.718:UTC] 0x000006bf 00000000 GlobalRouter SW INFO rcStart: Security Stack Log init SUCCESSFUL - 0
IO1 [02/06/15 22:38:26.960:UTC] 0x0011054a 00000000 GlobalRouter COP-SW INFO Detected Master CP in slot 1

--More-- (q = quit)

Switch:1(config)#show logging file module SNMP
CP1 [02/06/15 22:39:58.530:UTC] 0x00004595 00000000 GlobalRouter SNMP INFO Booted with file
CP1 [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1 [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1 [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR GlobalRouter SNMP INFO 2k card up(CardNum=3 AdminStatus=1 OperStatus=1)
CP1 [02/06/15 22:40:45.839:UTC] 0x000045e5 00400005 DYNAMIC SET GlobalRouter SNMP INFO Sending Cold-Start Trap

```

Variable Definitions

The following table defines parameters for the **show logging file** command.

Variable	Value
<i>alarm</i>	Displays alarm log entries.
<i>CPU WORD <0-100></i>	Filters and lists the logs according to the CPU that generated the message. Specify a string length of 0-25 characters. To specify multiple filters, separate each CPU by the vertical bar (), for example, CPU1 CPU2.

Variable	Value
<i>detail</i>	Displays CLI and SNMP logging information.
<i>event-code</i> <i>WORD</i> <0-10>	Specifies a number that precisely identifies the event reported.
<i>module</i> <i>WORD</i> <0-100>	Filters and lists the logs according to module. Specifies a string length of 0-100 characters. Categories include SNMP, EAP, RADIUS, RMON, WEB, HW, MLT, FILTER, QOS, CLILOG, SW, CPU, IP, VLAN, IPMC, and SNMPLOG. To specify multiple filters, separate each category by the vertical bar (), for example, <i>FILTER</i> <i>QOS</i> .
<i>name-of-file</i> <i>WORD</i> <1-99>	Displays the valid logs from this file. For example, /intflash/logcopy.txt. You cannot use this command on the current log file, the file into which the messages are currently logged. Specify a string length of 1 to 99 characters. If you enable enhanced secure mode, the system encrypts the entire log file. After you use the show log file name-of-file WORD<1-99> command, the system takes the encrypted log file name as input, then decrypts it, and prints the output to the screen. You can then redirect the decrypted output to a file that you can store onto the flash. If enhanced secure mode is disabled, the system only encrypts the proprietary portion of the log file.
<i>save-to-file</i> <i>WORD</i> <1-99>	Redirects the output to the specified file and removes all encrypted information. You cannot use the <i>tail</i> option with the <i>save-to-file</i> option. Specify a string length of 1-99 characters.
<i>severity</i> <i>WORD</i> <0-25>	Filters and lists the logs according to severity. Choices include INFO, ERROR, WARNING, and FATAL. To specify multiple filters, separate each severity by the vertical bar (), for example, <i>ERROR</i> <i>WARNING</i> <i>FATAL</i> .
<i>tail</i>	Shows the last results first.
<i>vrf</i> <i>WORD</i> <0-32>	Specifies the name of a VRF instance to show log messages that only pertain to that VRF.

Configure CLI Logging

Use CLI logging to track configuration changes made using CLI commands on the switch. The CLI commands are logged to the system log file as CLILOG module.

About This Task



Note

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs CLI Log and SNMP Log information regardless of the logging level you configure. This is not the case for other INFO messages.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Enable CLI logging:
clilog enable
3. Ensure that the configuration is correct:
show clilog
4. View the CLI log:
show logging file module clilog

Example

Enable CLI logging, and view the CLI log:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#clilog enable
Switch:1(config)#show logging file module clilog
CP1 [02/13/13 17:29:30.628] 0x002c0600 00000000 GlobalRouter CLILOG INFO 4 CONSOLE
rwa show logging file modug
CP1 [02/14/13 19:39:11.648] 0x002c0600 00000000 GlobalRouter CLILOG INFO 5 CONSOLE
rwa ena
CP1 [02/14/13 19:39:13.420] 0x002c0600 00000000 GlobalRouter CLILOG INFO 6 CONSOLE
rwa conf t
CP1 [02/14/13 19:49:21.044] 0x002c0600 00000000 GlobalRouter CLILOG INFO 7 CONSOLE
rwa filter acl 2 enable
CP1 [02/14/13 19:50:08.540] 0x002c0600 00000000 GlobalRouter CLILOG INFO 8 CONSOLE
rwa filter acl 2 type inpol
CP1 [02/14/13 19:50:38.444] 0x002c0600 00000000 GlobalRouter CLILOG INFO 9 CONSOLE
rwa filter acl 2 type inpoe
CP1 [02/14/13 19:50:52.968] 0x002c0600 00000000 GlobalRouter CLILOG INFO 10 CONSOLE
rwa filter acl enable 2
CP1 [02/14/13 19:51:08.908] 0x002c0600 00000000 GlobalRouter CLILOG INFO 11 CONSOLE
rwa filter acl 2 enable
CP1 [02/15/13 06:50:25.972] 0x002c0600 00000000 GlobalRouter CLILOG INFO 14 CONSOLE
rwa ena
CP1 [02/15/13 06:50:30.288] 0x002c0600 00000000 GlobalRouter CLILOG INFO 15 CONSOLE
rwa conf t
CP1 [02/15/13 06:50:39.412] 0x002c0600 00000000 GlobalRouter CLILOG INFO 16 CONSOLE
rwa show vlan basic
CP1 [02/15/13 06:51:09.488] 0x002c0600 00000000 GlobalRouter CLILOG INFO 17 CONSOLE
rwa show isis spbm
CP1 [02/15/13 06:56:00.992] 0x002c0600 00000000 GlobalRouter CLILOG INFO 19 CONSOLE
rwa spbm 23 b-vid 2 primar1
CP1 [02/15/13 06:56:59.092] 0x002c0600 00000000 GlobalRouter CLILOG INFO 20 CONSOLE
rwa show isis
CP1 [02/15/13 07:10:54.928] 0x002c0600 00000000 GlobalRouter CLILOG INFO 21 CONSOLE
rwa show isis interface
CP1 [02/15/13 07:12:33.404] 0x002c0600 00000000 GlobalRouter CLILOG INFO 22 CONSOLE
rwa show isis spbm
CP1 [02/15/13 07:45:28.596] 0x002c0600 00000000 GlobalRouter CLILOG INFO 23 CONSOLE
rwa ena
CP1 [02/15/13 07:45:30.236] 0x002c0600 00000000 GlobalRouter CLILOG INFO 24 CONSOLE
rwa conf t
CP1 [02/15/13 07:46:29.456] 0x002c0600 00000000 GlobalRouter CLILOG INFO 25 CONSOLE
rwa interface gigabitEther0
CP1 [02/15/13 07:47:28.476] 0x002c0600 00000000 GlobalRouter CLILOG INFO 26 CONSOLE
```

```
rra encapsulation dot1q
--More-- (q = quit)
```

Variable Definitions

The following table defines parameters for the **cliilog** command.

Variable	Value
<i>enable</i>	Activates CLI logging. To disable, use the no cliilog enable command.

Configure Email Notifications for System Events

Configure the SMTP feature to generate email notifications for component failures, critical conditions, or general system status.

Before You Begin

To identify the SMTP server by hostname, you must first configure a DNS client on the switch.

About This Task

The SMTP feature is disabled by default.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the TCP port the client uses to open a connection with the SMTP server:

```
smtp port <1-65535>
```



Note

The port you specify must match the port that the SMTP server uses.

3. Configure email recipients:

```
smtp receiver-email add WORD<3-1274>
```

```
smtp receiver-email remove WORD<3-1274>
```



Note

You must configure at least one recipient.

4. Configure the SMTP server hostname or IPv4 address:

```
smtp server WORD<1-256>
```

5. (Optional) Configure a sender email address:

```
smtp sender-email WORD<3-254>
```

6. (Optional) Add or remove log events to the default list that generate email notification:

```
smtp event-id add WORD<1-1100>
```

```
smtp event-id remove WORD<1-1100>
```

7. (Optional) Configure the status update interval:

```
smtp status-send-timer <0 | 30-43200>
```

8. Enable the SMTP client:

```
smtp enable
```

9. Configure an SMTP domain name:

```
smtp domain-name WORD<1-254>
```

10. Verify the configuration:

```
show smtp [event-id]
```

11. (Optional) Display information about the specific event:

```
show logging file event-code WORD<1-10>
```

Examples

Configure the SMTP client to use TCP port 26 to communicate with an SMTP server that is using port 26. Add two receiver email addresses, configure the server information using an IPv4 address, and enable the SMTP feature. Finally, configure an SMTP domain name, and then verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#smtp port 26
Switch:1(config)#smtp receiver-email add test1@default.com,test2@default.com
Switch:1(config)#smtp server 192.0.2.1
Switch:1(config)#smtp enable
Switch:1(config)#smtp domain-name test mailer
Switch:1(config)#show smtp
=====
SMTP Information
=====
SMTP Status: Enabled
Server Address: 192.0.2.1
Server Port: 26
Status send Timer: 30 (seconds)
Sender Email: LabSwitch@default.com
Domain Name: test mailer
Receiver Emails: test1@default.com
                  test2@default.com
```

Add an event ID to the list for which the switch sends email notification on a log event.

```
Switch:1(config)#smtp event-id add 0x0000c5ec
```

Verify the configuration.

```
Switch>show smtp event-id
=====
SMTP Event IDs Information
=====
Log Event IDs: (total: 50)
                ,0x000106d0
```



```

,0x000106da
,0x000107f5
,0x0000c5ec
,0x00040506
,0x00088524
,0x000e4602
,0x000e4608
,0x00210596
,0x0027458d

Default Event IDs: (total: 50)
,0x000106d0
,0x000106da
,0x000107f5
,0x00040506
,0x00088524
,0x000e4602
,0x000e4608
,0x00210596
,0x0027458d

Remove From Default: (total: 0)

Add List: (total: 1)
0x0000c5ec
```

Display information about a specific event:

```
Switch:1#show logging file event-code 0x00088524
1 2024-03-07T13:58:49.449Z BN1-7720-4BN-3 CP1 - 0x00088524 - 00000000 GlobalRouter SW
INFO Boot sequence successful
```

Variable Definitions

The following table defines parameters for the **smtp port** command.

Variable	Value
<1-65535>	<p>Specifies the TCP port on the switch that the SMTP client uses to communicate with the SMTP server. The default value is 25.</p> <p>Note: You must disable the SMTP feature before you can change an existing SMTP port configuration. The port you specify must match the port that the SMTP server uses.</p>

The following table defines parameters for the **smtp receiver-email** command.

Variable	Value
<i>add</i> <i>WORD<3-1274></i>	<p>Adds an email address to the recipient list. The recipients receive the email notification generated by the switch.</p> <p>You must configure at least one email recipient and can create a maximum of five email recipients. You can specify multiple addresses in a single command by separating them with a comma.</p> <p>You cannot use quotation marks (") or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC 5321.</p> <p>The maximum length for the address is 254 characters.</p>
<i>remove</i> <i>WORD<3-1274></i>	<p>Removes an email address from the recipient list. The recipients receive the email notification generated by the switch. You can specify multiple addresses in a single command by separating them with a comma.</p> <p>You cannot use quotation marks (") or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC 5321.</p> <p>The maximum length for the address is 254 characters.</p>

The following table defines parameters for the **smtp server** command.

Variable	Value
<i>WORD<1-256></i>	<p>Specifies the SMTP server address. You can use either a hostname or IPv4 address. If you use a hostname, you must configure the DNS client on the switch.</p>

The following table defines parameters for the **smtp sender-email** command.

Variable	Value
<i>WORD<3-254></i>	<p>Specifies the email address that the system displays it in the From field of the message that the switch generates. By default, the switch uses <i><SystemName>@default.com</i>.</p>

The following table defines parameters for the **smtp event-id** command.

Variable	Value
<i>add WORD<1-1100></i>	<p>Adds a log event to the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma.</p> <p>The event ID can be up to 10 digits in hexadecimal format.</p>
<i>remove</i> <i>WORD<1-1100></i>	<p>Removes a log event from the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma.</p> <p>The event ID can be up to 10 digits in hexadecimal format.</p>

The following table defines parameters for the **smtp status-send-timer** command.

Variable	Value
<0 30-43200>	Specifies the interval, in seconds, at which the switch sends status information. The default is 30 seconds. A value of 0 means the switch does not send status information.

The following table defines parameters for the **smtp domain-name** command.

Variable	Value
WORD<1-254>	Specifies the SMTP host name or IPv4 address (string length 1-254).

The following table defines parameters for the **show smtp** command.

Variable	Value
event-id	Shows a list of active event IDs for which the switch generates email notification. The command output includes the default list of IDs and IDs you specifically add or remove.

Log Configuration Using EDM

Use log files and messages to perform diagnostic and fault management functions. This section provides procedures to configure and use the logging system in Enterprise Device Manager (EDM).

Configure the System Log

About This Task

Configure the system log to track all user activity on the device. The system log can send messages of up to ten syslog hosts.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **System Log**.
3. In the **System Log** tab, select **Enable**.
4. Configure the maximum number of syslog hosts.
5. Configure the minimum version of TLS protocol.
6. Select **Apply**.

System Log Field Descriptions

Use the data in the following table to use the **System Log** tab.

Name	Description
Enable	Enables or disables the syslog feature. If you select this variable, this feature sends a message to a server on a network that is configured to receive and store diagnostic messages from this device. You can configure the type of messages sent. The default is enabled.
MaxHosts	Specifies the maximum number of remote hosts considered active and can receive messages from the syslog service. The range is 0–10 and the default is 5.
OperState	Specifies the operational state of the syslog service. The default is active.
TlsMinimumVersion	<p>Specifies the minimum version of TLS protocol supported by the syslog client.</p> <ul style="list-style-type: none">• <code>tlsv11</code> - configures TLS version 1.1• <code>tlsv12</code> - configures TLS version 1.2 <p>If the switch operates in Enhanced Secure Mode, the default TLS minimum version for Syslog is <code>tlsv12</code>. Otherwise, the default is <code>tlsv11</code>.</p>
EncryptionType	Specifies the ciphers for preset version of TLS for the syslog.

Configure the System Log Table

About This Task

Use the system log table to customize the mappings between the severity levels and the type of alarms.

You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses with no difference in functionality or configuration except in the following case. When you configure the system log table, under the **System Log Table** tab, you must select **ipv4** or **ipv6**, in **AddressType**. **Address** supports both IPv4 and IPv6 addresses.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **System Log**.
3. Select the **System Log Table** tab.
4. Select **Insert**.
5. Configure the parameters as required.
6. Select **Insert**.
7. To modify mappings, double-click a parameter to view a list of options.

8. Select **Apply**.

System Log Table *Field Descriptions*

Use the data in the following table to use the **System Log Table** tab.

Name	Description
Id	Specifies the ID for the syslog host. The range is 1–10.
AddressType	Specifies if the address is an IPv4 or IPv6 address.
Address	Specifies the IP address of the syslog host. You can log system log messages to external system log hosts with both IPv4 and IPv6 addresses.
UdpPort	Specifies the UDP port to use to send messages to the syslog host (514–530). The default is 514.
Enable	Enables or disables the sending of messages to the syslog host. The default is disabled.
HostFacility	Specifies the syslog host facility used to identify messages (local0 to local7). The default is local7.
Severity	Specifies the message severity for which syslog messages are sent. The default is info warning error fatal.
MapInfoSeverity	Specifies the syslog severity to use for INFO messages. The default is info.
MapWarningSeverity	Specifies the syslog severity to use for WARNING messages. The default is warning.
MapErrorSeverity	Specifies the syslog severity to use for ERROR messages. The default is error.
MapFatalSeverity	Specifies the syslog severity to use for FATAL messages. The default is emergency.
MapMfgSeverity	Specifies the syslog severity to use for Accelar manufacturing messages. The default is notice.
SecureForwardingTcpPort	Specifies the TCP port to use for secure forwarding for a particular host. The default is 1025.
SecureForwardingMode	Enables or disables secure forwarding of syslog over remote port forwarding. The supported values are tls and none. The default is none, which means that secure forwarding is disabled.
SecureForwardingServerCertificateName	Specifies the server certificate name. Certificate validation is done only if the server certificate name is configured.

Configure the Syslog Certificate Authority Common Name

Use the following procedure to add the common name of the Certificate Authority for the syslog server.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **System Log**.
3. Select the **Certificate Table** tab.
4. Select **Insert**.
5. In **CaCommonName**, type the common name of the Certificate Authority.
6. Select **Insert**.

Certificate Table Field Descriptions

Use the data in the following table to use the **Certificate Table** tab.

Name	Description
CaCommonName	Specifies the Certificate Authority common name for a syslog server.

Configure Email Notification for System Events

Configure the SMTP feature to generate email notifications for component failures, critical conditions, or general system status.

Before You Begin

To identify the SMTP server by hostname, you must first configure a DNS client on the switch.

About This Task

The SMTP feature is disabled by default.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **SMTP**.
3. Select the **Globals** tab.
4. In the **ServerAddress** field, configure the SMTP server address.
5. In the **ReceiverEmailsList** field, add email recipients.



Note

You must configure at least one recipient.

6. (Optional) In the **SenderEmail** field, configure a sender email address to use an address other than the default.
7. In the **DomainName** field, configure an SMTP domain name.
8. In the **Port** field, configure the TCP port that the client uses to open a connection with the SMTP server.
9. (Optional) In the **SystemStatusSendTimer** field, configure the status update interval.
10. Select **enable** to enable the SMTP client.

11. (Optional) In the **LogEventIds** field, add or remove log events to the default list that generates an email notification.
12. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
ServerAddressType	Specifies the type of server address as either an IPv4 address or a hostname. If you use a hostname, you must configure the DNS client on the switch.
ServerAddress	Specifies the SMTP server address. You can use either a hostname or an IPv4 address. If you use a hostname, you must configure the DNS client on the switch.
ReceiverEmailsList	<p>Specifies the recipient list. The recipients receive the email notification generated by the switch.</p> <p>You must configure at least one email recipient and can create a maximum of five email recipients. You can specify multiple addresses in a single command by separating them with a comma.</p> <p>You cannot use quotation marks (") or commas (,) in email addresses. Other restrictions for the format of the email address follow RFC5321.</p> <p>The maximum length for the address is 254 characters.</p>
NumOfEmails	Shows the total number of addresses in ReceiverEmailsList .
SenderEmail	Specifies the email address that the system displays it in the From field of the message that the switch generates. By default, the switch uses <i>SystemName@default.com</i> .
DomainName	Specifies the SMTP domain name. The maximum length is 254 characters.

Name	Description
Port	Specifies the TCP port on the switch that the SMTP client uses to communicate with the SMTP server. The default value is 25. Note: You must disable the SMTP feature before you can change an existing SMTP port configuration. The port you specify must match the port that the SMTP server uses.
SystemStatusSendTimer	Specifies the interval, in seconds, at which the switch sends status information. The default is 30 seconds. A value of 0 means the switch does not send status information.
Enable	Enables or disables the SMTP feature. By default, SMTP is disabled.
LogEventIds	Specifies the list of events that generate email notification. You can specify multiple event IDs in a single command by separating them with a comma. The event ID can be up to 10 digits in hexadecimal format.
NumOfEventIds	Shows the total number of IDs in LogEventIds .
DefaultLogEventIds	Shows the default list of event IDs that generate email notification.
NumOfDefaultEventIds	Shows the total number of IDs in DefaultLogEventIds .

SNMP Trap Configuration Using CLI

This section provides procedures to configure SNMP traps using the command line interface (CLI).

Configure an SNMP Host

Configure an SNMP host so that the system can forward SNMP traps to a host for monitoring. You can use SNMPv1, SNMPv2c, or SNMPv3. You configure the target table parameters (security name and model) as part of the host configuration.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```


2. Configure an SNMPv1 host:

```
snmp-server host WORD<1-256> [port <1-65535>] v1 WORD<1-32> [filter WORD<1-32>]
```

3. Configure an SNMPv2c host:

```
snmp-server host WORD<1-256> [port <1-65535>] v2c WORD<1-32> [inform [timeout <1-2147483647>] [retries <0-255>] [mms <0-2147483647>]] [filter WORD<1-32>]
```

4. Configure an SNMPv3 host:

```
snmp-server host WORD<1-256> [port <1-65535>] v3 {noAuthNoPriv|authNoPriv|AuthPriv} WORD<1-32> [inform [timeout <1-2147483647>] [retries <0-255>]] [filter WORD<1-32>]
```

5. Ensure that the configuration is correct:

```
show snmp-server host
```

Example

Configure the target table entry. Configure an SNMPv3 host.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server host 192.0.2.207 port 162 v2c ReadView inform timeout 1500
retries 3 mms 484
Switch:1(config)#snmp-server host 192.0.2.207 port 163 v3 authPriv Lab3 inform timeout
1500 retries 3
```

Variable Definitions

The following table defines parameters for the **snmp-server host** command.

Variable	Value
<i>inform [timeout <1-2147483647>] [retries <0-255>] [mms <0-2147483647>]</i>	Sends SNMP notifications as inform (rather than trap). To use all three options in one command, you must use them in the following order: <ol style="list-style-type: none"> <i>timeout <1-2147483647></i> specifies the timeout value in seconds with a range of 1-214748364. <i>retries <0-255></i> specifies the retry count value with a range of 0-255. <i>mms <0-2147483647></i> specifies the maximum message size as an integer with a range of 0-2147483647.
<i>filter WORD<1-32></i>	Specifies the filter profile to use.
<i>noAuthNoPriv authNoPriv AuthPriv</i>	Specifies the security level.
<i>port <1-65535></i>	Specifies the host server port number.
<i>WORD<1-32></i>	Specifies the security name, which identifies the principal that generates SNMP messages.
<i>WORD<1-256></i>	Specifies either an IPv4 or IPv6 address.

Configure an SNMP Notify Filter Table

Configure the notify table to select management targets to receive notifications, as well as the type of notification to send to each management target.

Before You Begin

- For more information about the notify filter table, see RFC3413.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Create a new notify filter table:

```
snmp-server notify-filter WORD<1-32> WORD<1-32>
```
- Ensure that the configuration is correct:

```
show snmp-server notify-filter
```

Example

```
Switch:1(config)#snmp-server notify-filter profile3 99.3.6.1.6.3.1.1.4.1
Switch:1(config)#show snmp-server notify-filter
```

```
=====
```

Notify Filter Configuration		
Profile Name	Subtree	Mask
profile1	+99.3.6.1.6.3.1.1.4.1	0x7f
profile2	+99.3.6.1.6.3.1.1.4.1	0x7f
profile3	+99.3.6.1.6.3.1.1.4.1	0x7f

Variable Definitions

The following table defines parameters for the **snmp-server notify-filter** command.

Variable	Value
<i>WORD<1-32></i> <i>WORD<1-32></i>	<p>Creates a notify filter table.</p> <p>The first instance of <i>WORD<1-32></i> specifies the name of the filter profile with a string length of 1–32.</p> <p>The second instance of <i>WORD<1-32></i> identifies the filter subtree OID with a string length of 1–32.</p> <p>If the subtree OID parameter uses a plus sign (+) prefix (or no prefix), this indicates include. If the subtree OID uses the minus sign (–) prefix, it indicates exclude.</p> <p>You do not calculate the mask because it is automatically calculated. You can use the wildcard character, the asterisk (*), to specify the mask within the OID. You do not need to specify the OID in the dotted decimal format; you can alternatively specify that the MIB parameter names and the OIDs are automatically calculated.</p>

Enable SNMP Trap Logging

Use SNMP trap logging to send a copy of all traps to the syslog server.

Before You Begin

- You must configure and enable the syslog server.

About This Task



Note

The platform logs CLILOG and SNMPLOG as INFO. Normally, if you configure the logging level to WARNING, the system skips all INFO messages. However, if you enable CLILOG and SNMPLOG the system logs CLI Log and SNMP Log information regardless of the logging level you set. This is not the case for other INFO messages.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Enable SNMP trap logging:

```
snmplog enable
```
- (Optional) Disable SNMP trap logging:

```
no snmplog enable
```

4. View the contents of the SNMP log:

```
show logging file module snmplog
```

Example

Enable SNMP trap logging and view the contents of the SNMP log:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmplog enable
Switch:1(config-app)#show logging file module snmp
CP1  [02/06/15 22:39:58.530:UTC] 0x00004595 00000000 GlobalRouter SNMP INFO Boot
ed with file
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=1 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=2 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:39:59.547:UTC] 0x00004603 00400003.67108870 DYNAMIC CLEAR Glob
alRouter SNMP INFO 2k card up(CardNum=3 AdminStatus=1 OperStatus=1)
CP1  [02/06/15 22:40:45.839:UTC] 0x000045e5 00400005 DYNAMIC SET GlobalRouter SN
MP INFO Sending Cold-Start Trap
```

Variable Definitions

The following table defines parameters for the **snmplog** command.

Variable	Value
<i>enable</i>	Enables the logging of traps. Use the command <code>no snmplog enable</code> to disable the logging of traps.

SNMP Trap Configuration Using EDM

This section provides procedures to configure SNMP traps using Enterprise Device Manager (EDM).

Configure an SNMP Host Target Address

Configure a target table to specify the list of transport addresses to use in the generation of SNMP messages.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **Target Table**.
3. Select the **Target Table** tab.
4. Select **Insert**.
5. In **Name**, type a unique identifier.
6. In **TDomain**, select the transport type of the address, either **ipv4Tdomain** or **ipv6Tdomain**.
7. In **TAddress**, type the transport address and User Datagram Protocol (UDP) port.
8. In **Timeout**, type the maximum round trip time.

9. In **RetryCount**, type the number of retries to be attempted.
10. In **TagList**, type the list of tag values.
11. In **Params**, type the SnmpAdminString.
12. In **TMask**, type the mask.
13. In **MMS**, type the maximum message size.
14. Select **Insert**.

Target Table Field Descriptions

Use the data in the following table to use the **Target Table** tab.

Name	Description
Name	Specifies a unique identifier for this table. The name is a community string.
TDomain	Specifies the transport type of the address. ipv4Tdomain specifies the transport type of address is an IPv4 address. ipv6Tdomain specifies the transport type of address is IPv6. The default is ipv4Tdomain .
TAddress	Specifies the transport address in xx.xx.xx.xx:port format, for example: 192.1.2.12:162, where 162 is the trap listening port on the system 192.1.2.12.
Timeout	Specifies the maximum round trip time required to communicate with the transport address. The value is in 1/100 seconds from 0–2147483647. The default is 1500. After the system sends a message to this address, if a response (if one is expected) is not received within this time period, you can assume that the response is not delivered.
RetryCount	Specifies the maximum number of retries if a response is not received for a generated message. The count can be in the range of 0–255. The default is 3.
TagList	Contains a list of tag values used to select target addresses for a particular operation. A tag refers to a class of targets to which the messages can be sent.
Params	Contains SNMP parameters used to generate messages to send to this transport address. For example, to receive SNMPv2C traps, use TparamV2.
TMask	Specifies the mask. The value can be empty or in six-byte hex string format. Tmask is an optional parameter that permits an entry in the TargetAddrTable to specify multiple addresses.
MMS	Specifies the maximum message size. The size can be zero, or 484–2147483647. The default is 484. Although the maximum message size is 2147483647, the device supports the maximum SNMP packet size of 8192.

Configure Target Table Parameters

About This Task

Configure the target table to configure the security parameters for SNMP. Configure the target table to configure parameters such as SNMP version and security levels.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **Target Table**.
3. Select the **Target Params Table** tab.
4. Select **Insert**.
5. In **Name**, type a target table name.
6. From **MPModel**, select an SNMP version.
7. From **Security Model**, select the security model.
8. In **SecurityName**, type `readview` or `writeview`.
9. From **SecurityLevel**, select the security level for the table.
10. Select **Insert**.

Target Params Table Field Descriptions

Use the data in the following table to use the **Target Params Table** tab.

Name	Description
Name	Identifies the target table.
MPModel	Specifies the message processing model to use to generate messages: SNMPv1, SNMPv2c, or SNMPv3/USM.
SecurityModel	Specifies the security model to use to generate messages: SNMPv1, SNMPv2c, or USM. You can receive an inconsistentValue error if you try to configure this variable to a value for a security model that the implementation does not support.
SecurityName	Identifies the principal on whose behalf SNMP messages are generated.
SecurityLevel	Specifies the security level used to generate SNMP messages: noAuthNoPriv, authNoPriv, or authPriv.

Configure SNMP Notify Filter Profiles

About This Task

Configure the SNMP table of filter profiles to determine whether particular management targets receive particular notifications.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **Notify Table**.

3. Select the **Notify Filter Table** tab.
4. Select **Insert**.
5. In **NotifyFilterProfileName**, type a name for the notify filter profile.
6. In **Subtree**, type the subtree location information in x.x.x.x.x.x.x.x format.
7. In **Mask**, type the mask location in hex string format.
8. From **Type**, select **included** or **excluded**.
9. Select **Insert**.

Notify Filter Table *Field Descriptions*

Use the data in the following table to use the **Notify Filter Table** tab.

Name	Description
NotifyFilterProfileName	Specifies the name of the filter profile used to generate notifications.
Subtree	Specifies the MIB subtree that, if you combine it with the mask, defines a family of subtrees, which are included in or excluded from the filter profile. For more information, see RFC 2573.
Mask	Specifies the bit mask (in hexadecimal format) that, in combination with the subtree, defines a family of subtrees, which are included in or excluded from the filter profile.
Type	Indicates whether the family of filter subtrees are included in or excluded from a filter. The default is included.

Configure SNMP Notify Filter Profile Table Parameters

Before You Begin

- The notify filter profile exists.

About This Task

Configure the profile table to associate a notification filter profile with a particular set of target parameters.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **Notify Table**.
3. Select the **Notify Filter Profile Table** tab.
4. Select **Insert**.
5. In **TargetParamsName**, type a name for the target parameters.
6. In **NotifyFilterProfileName**, type a name for the notify filter profile.
7. Select **Insert**.

Notify Filter Profile Table *Field Descriptions*

Use the data in the following table to use the **Notify Filter Profile Table** tab.

Name	Description
TargetParamsName	Specifies the unique identifier associated with this entry.
NotifyFilterProfileName	Specifies the name of the filter profile to use to generate notifications.

Enable Authentication Traps

About This Task

Enable the SNMP agent process to generate authentication-failure traps.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **General**.
3. Select the **Error** tab.
4. Select **AuthenticationTraps**.
5. Select **Apply**.

Error *Field Descriptions*

Use the data in the following table to use the **Error** tab.

Name	Description
AuthenticationTraps	Enables or disables the sending of traps after an error occurs. The default is disabled.
LastErrorCode	Specifies the last reported error code.
LastErrorSeverity	Specifies the last reported error severity: 0= Informative Information 1= Warning Condition 2= Error Condition 3= Manufacturing Information 4= Fatal Condition

Traps Reference

The switch generates alarms, traps, and logs. This section provides information about traps.

**Note**

The OID values of the rclsisTrap (OID 1.3.6.1.4.1.2272.1.63.9) table are populated only when the Duplicate ISIS Sys-id & Nickname condition is present. This MIB table captures this specific condition only. Use other tables such as rclsisAdjTable (OID 1.3.6.1.4.1.2272.1.63.10) or isisISAdjTable (OID 1.3.6.1.3.37.1.6.1) to gather IS-IS information.

Proprietary Traps

The following tables describe proprietary traps for the switch. All of the following traps have a status of current.

Table 190: 1.3.6.1.4.1.45.5.17.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.17.0.1	bsDhcpSnoopingBindingTableFull	bsDhcpSnoopingNotificationClientMACAddr	A bsDhcpSnoopingBindingTableFull notification is generated when you try adding a new DHCP binding entry, and the binding table is full. The value of bsDhcpSnoopingNotificationClientMACAddr gives the MAC address that cannot be added to the binding table. This notification also indicates that additional untrusted DHCP packets will not be added to the binding table, and will be dropped.
1.3.6.1.4.1.45.5.17.0.2	bsDhcpSnoopingTrap	<ul style="list-style-type: none"> bsDhcpSnoopingNotificationSourcePort bsDhcpSnoopingNotificationMsgType bsDhcpSnoopingNotificationSourceMACAddr bsDhcpSnoopingNotificationClientMACAddr bsDhcpSnoopingIfTrusted 	A bsDhcpSnoopingTrap notification is generated when a DHCP packet is dropped.
1.3.6.1.4.1.45.5.17.0.4	bsDhcpSnoopingStaticEntryMACConflict	bsDhcpSnoopingNotificationSourceMACAddr bsDhcpSnoopingIfIndex	A bsDhcpSnoopingStaticEntryMACConflict notification is generated when a DHCP packet is dropped, because a static entry with the same MAC address was found in the binding table.

Table 191: 1.3.6.1.4.1.45.5.18.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.18.0.1	bsaiArpPacketDroppedOnUntrustedPort	<ul style="list-style-type: none"> bsArpInspectionIfTrusted bsArpInspectionNotificationSourceMACAddr 	A bsaiArpPacketDroppedOnUntrustedPort trap signifies that an ARP packet is dropped on an untrusted port, due to an invalid IP or MAC address binding. The port is identified by the instance of bsArpInspectionIfTrusted generated by the trap.

Table 192: 1.3.6.1.4.1.45.5.20.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.20.0.2	bsSourceGuardCannotEnablePort	bsSourceGuardConfigMode	A bsSourceGuardCannotEnablePort trap signifies that there are insufficient resources to enable IP Source Guard checking on a port because of internal state changes within the system such as system initialization. The port is identified by the instance of bsSourceGuardConfigMode generated by the trap.
1.3.6.1.4.1.45.5.20.0.1	bsSourceGuardReachedMaxIpEntries	bsSourceGuardConfigMode	A bsSourceGuardReachedMaxIpEntries trap signifies that the maximum number of IP address entries on a port has been reached. The port is identified by the instance of bsSourceGuardConfigMode generated by the trap.

Table 193: 1.3.6.1.4.1.45.5.34.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.34.0.1	bsnesGloballyEnabled		A bsnesGloballyEnabled trap signifies that the Energy Saver feature was enabled globally.
1.3.6.1.4.1.45.5.34.0.2	bsnesGloballyDisabled		A bsnesGloballyDisabled trap signifies that the Energy Saver feature was disabled globally.
1.3.6.1.4.1.45.5.34.0.3	bsnesManuallyActivated		A bsnesManuallyActivated trap signifies that the Energy Saver was activated manually.
1.3.6.1.4.1.45.5.34.0.4	bsnesManuallyDeactivated		A bsnesManuallyDeactivated trap signifies that the Energy Saver feature was deactivated manually.
1.3.6.1.4.1.45.5.34.0.5	bsnesScheduleNotApplied		A bsnesScheduleNotApplied trap signifies that a schedule was not applied because SNTP is not synchronized.
1.3.6.1.4.1.45.5.34.0.6	bsnesScheduleApplied		A bsnesScheduleApplied trap signifies that SNTP is synchronized and the schedule is being applied.
1.3.6.1.4.1.45.5.34.0.7	bsnesActivated		A bsnesActivated trap signifies that the Energy Saver feature was activated by the schedule.
1.3.6.1.4.1.45.5.34.0.8	bsnesDeactivated		A bsnesDeactivated trap signifies that the Energy Saver feature was deactivated by the schedule.

Table 194: 1.3.6.1.4.1.45.5.43.0.xx series

ODID	Notification type	Objects	Description
1.3.6.1.4.1.45.5.43.0.1	bsLstInterfaceStatusChanged	<ul style="list-style-type: none"> • ifIndex • bsLstInterfaceStatus • bsLstGroupIndex 	A bsLstInterfaceStatusChanged trap signifies that a physical or logical interface changed status in a Link-state tracking (LST) tracking group.
1.3.6.1.4.1.45.5.43.0.2	bsLstInterfaceOperStateChanged	<ul style="list-style-type: none"> • ifIndex • bsLstInterfaceStatus • bsLstGroupIndex 	A bsLstInterfaceOperState Changed trap signifies that the operational status of an LST group changed due to an interface status change. For example, when the last interface in an LST group is down, the operational status of the LST group changes to down.

Table 195: 1.3.6.1.4.1.2272.1.4.10.1.1.xxx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.4.10.1.1.114	rcPortShutdownReason	<ul style="list-style-type: none"> • rcPortVLacpTotalFlapCount • rcPortVLacpFirstFlapTimeStamp • rcPortVLacpLastFlapTimeStamp 	An rcPortShutdownReason trap signifies that the specific VLACP link port status is down as VLACP flaps are detected on that port.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.3	rcnErrorNotification	rcErrorLevel rcErrorCode rcErrorText	An rcnErrorNotification trap signifies that the SNMPv2 entity, acting in an agent role, has detected an error condition.
1.3.6.1.4.1.2272.1.21.0.4	rcnStpNewRoot	rcStgId	An rcnStpNewRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Spanning Tree Protocol declared the device to be the new root of the spanning tree.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.5	rcnStpTopologyChange	rcStgId rcPortIndex	An rcnStpTopologyChange trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Spanning Tree Protocol has gone due to a topology change event.
1.3.6.1.4.1.2272.1.21.0.6	rcnChasPowerSupplyDown	<ul style="list-style-type: none"> rcChasPowerSupplyId rcChasPowerSupplyOperStatus 	An rcnChasPowerSupplyDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasPowerSupplyOperStatus object for one of its power supply units is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.7	rcnChasFanDown	<ul style="list-style-type: none"> rcChasFanId rcChasFanOperStatus 	An rcnChasFanDown trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasFanOperStatus object for one of its power supply units is about to transition into the down state.
1.3.6.1.4.1.2272.1.21.0.8	rcnLinkOscillation	rcPortIndex	An rcnLinkOscillation trap signifies that the SNMPv2 entity, acting in an agent role, has detected an excessive number of link state transitions on the specified port.
1.3.6.1.4.1.2272.1.21.0.9	rcnMacViolation	<ul style="list-style-type: none"> rcErrorText rcPortIndex 	An rcnMacViolation trap signifies that the SNMPv2 entity, acting in an agent role, has received a PDU with an invalid source MAC address.
1.3.6.1.4.1.2272.1.21.0.13	rcn2kTemperature	rc2kChassisTemperature	An rcn2kTemperature trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the chassis is overheating.
1.3.6.1.4.1.2272.1.21.0.14	rcnChasPowerSupplyUp	<ul style="list-style-type: none"> rcChasPowerSupplyId rcChasPowerSupplyOperStatus 	An rcnChasPowerSupplyUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasPowerSupplyOperStatus object for one of its power supply units is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.16	rcnStpTCN	<ul style="list-style-type: none"> rcStgId rcPortIndex rcStgBridgeAddresses 	An rcnStpTCN trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Spanning Tree Protocol has gone due to a topology change event.
1.3.6.1.4.1.2272.1.21.0.17	rcnSmltLstLinkUp		An rcnSmltLstLinkUp trap signifies that the split MLT link is from down to up.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.18	rcnSmltIstLinkDown	rcnSmltIstLinkDown	An rcnSmltIstLinkDown trap signifies that the split MLT link is from up to down.
1.3.6.1.4.1.2272.1.21.0.19	rcnSmltLinkUp	rcMltSmltId	An rcnSmltLinkUp trap signifies that the split SMLT link is up.
1.3.6.1.4.1.2272.1.21.0.20	rcnSmltLinkDown	rcMltSmltId	An rcnSmltLinkDown trap signifies that the split SMLT link is down.
1.3.6.1.4.1.2272.1.21.0.21	rcnChasFanUp	<ul style="list-style-type: none"> rcChasFanId rcChasFanOperStatus 	An rcnChasFanUp trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the rcChasFanOperStatus object for one of its power supply units is about to transition into the up state.
1.3.6.1.4.1.2272.1.21.0.22	rcnPasswordChange	<ul style="list-style-type: none"> rcCliPasswordChange rcCliPassChangeResult 	An rcnPasswordChange trap signifies that the SNMPv2 entity, acting in an agent role, has detected that one of the CLI passwords is changed.
1.3.6.1.4.1.2272.1.21.0.23	rcnEmError	<ul style="list-style-type: none"> rc2kCardIndex rcChasEmModeError 	An rcnEmError trap signifies that the SNMPv2 entity, acting in an agent role, has detected an Em error.
1.3.6.1.4.1.2272.1.21.0.26	rcnSmartCpldTimerFired	rc2kCardIndex	An rcnSmartCpldTimerFired trap signifies that the CP ID timer fired.
1.3.6.1.4.1.2272.1.21.0.27	rcnCardCpldNotUpDate	rc2kCardIndex	An rcnCardCpldNotUpDate trap signifies that the CP ID is not up-to-date.
1.3.6.1.4.1.2272.1.21.0.28	rcnIgapLogFileFull		An rcnIgapLogFileFull trap signifies that the IGAP accounting time-out Log File has reached the maximum.
1.3.6.1.4.1.2272.1.21.0.30	rcnSshServerEnabled	rcSshGlobalPort	An rcnSshServerEnabled trap signifies that the SSH server is enabled.
1.3.6.1.4.1.2272.1.21.0.31	rcnSshServerDisabled	rcSshGlobalPort	An rcnSshServerDisabled trap signifies that the SSH server is disabled.
1.3.6.1.4.1.2272.1.21.0.37	rcnSaveConfigAction	rcSysActionL1	An rcnSaveConfigAction trap indicates that the switch run time or boot configuration is being saved.
1.3.6.1.4.1.2272.1.21.0.38	rcnLoopDetectOnPort	<ul style="list-style-type: none"> rcVlanId rcPortIndex 	An rcnLoopDetectOnPort trap indicates that a loop has been detected on a port. The VLAN on that port will be disabled.
Note: This trap does not apply to all platforms.			

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.41	rcnAggLinkUp	rcMltId	An rcnAggLinkUp trap is generated when the operational state of the aggregator changes from down to up.
1.3.6.1.4.1.2272.1.21.0.42	rcnAggLinkDown	rcMltId	An rcnAggLinkDown trap is generated when the operational state of the aggregator changes from up to down.
1.3.6.1.4.1.2272.1.21.0.59	rcnFdbProtectViolation	<ul style="list-style-type: none"> rcVlanId rcPortIndex 	The rcnFdbProtectViolation trap signifies that the port has violated the user-configured limit for total number of fdb-entries learned on that port.
1.3.6.1.4.1.2272.1.21.0.60	rcnLogMsgControl	<ul style="list-style-type: none"> rcSysMsgLogFrequency rcSysMsgLogText 	An rcnLogMsgControl trap signifies whether the number of times of repetition of the particular Log message has exceeded the particular frequency or count.
1.3.6.1.4.1.2272.1.21.0.61	rcnSaveConfigFile	<ul style="list-style-type: none"> rcSysActionL1 rcSysConfigFileName 	An rcnSaveConfig trap signifies that either the runtime config or the boot config has been saved on the switch.
1.3.6.1.4.1.2272.1.21.0.62	rcnDNSRequestResponse	<ul style="list-style-type: none"> rcSysDnsServerListIpAddr rcSysDnsRequestType 	An rcnDnsRequestResponse trap signifies that the switch sent a query to the DNS server or that it received a successful response from the DNS Server.
1.3.6.1.4.1.2272.1.21.0.63	rcnDuplicateIpAddress	<ul style="list-style-type: none"> ipNetToMediaNetAddress ipNetToMediaPhysAddress 	An rcnDuplicateIpAddress trap signifies that a duplicate IP address is detected on the subnet.
1.3.6.1.4.1.2272.1.21.0.64	rcnLoopDetectPortDown	<ul style="list-style-type: none"> rcPortIndex ifAdminStatus ifOperStatus 	An rcnLoopDetectPortDown trap signifies that a loop has been detected on a port and the port is going to shut down.
1.3.6.1.4.1.2272.1.21.0.67	rcnLoopDetectMacDiscard	<ul style="list-style-type: none"> rcPortIndex rcSysMacFlapLimitTime rcSysMacFlapLimitCount 	An rcnLoopDetectMacDiscard trap signifies that a loop has been detected on a port and the MAC address will be discarded on all ports in that VLAN.
1.3.6.1.4.1.2272.1.21.0.68	rcnAutoRecoverPort	rcPortIndex	An rcnAutoRecoverPort trap signifies that autorecovery has reenabled a port that was previously disabled by link flap.
1.3.6.1.4.1.2272.1.21.0.69	rcnAutoRecoverLoopDetectedPort	rcVlanNewLoopDetectedAction	An rcnAutoRecoverPort trap signifies that autorecovery has cleared the loop detect action that was taken on a port.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.74	rcnTacacsAuthFailure	rcTacacsGlobalLastUserName	An rcnTacacsAuthFailure trap signifies that TACACS+ authentication failed for a user.
1.3.6.1.4.1.2272.1.21.0.75	rcnTacacsNoServers		An rcnTacacsNoServers trap signifies that you are unable to use any TACACS+ servers for authentication.
1.3.6.1.4.1.2272.1.21.0.76	rcnTacacsRxUnsupportedFrame	<ul style="list-style-type: none"> rcTacacsGlobalLastAddressType rcTacacsGlobalLastAddress 	An rcnTacacsRxUnsupportedFrame trap signifies that an unsupported frame was received from the TACACS+ server.
1.3.6.1.4.1.2272.1.21.0.77	rcnTacacsExceededMaxLogins		An rcnTacacsExceededMaxLogins trap signifies that there was an attempt to exceed the maximum number of allowed TACACS+ logins.
1.3.6.1.4.1.2272.1.21.0.78	rcnTacacsClientFailure		An rcnTacacsClientFailure trap signifies that the TACACS+ Client application is down.
1.3.6.1.4.1.2272.1.21.0.80	rcnVlcpPortDown	rcPortIndex	An rcnVlcpPortDown trap signifies that VLACP is down on the port specified.
1.3.6.1.4.1.2272.1.21.0.81	rcnVlcpPortUp	rcPortIndex	An rcnVlcpPortUp trap signifies that VLACP is up on the port specified.
1.3.6.1.4.1.2272.1.21.0.83	rcnEapMacIntrusion	<ul style="list-style-type: none"> rcSysIpAddr rcRadiusPaePortNumber rcRadiusEapLastAuthMac rcRadiusEapLastRejMac 	An rcnEapMacIntrusion trap signifies that an EAP MAC intrusion has occurred on this port.
1.3.6.1.4.1.2272.1.21.0.110	rcnMaxRouteWarnClear	rcVrfName	An rcnMaxRouteWarnClear trap signifies that the number of routes in the routing table of the virtual router has dropped below the warning threshold.
1.3.6.1.4.1.2272.1.21.0.111	rcnMaxRouteWarnSet	rcVrfName	An rcnMaxRouteWarnSet trap signifies that the routing table for the virtual router is reaching its maximum size. Take action to prevent this.
1.3.6.1.4.1.2272.1.21.0.112	rcnMaxRouteDropClear	rcVrfName	An rcnMaxRouteDropClear trap signifies that the routing table for the virtual router is no longer dropping new routes as it is below the maximum size.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.113	rcnMaxRouteDropSet	rcVrfName	An rcnMaxRouteDropSet trap signifies that the routing table for the virtual router has reached the maximum size, and is now dropping all new nonstatic routes.
1.3.6.1.4.1.2272.1.21.0.117	rcnMstpNewCistRoot	rcStgBridgeAddress	An rcnMstpNewCistRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Multiple Spanning Tree Protocol has declared the device to be the new root of the common internal spanning tree.
1.3.6.1.4.1.2272.1.21.0.118	rcnMstpNewMstiRoot	<ul style="list-style-type: none"> rcStgBridgeAddresses rcStgId 	An rcnMstpNewMstiRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Multiple Spanning Tree Protocol has declared the device to be the new root of the spanning tree instance.
1.3.6.1.4.1.2272.1.21.0.119	rcnMstpNewCistRegionalRoot	rcStgBridgeAddress	An rcnMstpNewCistRegionalRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Multiple Spanning Tree Protocol has declared the device to be the new regional root of the common internal spanning tree.
1.3.6.1.4.1.2272.1.21.0.120	rcnRstpNewRoot	rcStgBridgeAddress	An rcnRstpNewRoot trap signifies that the SNMPv2 entity, acting in an agent role, has detected that the Rapid Spanning Tree Protocol has declared the device to be the new root of the spanning tree.
1.3.6.1.4.1.2272.1.21.0.124	rcnRsmltEdgePeerModified	rcVlanId	An rcnRsmltEdgePeerModified trap signifies that the RSMLT peer address is different from that of the stored address. You must save the configuration if EdgeSupport has to use this information on the next restart.
1.3.6.1.4.1.2272.1.21.0.143	rcn2kGbicRemoved	<ul style="list-style-type: none"> rc2kCardIndex rcPortIndex rcPortUserLabel1 rcPortUserLabel2 rc2kChassisUserLabel1 	An rcGbicRemoved trap signifies that the SNMPv2 entity, acting in an agent role has detected that an XFP or SFP is removed from the specified slot or port.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.144	rcn2kGbicInserted	<ul style="list-style-type: none"> rc2kCardIndex rcPortIndex rcPortUserLabel1 rcPortUserLabel2 rc2kChassisUserLabel1 	An rcGbicInserted trap signifies that the SNMPv2 entity, acting in an agent role has detected that the an XFP or SFP is inserted in the specified slot or port.
1.3.6.1.4.1.2272.1.21.0.167	rcnChasPowerSupplyNoRedundancy		An rcnChasPowerSupplyNoRedundancy trap signifies that the chassis is running on a power supply without redundancy.
1.3.6.1.4.1.2272.1.21.0.168	rcnChasPowerSupplyRedundancy		An rcnChasPowerSupplyRedundancy trap signifies that the chassis is running on a power supply with redundancy.
1.3.6.1.4.1.2272.1.21.0.171	rcnLicenseTrialPeriodExpired		An rcnLicenseTrialPeriodExpired trap signifies that the Trial Period License has expired.
1.3.6.1.4.1.2272.1.21.0.172	rcnLicenseTrialPeriodExpiry	rcSysLicenseTrialDaysLeft	An rcnLicenseTrialPeriodExpiry trap signifies the time remaining, in days, before the Trial Period License expires.
1.3.6.1.4.1.2272.1.21.0.173	rcnVrfUp	rcVrfName rcVrfOperStatus	This notification is generated when the operational status of the specified VRF is toggled from down to up.
1.3.6.1.4.1.2272.1.21.0.174	rcnVrfDown	rcVrfName rcVrfOperStatus	This notification is generated when the operational status of the specified VRF is toggled from up to down.
1.3.6.1.4.1.2272.1.21.0.175	rcnMrouteIngressThresholdExceeded	<ul style="list-style-type: none"> rcIpResourceUsageGlobalIngressRecInUse rcIpResourceUsageGlobalIngressThreshold 	This notification is generated when the number of mroute ingress records exceeds the ingress threshold.
1.3.6.1.4.1.2272.1.21.0.176	rcnMrouteEgressThresholdExceeded	<ul style="list-style-type: none"> rcIpResourceUsageGlobalIngressRecInUse rcIpResourceUsageGlobalIngressThreshold 	This notification is generated when the number of mroute egress records exceeds the egress threshold.
1.3.6.1.4.1.2272.1.21.0.185	rcnChasPowerSupplyRunningLow		An rcnChasPowerSupplyRunningLow trap signifies that the chassis is running on low power supply.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.192	rcnIsisPlsbMetricMismatchTrap	<ul style="list-style-type: none"> • rclsisLocalLspId • rclsisLocalL1Metric • rclsisNgbLspId • rclsisNgbL1Metric • rclsisPlsbTrapType • rclsisTrapIndicator • rclsisLocalHostName • rclsisNgbHostName 	An rcnIsisPlsbMetricMismatchTrap signifies that a Link State Packet (LSP) with a different value of Level 1 metric is received.
1.3.6.1.4.1.2272.1.21.0.193	rcnIsisPlsbDuplicateSysidTrap	<ul style="list-style-type: none"> • rclsisLocalSysId • rclsisLocalInterface • rclsisPlsbTrapType • rclsisTrapIndicator 	An rcnIsisPlsbDuplicateSysidTrap signifies that a Hello packet with a duplicate system ID is received.
1.3.6.1.4.1.2272.1.21.0.194	rcnIsisPlsbLsdbUpdateTrap	rclsisPlsbTrapType	An rcnIsisPlsbLsdbUpdateTrap signifies that link state database (LSDB) information has changed.
1.3.6.1.4.1.2272.1.21.0.196	rcnChasFanCoolingLow	<ul style="list-style-type: none"> • rcChasFanOperStatus • rcChasFanType • rcErrorLevel • rcErrorText 	An rcnaChasFanCoolingLow trap signifies that the chassis is running on low fan cooling.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.278	rcnIsisPlsbBvidMismatchTrap	<ul style="list-style-type: none"> • rclsisLocalSysId • rclsisLocalPrimaryBvid • rclsisLocalPrimaryTieBrkAlg • rclsisLocalSecondaryBvid • rclsisLocalSecondaryTieBrkAlg • rclsisNgbSysId • rclsisNgbPrimaryBvid • rclsisNgbPrimaryTieBrkAlg • rclsisNgbSecondaryBvid • rclsisNgbSecondaryTieBrkAlg • rclsisLocalBvidCounter • rclsisNgbBvidCounter • rclsisPlsbTrapType • rclsisTrapIndicator • rclsisNgbHostName 	An rcnIsisPlsbBvidMismatchTrap signifies when a backbone VLAN ID (BVID) Type-Length-Value (TLV) from a neighbor node does not match the local configuration.
1.3.6.1.4.1.2272.1.21.0.279	rcnIsisPlsbSmltVirtBmacMismatchTrap	<ul style="list-style-type: none"> • rclsisLocalVirtualBmac • rclsisPeerVirtualBmac • rclsisPlsbTrapType • rclsisTrapIndicator 	An rcnIsisPlsbSmltVirtBmacMismatchTrap signifies that the virtual Backbone MAC (BMAC) configured in the switch is different from the virtual BMAC configured on the interswitch trunking (IST) peer.
1.3.6.1.4.1.2272.1.21.0.280	rcnIsisPlsbSmltPeerBmacMismatchTrap	<ul style="list-style-type: none"> • rclsisSysId • rclsisSmltPeerSysId • rclsisPlsbTrapType • rclsisTrapIndicator 	<p>An rcnIsisPlsbSmltPeerBmacMismatchTrap signifies one of the following situations:</p> <ul style="list-style-type: none"> • The Split MultiLink Trunking (SMLT) peer Backbone MAC (BMAC) configured in the interswitch trunking (IST) peer is different from the Intermediate-System-to-Intermediate-System (IS-IS) System ID of the local switch. • The SMLT peer BMAC configured on the local switch is different from the IS-IS System ID of the IST peer.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.281	rcnIsisPlsbAdjStateTrap	<ul style="list-style-type: none"> • rclsisNgbSysId • rclsisLocalInterface • rclsisPlsbTrapType • rclsisAdjState • rclsisNgbHostName 	An rcnIsisPlsbAdjStateTrap signifies when IS-IS adjacency state changes.
1.3.6.1.4.1.2272.1.21.0.282	rcnIsisPlsbDuplicateNNameTrap	<ul style="list-style-type: none"> • rclsisNgbNickname • rclsisPlsbTrapType • rclsisTrapIndicator • rclsisNgbSysId • rclsisDuplicateNnameCounter • rclsisNgbHostName 	An rcnIsisPlsbDuplicateNNameTrap signifies that a Link State Packet (LSP) with a duplicate nickname is received. The trap should be generated by all the switches in the network.
1.3.6.1.4.1.2272.1.21.0.283	rcnIsisPlsbSmltSplitBebMismatchTrap	<ul style="list-style-type: none"> • rclsisLocalSmltSplitBeb • rclsisPeerSmltSplitBeb • rclsisPlsbTrapType • rclsisTrapIndicator 	An rcnIsisPlsbSmltSplitBebMismatchTrap signifies that the SMLT Split Backbone Edge Bridge (BEB) configured on the local switch and on the IST peer are the same. One IST switch must be configured as the primary Split BEB and the other IST peer must be configured as the secondary Split BEB.
1.3.6.1.4.1.2272.1.21.0.284	rcnIsisPlsbMultiLinkAdjTrap	<ul style="list-style-type: none"> • rclsisNgbSysId • rclsisLocalInterface • rclsisPrevInterface • rclsisPlsbTrapType • rclsisNgbHostName • rclsisTrapIndicator 	An rcnIsisPlsbMultiLinkAdjTrap signifies when the Intermediate-System-to-Intermediate-System (IS-IS) protocol forms more than one adjacency with the same IS-IS.
1.3.6.1.4.1.2272.1.21.0.285	rcnaSshSessionLogout	rcSshGlobalHostIpAddress	An rcnaSshSessionLogout trap signifies a Secure Shell (SSH) session logout.
1.3.6.1.4.1.2272.1.21.0.286	rcnaSshUnauthorizedAccess	rcSshGlobalHostIpAddress	An rcnaSshUnauthorizedAccess trap signifies that unauthorized access has occurred. It is deprecated by rcnaSshUnauthorizedAccess.
1.3.6.1.4.1.2272.1.21.0.287	rcnaAuthenticationSuccess	<ul style="list-style-type: none"> • rcLoginUserName • rcLoginHostIpAddress 	An rcnaAuthenticationSuccess trap signifies that a login is successful. The Trap includes the login user name and the host IP address. It is deprecated by rcnaAuthenticationSuccess.
1.3.6.1.4.1.2272.1.21.0.288	rcnaSshSessionLogin	rcSshGlobalHostIpAddress	An rcnaSshSessionLogin trap signifies that there is a Secure Shell (SSH) session login.

Table 196: 1.3.6.1.4.1.2272.1.21.0.xx series (continued)

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.21.0.305	RclsisPlsbSmltVirtBmacMisconfigTrap	<ul style="list-style-type: none"> • rclsisSmltVirtBmacMisconfigNodeSysId • rclsisPlsbTrapType • rclsisSmltVirtBmacMisconfigNodeHostName • rclsisTrapIndicator 	An SPBM ISIS trap signifies that SMLT virtual BMAC has been used by nodes other than the SMLT nodes as system-id or MAC.
1.3.6.1.4.1.2272.1.21.0.306 Note: This trap does not apply to all platforms.	rcnPortChannelizedStateChangeTrap	<ul style="list-style-type: none"> • rcPortIndex • rcChannelizedPortAdminMode 	An rcnPortChannelizedStateChangeTrap notification signifies that a port channelized state has changed by administratively enabling or disabling.
1.3.6.1.4.1.2272.1.21.0.335	rcnSystemUsbInternalAccessErrorTrap		An rcnSystemUsbInternalAccessErrorTrap notification signifies that the system has lost internal access to the USB. This trap only applies to platforms that require the USB as part of the operating system.
1.3.6.1.4.1.2272.1.21.0.341	rcnDvrVistPeerDomainMismatchErrorTrap	rclsisPeerVirtualBmac	An rcnDvrVistPeerDomainMismatchErrorTrap notification is generated when the VIST link comes up between two DvR peers that are in different DvR domains.
1.3.6.1.4.1.2272.1.21.0.342	rcnDvrVistPeerDomainMismatchErrorClearTrap	rclsisPeerVirtualBmac	An rcnDvrVistPeerDomainMismatchErrorTrap notification is generated when the error condition of having a VIST link up between two DvR peers from different DvR domains is cleared.
1.3.6.1.4.1.2272.1.21.0.351	rcnIsisPlsbIsisEnabledWithZeroNicknameTrap	<ul style="list-style-type: none"> • rclsisLocalSysId • rclsisPlsbTrapType • rclsisTrapIndicator 	An rcnIsisPlsbIsisEnabledWithZeroNicknameTrap notification is generated when the IS-IS is enabled with a zero nickname.
1.3.6.1.4.1.2272.1.21.0.352	rcnRestConfServerOperationStatusTrap	rcnRestConfServerOperationStatus	An rcnRestConfServerOperationStatusTrap notification is generated when the RESTCONF server is enabled or disabled to indicate the operational status of the RESTCONF server.

Table 197: 1.3.6.1.4.1.2272.1.206.x.x.x series

OID	Notification type	Objects	Description
1.3.6.1.4.1.2272.1.206.1.0.1	rcVrrpTmpTrapNewMaster	<ul style="list-style-type: none"> rcVrrpTmpOperationsMasterIpAddr rcVrrpTmpNewMasterReason 	This notification is generated when Virtual Router Redundancy Protocol (VRRP) transitions to the master.
1.3.6.1.4.1.2272.1.206.2.2.1	rcVrrpExtTrapStateTransition	<ul style="list-style-type: none"> ifIndex rcVrrpExtTrapStateTransitionType rcVrrpExtTrapStateTransitionCause rcVrrpExtOperationSVrId rcVrrpTmpOperationsPrimaryIpAddr rcVrrpTmpOperationsMasterIpAddr 	This notification is generated when a transition happens in the state of Virtual Router Redundancy Protocol (VRRP), for instance, a transition from master to backup when shutdown is received.

Standard Traps

The following table describes standard traps that the switch can generate.

Table 198: Standard traps

OID	Notification type	Objects	Description
1.3.6.1.2.1.16.0.1	risingAlarm	alarmIndex alarmVariable alarmSampleType alarmValue alarmRisingThreshold	The SNMP trap that is generated after an alarm entry crosses the rising threshold and generates an event that is configured to send SNMP traps. TRAP TYPE ENTERPRISE rmon
1.3.6.1.2.1.16.0.2	fallingAlarm	alarmIndex alarmVariable alarmSampleType alarmValue alarmFallingThreshold	The SNMP trap that is generated after an alarm entry crosses the falling threshold and generates an event that is configured to send SNMP traps. TRAP TYPE ENTERPRISE rmon
1.3.6.1.2.1.46.1.3.0.3	vrrpTrapStateTransition	ifIndex vrrpTrapStateTransitionType vrrpTrapStateTransitionCause vrrpOperVrId vrrpOperIpAddr ipAdEntAddr	A vrrpTrapStateTransition trap signifies a state transition has occurred on a particular Virtual Router Redundancy Protocol (VRRP) interface. Implementation of this trap is optional. vrrpOperIpAddr contains the IP address of the VRRP interface while ipAdEntAddr contains the IP address assigned to the physical interface.
1.3.6.1.2.1.68.0.1	vrrpTrapNewMaster	vrrpOperMasterIpAddr	The newMaster trap indicates that the sending agent has transitioned to Master state.

Table 198: Standard traps (continued)

OID	Notification type	Objects	Description
1.3.6.1.2.1.68.0.2	vrrpTrapAuthFailure	vrrpTrapPacketSrc vrrpTrapAuthErrorType	A vrrpAuthFailure trap signifies that a packet has been received from a router whose authentication key or authentication type conflicts with the authentication key or authentication type of this router.
1.3.6.1.2.1.80.0.1	pingProbeFailed	pingCtlTargetAddressType pingCtlTargetAddress pingResultsOperStatus pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponse pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	This trap is generated after a probe failure is detected when the corresponding pingCtlTrapGeneration object is configured to probeFailure(0) subject to the value of pingCtlTrapProbeFailureFilter. The object pingCtlTrapProbeFailureFilter can specify the number of successive probe failures required before this notification can be generated.
1.3.6.1.2.1.80.0.2	pingTestFailed	pingCtlTargetAddressType pingCtlTargetAddress pingResultsOperStatus pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponses pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	This trap is generated after a ping test fails when the corresponding pingCtlTrapGeneration object is configured to testFailure(1). In this instance pingCtlTrapTestFailureFilter specifies the number of probes in a test required to fail to consider the test as failed.

Table 198: Standard traps (continued)

OID	Notification type	Objects	Description
1.3.6.1.2.1.80.0.3	pingTestCompleted	pingCtlTargetAddressType pingCtlTargetAddress pingResultsOperStatus pingResultsIpTargetAddressType pingResultsIpTargetAddress pingResultsMinRtt pingResultsMaxRtt pingResultsAverageRtt pingResultsProbeResponses pingResultsSentProbes pingResultsRttSumOfSquares pingResultsLastGoodProbe	This trap is generated at the completion of a ping test when the corresponding pingCtlTrapGeneration object is configured to testCompletion(4).
1.3.6.1.2.1.81.0.1	traceRoutePathChange	traceRouteCtlTargetAddressType traceRouteCtlTargetAddress traceRouteResultsIpTgtAddrType traceRouteResultsIpTgtAddr	This trap is generated after the path to a target changes.
1.3.6.1.2.1.81.0.2	traceRouteTestFailed	traceRouteCtlTargetAddressType traceRouteCtlTargetAddress traceRouteResultsIpTgtAddrType traceRouteResultsIpTgtAddr	This trap is generated is traceroute cannot determine the path to a target (traceRouteNotifications 2).
1.3.6.1.2.1.81.0.3	traceRouteTestCompleted	traceRouteCtlTargetAddressType traceRouteCtlTargetAddress traceRouteResultsIpTgtAddrType traceRouteResultsIpTgtAddr	This trap is generated after the path to a target is determined.
1.3.6.1.6.3.1.1.5.1	coldStart		A coldStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing and that its configuration may have been altered.

Table 198: Standard traps (continued)

OID	Notification type	Objects	Description
1.3.6.1.6.3.1.1.5.2	warmStart		A warmStart trap signifies that the SNMPv2 entity, acting in an agent role, is reinitializing such that its configuration is unaltered.
1.3.6.1.6.3.1.1.5.3	linkDown		A linkDown trap signifies that the sending protocol entity recognizes a failure in one of the communication links represented in the agent configuration. TRAP-TYPE ENTERPRISE snmp
1.3.6.1.6.3.1.1.5.4	linkUp		A linkUp trap signifies that the sending protocol entity recognizes that one of the communication links represented in the agent configuration has come up. TRAP-TYPE ENTERPRISE snmp
1.3.6.1.6.3.1.1.5.5	authenticationFailure		



MACsec

[MACsec Fundamentals](#) on page 2601

[MACsec Configuration Using CLI](#) on page 2615

[MACsec Configuration using EDM](#) on page 2636

[MACsec Performance](#) on page 2645

Table 199: MACsec product support

Feature	Product	Release introduced
MACsec 2AN mode (static)	4220 Series	Not Applicable
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Fabric Engine 9.0 Supported on 7520-48YE-8CE For MACsec Key Agreement (MKA)
	7720 Series	Not Supported
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	Not Supported

Table 199: MACsec product support (continued)

Feature	Product	Release introduced
MACsec 4AN mode (static)	4220 Series	Not Applicable
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-4X-XT or 5320-16P-2MXT-2X
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.2.5 Not supported on 5520-24X or 5520-VIM-4X 5520-24T and 5520-24W on all fixed ports except 1/25 and 1/26 5520-12MW-36W, 5520-48SE, 5520-48T, and 5520-48W on all fixed ports except 1/49 and 1/50
	5720 Series	Fabric Engine 8.7 5720-24MW and 5720-24MXW on all fixed ports except 1/25 and 1/26 5720-48MW and 5720-48MXW on all fixed ports except 1/49 and 1/50
	7520 Series	Fabric Engine 9.0 Supported on 7520-48YE-8CE
	7720 Series	Not Supported
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	Not Supported
Automatic QoS priority for MACsec packets on intermediate switches	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.9
	5420 Series	Fabric Engine 8.9
	5520 Series	Fabric Engine 8.9
	5720 Series	Fabric Engine 8.9
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.9
	VSP 7400 Series	VOSS 8.9

Table 199: MACsec product support (continued)

Feature	Product	Release introduced
MACsec Key Agreement (MKA)	4220 Series	Not Applicable
	5320 Series	Fabric Engine 9.4 Not supported on 5320-16P-2MXT-2X or 5320-24T-4X-XT
	5420 Series	Fabric Engine 9.4
	5520 Series	VOSS 8.5 Not supported on 5520-24X or 5520-VIM-4X
	5720 Series	Fabric Engine 8.7 5720-24MW and 5720-24MXW on all fixed ports except 1/25 and 1/26 5720-48MW and 5720-48MXW on all fixed ports except 1/49 and 1/50
	7520 Series	9.0 Supported on 7520-48YE-8CE
	7720 Series	Not Supported
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	Not Supported

Table 199: MACsec product support (continued)

Feature	Product	Release introduced
MACsec Key Agreement (MKA) Keychain support	4220 Series	Not Supported
	5320 Series	Fabric Engine 9.4 Not supported on 5320-16P-2MXT-2X or 5320-24T-4X-XT
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.3 Not supported on 5520-24X and 5520-VIM-4X
	5720 Series	Fabric Engine 9.3 5720-24MW and 5720-24MXW on all fixed ports except 1/25 and 1/26 5720-48MW and 5720-48MXW on all fixed ports except 1/49 and 1/50
	7520 Series	Fabric Engine 9.3 Supported on 7520-48YE-8CE
	7720 Series	Not Supported
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.3
	VSP 7400 Series	Not Supported

MACsec Fundamentals

MAC Security (MACsec) is based on the IEEE 802.1ae standard that allows authorized systems in a network to transmit data confidentially and to protect against data transmitted or modified by unauthorized devices.

You can use MACsec for core and enterprise edge switches to do the following:

- Secure site-to-site connectivity between data centers.
- Provide data security on links that run over public ground.
- Provide data security on links that run outside the physically secure boundaries of a site.

You can use MACsec on access switches to secure host-to-switch connectivity, and host-to-switch connectivity in an environment where both trusted and untrusted hosts coexist.

In addition to host level authentication, MACsec capable LANs provide data origin authentication, data confidentiality, and data integrity between authenticated hosts or

systems. MACsec protects data from external hacking while the data passes through the public network to reach a receiver host.

MACsec enabled hosts encrypt and decrypt every frame exchanged between them using a MACsec key. The source MACsec host encrypts data frames, and the destination MACsec host decrypts the frames, ensuring delivery of the frame in its original condition to the recipient host. This ensures secure data communication.

**Note**

Before you enable MACsec on the 5320 Series or 5420 Series you must configure the macsec boot flag.

You can configure MACsec encryption over any type of point-to-point Ethernet or emulated Ethernet connection, which includes:

- Dark fiber
- Conventional wavelength-division multiplexing/dense wavelength-division multiplexing (CWDM/DWDM) service
- Multiprotocol label switching (MPLS) point-to-point (ELINE)
- Provider Backbone Bridge Traffic Engineering (PBB-TE)

You can configure MACsec on physical ports only. However, the physical ports can belong to an MLT trunk group that includes: Split MultiLink Trunking (SMLT), distributed MultiLink Trunking (DMLT), or Link aggregate group (LAG).

You configure a pre-shared key on either end of the MACsec link. The pre-shared key is an interface parameter, not a switch-wide parameter.

MACsec encrypts all packets. If you configure MACsec on one or more MultiLink Trunking (MLT) port members on one side, you must configure MACsec on the same port members on the other side. If you do not do this, the port can physically be enabled, but any overlying protocols can be disabled. You do not have to provision MACsec on all MLT port members, but if you configure MACsec on an MLT port member on one side, you must also provision MACsec on the corresponding MLT port on the other side.

One way to detect a mismatch of MACsec configuration is to use Virtual Link Aggregation Control Protocol (VLACP) on the links. If VLACP is enabled on an MKA-enabled link, it takes approximately 30 seconds for the VLACP session to begin.

MACsec provides security at the data link layer or the physical layer. It provides enhancements at the MAC service sub layer for its operation and services to the upper layer.

MACsec is an interface-level feature and is disabled by default.

**Note**

On 5320 Series and 5420 Series, Fabric Extend is not available if you boot the switch when the MACsec boot flag is enabled.

MACsec Keys

MACsec provides industry-standard security through secure point-to-point Ethernet links. The point-to-point links are secured after matching security keys.

Security keys are of two types:

- Connectivity Association Key (CAK), which is a configured pre-shared key.



Important

The switch supports the configuration of a pre-shared key to enable MACsec using the static connectivity association key (CAK) security mode.

The CAK must be identical across both ends of MACsec links.

- Secure Association Key (SAK):
 - Static SAKs: SAKs are static short-lived keys derived from the CAK or pre-configured for a particular secure channel (SC). MACsec uses a timer to refresh these keys so that the key and the session are secure.
 - Dynamic SAKs: MACsec Key Agreement (MKA) protocol generates SAKs. The MKA protocol determines which switch on the point-to-point link becomes the key server. The key server then generates SAKs and distributes them to the switch at the other end of the point-to-point link.

MACsec uses derived keys to encrypt or decrypt data at each end of the MACsec links.

Integrity Check Verification (ICV)

MACsec ensures data integrity using Integrity Check Verification (ICV). MACsec introduces an 8-byte or 16-byte SecTag after the Ethernet header, and an 8-byte or 16-byte calculated ICV after the Encrypted Payload. MACsec computes the ICV for the entire frame, starting from the Ethernet header, SecTag until the Checksum. The receiving side recalculates the ICV after data decryption, and verifies if the received ICV and computed ICV match. If the ICVs do not match, it indicates that data is modified, and MACsec drops the frame.

MACsec Security Modes

The static Connectivity Association Key (CAK) security mode is the only supported MACsec security mode on the platform, and is also the most common mode to enable MACsec.

When you use the static CAK security mode to enable MACsec, you configure a connectivity association on both ends of the link. Secure Association Keys (SAK) establish the MACsec relationship between the switches on each end of the Ethernet link. The SAKs include a connectivity association key name (CKN) and its own CAK. The MACsec CKN and CAK are configured in a connectivity association, and the CAK must match on both ends of the link to initially enable MACsec.

To ensure link security, the system periodically refreshes keys based on traffic volume and link speed.

To enable MACsec at the port level, you must first associate the port to the connectivity association. You complete the configuration within the connectivity association, but outside of the secure channel.

**Note**

If you use MKA, you must apply MKA profile to a port before you associate it with a Connectivity Association (CA). After you associate the port with a CA, you cannot enable MKA on the port.

When you use the static CAK security mode, the system automatically creates two secure channels, one for inbound traffic and another for outbound traffic. You cannot configure any parameters in the automatically created secure channels.

The CAK security mode ensures security by frequently refreshing to a new random security key, and by only sharing the security key between the two devices on the MACsec-secured point-to-point link.

MACsec provides options to encrypt all data, or configure a confidentiality offset, which specifies the number of unencrypted bytes in a frame that precede MACsec encryption.

You can configure the following optional features:

- Data encryption — If you disable encryption, MACsec forwards traffic in clear text. You can view that data that is not encrypted in the Ethernet frame that travels across the link. Even if you disable encryption the MACsec header applies to the frame and integrity checks make sure that traffic has not been tampered with.
- Confidentiality offset — If encryption is enabled, and an offset is not configured, all traffic in the connectivity is encrypted. The confidentiality offset specifies the bytes after the Ethernet header from which data encryption begins. Valid values are 30 bytes and 50 bytes. The confidentiality offset facilitates traffic flow inspection and classification on intermediate devices by not encrypting the Network Layer header for IPv4 or IPv6. For instance, if you configure the offset to 30, the IPv4 header and the TCP/UDP header are not encrypted. If you configure the offset to 50, the IPv6 header and the TCP/UDP header are not encrypted.

Connectivity Associations and Secure Channels

You configure MACsec in connectivity associations (CA). You can enable MACsec after you associate a connectivity association with an interface. To use the static Connectivity Association Key (CAK) security mode to enable MACsec, you must create and configure connectivity associations on both ends of the link.

A CA is a logical representation of a MACsec domain within a network. Each connectivity association is associated with a CAK. MACsec links are associated with a CA to establish end-to-end MACsec communication. Every MACsec enabled interface is a member of one CA. Switch ports are members of a CA, and can only be a member of one CA.

A secure channel (SC) is a unidirectional channel that connects two endpoints of MACsec. A secure channel is a long-term relationship that persists through the sequence of secure associations.

A secure association (SA) is a short-lived relationship within an SC. MACsec identifies each SA by AN, and supported secure association key (SAK), which is derived from the CAK or generated by the MKA key server. Both ends of the MACsec link use the SAK to encrypt and decrypt the frames. SAKs are frequently refreshed for security reasons. Periodically changing SAs allows the use of fresh keys without terminating the SC relationship.

For static MACsec, you configure CAs. SCs and SAs are internally created in the hardware.

For MKA MACsec, you configure CAs. SCs and SAs are internally created in the hardware based on the information provided by the MKA key server.

MACsec 2AN and 4AN mode

MACsec 2AN mode implementations use two security associations (SA) for each secure channel (SC) and symmetric keys on both MACsec endpoints. The keys are symmetric because they are both derived from the same connectivity association key (CAK).

MACsec 4AN mode generates four Secure Associations Keys (SAK) per secure channel. It uses enhanced hashing algorithm to derive eight SAKs, and uses asymmetric keys on both ends. You can use the **macsec connectivity-association** command to configure different (asymmetric) transmit keys for each endpoint by using the *key-parity* keyword. If you do not specify a value for *key-parity*, the connectivity association is created in 2AN mode. For more information about configuring MACsec transmit keys, see [Configuring a connectivity association](#).

MACsec Components

MACsec has three major components:

- Security entity (SecY)

SecY is the entity that operates the MACsec protocol within the system. You configure a secure connectivity association (CA) to meet the requirements of MACsec for connectivity between stations that attach to an individual LAN. Unidirectional secure channels (SC) support each CA. Each SC supports secure transmission of frames by use of symmetric key cryptography from one of the systems to all others in the CA.

Each SecY transmits frames conveying secure MACsec service requests on a single SC, and receives frames conveying secure service indications on separate SCs, one for each of the other SecYs that participate in the secure CA.

A connectivity association (CA) is a logical representation of a MACsec domain within a network. Each connectivity association is associated with a connectivity association key (CAK). MACsec links are associated with a CA to establish end-to-end MACsec communication. Every MACsec enabled interface is a member of one connectivity association. Switch ports are members of a connectivity association, and can only be a member of one connectivity association.

A secure channel (SC) is a unidirectional channel that connects two endpoints of MACsec. A secure channel is a long-term relationship that persists through the sequence of secure associations. An SC is a unidirectional point-to-multipoint communication, and can persist through Secure Association Key (SAK) changes. A sequence of Secure Associations (SAs) support each SC and allow for the periodic use of fresh keys without terminating the relationship. A single secret key or a set of keys support each SA, where the cryptographic operations used to protect one frame require more than one key. A Secure Channel Identifier (SCI) identifies each SC. An SCI is comprised of a unique 48-bit universally administered MAC address, identifying the system to which the transmitting SecY belongs, concatenated with a 16-bit port number, identifying the SecY within that system.

The SCI concatenated with a two-bit AN identifies each SA. The Secure Association Identifier (SAI) that is created allows the receiving SecY to identify the SA. It also allows the SAK used to decrypt and authenticate the received frame. The AN, and hence the SAI, are unique only for the SAs that can be used or recorded by participating SecYs at any instant.

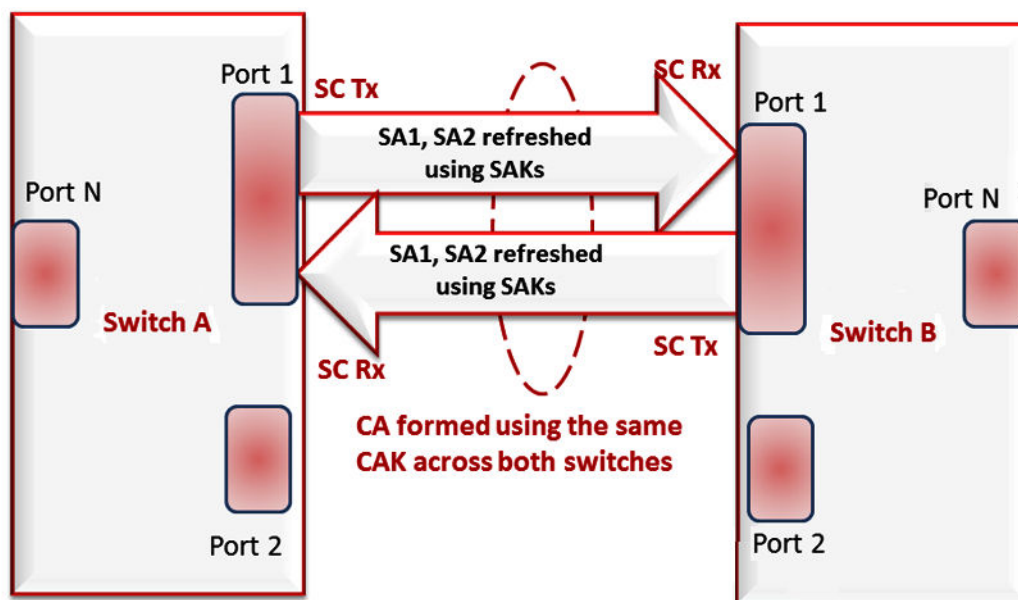


Figure 178: MACsec relationship

- Key agreement entity (KaY)

The KaY in MACsec is responsible for CAK and SAK computations, distributions and maintenance. CAK is a global key that is persistent until the CA exists. When you configure the CAK, ensure that it is identical across MACsec links.

In static SAK mode, SAKs are short-lived keys derived from the CAK, or pre-configured for a particular SC. MACsec uses a timer to refresh SAKs so that the key, as well the session, is secure.


In dynamic SAK mode, the MKA key server generates SAKs. The key server maintains the ethernet link by periodically generating and distributing SAKs across the point-to-point link as long as MACsec is enabled.

- Integrity check verification or Cryptographic entity

The Cryptographic entity provides integrity check protection and validation for frames transmitted or received through the SecY layer. The integrity check verification (ICV) is calculated for the frame source address/destination address (SA/DA), SecTag, User Payload, and Cyclic Redundancy Check (CRC). The calculated ICV is appended at the end-of-frame, recalculated at the receiver side of MACsec link and validated to see if they are equal. The frames that pass the integrity check are further processed, while the system drops the frames that fail the integrity check.

MACsec configuration provides an option to encrypt the user payload. There is also the option to start the encryption from N bytes after the Ethernet header.

In the following figure, CA connects switches A, B, and C by their respective SC and SAK. Switch D cannot participate in the secure communication between A, B, or C as switch D is not configured with the SAK.



MACsec Key Agreement Protocol

MKA protocol performs key server election and generates Secure Association Keys (SAK). SAKs and other MKA information is distributed in MACsec Key Agreement Protocol Data Units (MKPDU) between peers in the Connectivity Association (CA).

Initially you configure pre-shared keys on both ends of an Ethernet link, including values for the CAK and the connectivity association key name (CKN). The CAK and CKN values must match at both ends of the link.

You enable the MKA process by configuring pre-shared keys. MKA performs peer detection, identifies a live peer, and elects the peer with the highest priority as the key server.

The key server generates and distributes SAKs. After the key server and the peer successfully install the generated SAKs, the link can securely transmit encrypted data. The key server maintains the secure link by periodically generating and distributing SAKs for as long as MACsec is enabled.

The following figure illustrates the deployment of MACsec using MKA protocol.

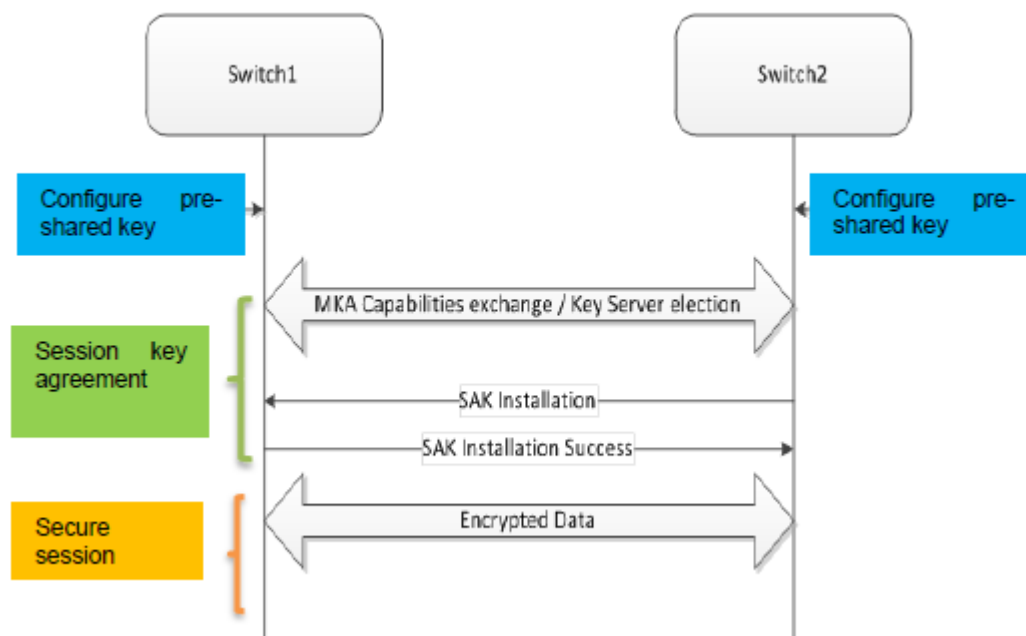


Figure 179: Switch to switch MACsec deployment scenario

You can create and configure an MKA profile and then apply that profile to a port. After applying the profile to the port and associating the port with a connectivity association, you can enable MKA on the port and optionally assign a value for actor priority.



Note

If you enable MKA on a port, traffic is not sent or received on that link until the MKA session is active.

You can configure an MKA actor priority value for each MKA participant. You select priority values from the range 0x00 to 0xff, where lower numbers indicate higher priority. Each participant advertises an actor priority value, and the participant

advertising the highest priority is elected as the key server. If there is a tie for the highest priority, the participant with the highest priority MAC address is selected.

**Note**

Do not configure both peers in an MKA session with an actor priority value of 0xff. If both peers are configured with an actor priority value of 0xff, key server election fails.

You can enable replay protect and configure a replay window size to protect against out-of-sequence packets. Window size specifies the maximum acceptable difference in packet ID numbers between out of order packets. If a packet ID number differs from the ID number of the previously received packet by more than the specified window size, the packet is dropped.

Confidentiality offset specifies the bytes after the Ethernet header from which data encryption begins. Valid values are 30 and 50. Configuring the offset to 30 allows an IPv4 header and TCP/UDP header to remain unencrypted, while configuring the offset to 50 allows an IPv6 header and TCP/UDP header to remain unencrypted.

MKA Interoperability with ExtremeXOS and Switch Engine

Switches configured with MKA MACsec in Fabric Engine 8.6, or later, can interoperate with EXOS 30.3, or later, and Switch Engine 31.6, or later, switches.

Interoperability issues can occur between VOSS/Fabric Engine switches and EXOS/Switch Engine switches when you use MACsec MKA and disable SCI tagging on both ends. Disabling SCI tagging on both ends works for EXOS/Switch Engine if the VOSS/Fabric Engine version is earlier than 8.7.

**Note**

Traffic loss occurs on the ExtremeXOS or Switch Engine to Fabric Engine MKA MACsec link when interoperating with EXOS and Switch Engine, using the Fabric Engine devices as the keyserver. As a best practice, use EXOS or Switch Engine as the keyserver.

MACsec MKA Keychains

Keychains provide an alternative to configuring a single PSK on each MACsec interface. You can use MACsec Key Agreement (MKA) keychains to store and manage multiple pre-shared keys (PSK) used in MACsec encryption. A keychain includes a configured pair of Connectivity Association Name (CKN), Connectivity Association Key (CAK), and a key lifetime. Keychains provide hitless key rollover, which enables the system to switch to the next configured key in the keychain when the key lifetime expires without traffic loss.

Operational Considerations

The following section describes operational considerations for using MKA keychains.

- An MKA keychain supports either AES-128, AES-256, AES-XPB-128, or AES-XPB-256 encryption algorithms. You can configure only one algorithm per keychain. You cannot combine different algorithms in the same keychain.
- Each MKA keychain can contain up to 10 key associations CKN/CAK pairs.
- Ensure there is at least one minute between activating two successive MKA keys.
- Keychain key expiry is expressed in Coordinated Universal Time (UTC) regardless of the switch time-zone.
- When the key expiry is unlimited, the MKA connection automatically switches to the next key. The key expiry for the last key in the keychain is always unlimited.
- You must enable and synchronize Network Time Protocol (NTP) on both endpoints to ensure that keys rollover at the same time. However, you can pre-configure MACsec keychains when the system clock is not synchronized over NTP.

The following warning displays:

```
Switch:1(config-if)#macsecmka keychain KeyChain1
Warning: NTP is not synchronized. Keychain require clock to be synchronized.
Switch:1(config-if)#macsec mka enable
```


MACsec Encryption Cipher Suites

Table 200: MACsec encryption cipher suites product table

Feature	Product	Release introduced
MACsec encryption cipher suites with 32-bit Packet Numbering	4220 Series	Not Applicable
	5320 Series	Fabric Engine 8.6 Both 128 bits and 256 bits Not supported on 5320-24T-4X-XT or 5320-16P-2MXT-2X
	5420 Series	VOSS 8.5 Both 128 bits and 256 bits
	5520 Series	VOSS 8.2.5 Both 128 bits and 256 bits Not supported on 5520-24X and 5520-VIM-4X 5520-24T and 5520-24W on all fixed ports except 1/25 and 1/26 5520-12MW-36W, 5520-48SE, 5520-48T, and 5520-48W on all fixed ports except 1/49 and 1/50
	5720 Series	Fabric Engine 8.7 Both 128 bits and 256 bits 5720-24MW and 5720-24MXW on all fixed ports except 1/25 and 1/26 5720-48MW and 5720-48MXW on all fixed ports except 1/49 and 1/50
	7520 Series	9.0 Supported on 7520-48YE-8CE Both 128 bits and 256 bits
	7720 Series	Not Supported
	7830 Series	Fabric Engine 9.3 Both 128 bits and 256 bits
	VSP 4900 Series	VOSS 8.1 VSP4900-48P supports 128-bit on all fixed ports. VSP4900-24XE supports 128-bit and 256-bit on all fixed ports. VSP4900-12MXU-12XE supports 128-bit and 256-bit on the SFP+ ports.

Table 200: MACsec encryption cipher suites product table (continued)

Feature	Product	Release introduced
MACsec encryption cipher suites with 64-bit Extended Packet Numbering		VIM5-4XE and VIM5-4YE support 128-bit and 256-bit cipher suite.
	VSP 7400 Series	Not Supported
	4220 Series	Not Applicable
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Fabric Engine 9.3 Supported on 5720-VIM-2CE Both 128 bits and 256 bits
	7520 Series	Fabric Engine 9.3 Supported on 7520-48YE-8CE Both 128 bits and 256 bits
	7720 Series	Not Supported
	7830 Series	Fabric Engine 9.3 Both 128 bits and 256 bits
	VSP 4900 Series	Not Supported
	VSP 7400 Series	Not Supported

MACsec cipher suites specify a set of encryption algorithms used to encrypt traffic on an Ethernet link that is secured with Media Access Control Security (MACsec).

Cipher suite	Max key length	Packet Number (PN)
GCM-AES-128	128 bits	32 bits
GCM-AES-256	256 bits	32 bits
GCM-AES-XPB-128	128 bits	64 bits eXtended
GCM-AES-XPB-256	256 bits	64-bits eXtended

The 256-bit algorithm provides enhanced data security and also includes the security provided by the 128-bit algorithm. Both the GCM-AES-128 and GCM-AES-256 cipher suites use a 32-bit packet number (PN) as part of the unique initial value for every packet transmitted with a given secure association key (SAK). The system refreshes the SAK when all the permutations of the 32-bit PN are exhausted.

Additionally, 7830 Series, 5720-VIM-2CE, and 7520-48YE-8CE support GCM-AES-XPB-128 and GCM-AES-XPB-256 cipher suites, which use 64-bit eXtended Packet Numbering (XPB) as part of the unique initial value for every packet transmitted with a given secure

association key (SAK). These cipher suites eliminate the need for frequent SAK rekeying that can occur in high speed links.



Note

GCM-AES-XPB-128 and GCM-AES-XPB-256 cipher suites are supported with MACsec Key Agreement.

The default cipher suite is the GCM-AES-128.



Note

Not all products support both a 128-bit cipher suite and a 256-bit cipher suite. For information about product support, see [Table 200](#) on page 2611.

You typically configure a MACsec cipher suite at the port level on the switch. The configuration is optional. When you configure a cipher suite, ensure that you configure the same cipher suite on both MACsec peers.

5320 Series and 5420 Series Encryption and Decryption

These products limit encryption and decryption bandwidth. 5420 Series supports a maximum of 50 Gbps. The 5320 Series models support differing encryption and decryption rates, as identified in the following table.

Table 201: 5320 Series MACsec encryption and decryption rates

Model	Encryption rate	Decryption rate
48-port models	Up to 50 Gbps	Up to 50 Gbps
24-port models	Up to 25 Gbps	Up to 25 Gbps
16-port models	Up to 25 Gbps	Up to 25 Gbps

MACsec Operation

As shown in the following figure, a host that connects to Switch A sends an Ethernet frame to a host that connects to Switch B. Switch A encrypts the frame, excluding the Ethernet header and optionally the 802.1Q header. Switch A also appends MACsec information, such as SecTag and ICV to the encrypted payload and transmits the frame using normal frame transmission. This process ensures data confidentiality.

Switch B decrypts the frame. Switch B recalculates the ICV using a MACsec key and the SecTag present in the frame. If the ICV present in the received frame matches the recalculated ICV, the switch processes the frame. If the two ICVs do not match, the switch discards the frame. This process ensures data origin authenticity and data integrity.

The encryption and decryption algorithms follow either the AES-GCM-128 standard or the AES-GCM-256 standard depending on the configured cipher suite. The default is the AES-GCM-128 standard.

The MACsec connectivity association key (CAK) between switches A and B are statically pre-configured.

MACsec operates between two switches across Point-to-Point Connectivity only when the switches are either directly connected or across a network cloud that provides P2P connectivity between the two switches.

For example, in the following figure you can enable MACsec between two switches across a network cloud where P2P connectivity between the switches is provided using services such as P2P, MPLS, Layer 2 VPN (ELINE), or connectivity across Dark Fiber.




Figure 180: MACsec operation

Automatic QoS Priority for MACsec Packets on Intermediate Switches

In certain situations where MACsec encrypted packets traverse intermediate non-MAC switches, QoS visibility is lost.

This feature uses confidentiality-offset to specify that the first 30 to 50 bytes within the MACsec frame transmit without encryption, thus leaving the 802.1Q VLAN tag p-bits in the clear so that the intermediate switch can differentiate between encrypted traffic. With the 802.1Q p-bits in the clear, internal QoS priority for MACsec packets on intermediate switches can be automatically assigned.



Note

This feature is always enabled.

Operational Considerations

The following list describes the operational considerations for automatic QoS prioritization for MACsec packets on intermediate switches:

- The MACsec port on the end point switch must have confidentiality offset configured to leave the 802.1Q VLAN tag p-bits in the clear.
- The intermediate switch port must be a Layer 2 trusted port.

- If the port on an intermediate switch is an untrusted port or if the packet is not originally VLAN tagged, the switch assigns MACsec packets to the default QoS level.
- ACL QoS rules can override the automatic QoS prioritization for MACsec packets.
- Internal QoS level assignment is on ingress. The egress queue and remarking the packet is derived from the internal QoS level.
- The priority extraction is performed from the original VLAN Tag only.
- If a switch is the endpoint of a MACsec flow, it decrypts the packets and uses the standard priority assignment.

MACsec Configuration Using CLI

This section provides procedures to configure MACsec using the command line interface (CLI).

Before You Begin

Product Notice: Before you enable MACsec on the 5320 Series and 5420 Series, you must enable the boot configuration flag: **boot config flags macsec**. The **boot config flags ipv6-egress-filter** and **boot config flags macsec** commands are mutually exclusive.

Configure MACsec Encryption on a Port

Use the following procedure to enable or disable encryption on a MACsec capable port. The default is disabled.

About This Task

If you disable encryption, MACsec forwards traffic in clear text. You can view that data that is not encrypted in the Ethernet frame that travels across the link. Even if you disable encryption the MACsec header applies to the frame and integrity checks make sure that traffic has not been tampered with.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure MACsec encryption on the port using one of the following commands:

- To enable MACsec, use `macsec encryption enable`
- To disable MACsec, use `no macsec encryption enable`

Example

Configure MACsec encryption on a port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabit 1/2
Switch:1(config-if)#macsec encryption enable
```

Configure a MACsec Cipher Suite on a Port

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format `slot/port/sub-port`.

2. Configure a MACsec encryption cipher suite:

```
macsec cipher-suite {gcm-aes-128 | gcm-aes-256 | gcm-aes-xpn-128 |
gcm-aes-xpn-256}
```

The default cipher suite is GCM-AES-128.



Note

`gcm-aes-xpn-128` and `gcm-aes-xpn-256` ciphers are supported with MKA only.

Ensure that you configure the same cipher suite on both MACsec peers.

3. (Optional) [Associate an MKA Profile to a Port](#) on page 2630.

4. Verify the configuration:

```
show macsec status {slot/port[/sub-port] [-slot/port[/sub-port]][,...]
[slot/all][all]}
```

Example

Configure the 256-bit MACsec cipher suite on the port 1/3 and verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config)#interface gigabitEthernet 1/3
Switch:1(config-if)#macsec cipher-suite gcm-aes-256
Switch:1#show macsec status 1/3
```

```
=====
MACSEC Port Status
=====
MACSEC Encryption Replay  Replay  Encryption Cipher CA  MKA-Profile MKA
PortId Status Status   Protect Protect
                        W'dow  Offset    Suite  Name  Name      Connect
-----
1/3   enabled disabled enabled  50 ipv4Offset(30) AES-256 mkanka extreme  pending
```

The system displays the following error message if you attempt to configure a cipher suite on a port that is not MACsec capable.

```
Switch:1>enable
Switch:1(config)#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/2
Switch:1(config-if)#macsec cipher suite gcm-aes-256

Error: port 1/2, Port is not MACSec capable. No MACSec configurations allowed on port
```

The system displays the following error message if your hardware does not support the MACsec 256-bit cipher suite.

Configure an Extended Packet Numbering (XPN) cipher with MKA on a port:

```
Switch:1>enable
Switch:1(config)#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#macsec cipher suite gcm-aes-xpn-256
Switch:1(config-if)#macsec mka profile p1
NOTE: Both ends of the MKA Macsec link must have the same Include-SCI value configured.
Switch:1(config-if)#macsec mka enable
```

Display the MACsec MKA profile:

```
Switch:1(config-if)#show macsec mka profile
```

```
=====
MACsec MKA Profile
=====
Profile      Profile  Cipher      Confidentiality  Replay  Window  Port
Include-SCI  Id       Suite      Offset          Protect  Size
-----
p1           1        gcm-aes-xpn-256 0          Disabled  0        1/1
true
-----
```

Variable Definitions

The following table defines parameters for the **macsec cipher-suite** command.

Variable	Definition
<code>{gcm-aes-128 gcm-aes-256 gcm-aes-xpn-128 gcm-aes- xpn-256}</code> Note: Exception: gcm-aes-xpn-128 and gcm-aes-xpn-256 ciphers are supported only on 7830 Series, 7520-48YE-8CE, and 5720- VIM-2CE with MKA.	Configures the cipher suite for encrypting traffic with MACsec containing a 32-bit packet number (PN) or 64-bit extended packet number (XPN). The supported cipher suites are: <ul style="list-style-type: none"> • gcm-aes-128 — maximum key length of 128 bits • gcm-aes-256 — maximum key length of 256 bits • gcm-aes-xpn-128 — maximum key length of 128 bits • gcm-aes-xpn-256 — maximum key length of 256 bits The default is the gcm-aes-128 cipher suite.

Configure a Connectivity Association

Use the following procedure to configure a connectivity association (CA) in static Connectivity Association Key (CAK) security mode with static Secure Association Keys (SAK).



Important

For static MACsec, you can configure a different connectivity association name for local and peer nodes but you must configure the same value for the connectivity association key at both ends of the link with, either even or odd mode.

For MKA MACsec, you must configure the same value for the connectivity association name and the connectivity association key for local and peer nodes. Even or odd mode does not apply to a connectivity association for MKA MACsec .

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a CA:

```
macsec connectivity association WORD <5-32> connectivity-association-  
key WORD<10-64> [key-parity even | odd]
```



Note

If you do not specify a key-parity value, the CA is created in 2AN mode. This applies only to platforms that support 2AN mode.

Example

Configure a connectivity association and enable MACsec on a port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#macsec connectivity-association canamel connectivity-association-key
1029384756abcdef key-parity even
```

Variable Definitions

The following table defines parameters for the **macsec** command.

Variable	Value
<i>connectivity-association WORD<5-32></i>	Specifies the connectivity-association name as an alpha-numeric ASCII string up to 32 characters long. The device uses this value for the connectivity-association key name (CKN).
<i>connectivity-association-key WORD<10-64></i>	Specifies the connectivity-association key (CAK) value as a 32-character (128-bit) or a 64 character (256-bit) hexadecimal string. Note: Always select the 128-bit CAK value for AES-GSM-128 and the 256-bit CAK value for AES-GSM_256.
<i>key-parity <even odd></i> Note: This parameter only applies to static MACsec configurations.	Specifies Tx key parity using the following values: <ul style="list-style-type: none"> • even — generates even-numbered keys for Tx • odd — generates odd-numbered keys for Tx Note: If you do not specify a key-parity value, the connectivity association (CA) is created in 2AN mode. This only applies to platforms that support 2AN mode.

The following table defines parameters for the **interface gigabitethernet** command.

Variable	Value
<i>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</i>	Specifies the port that you want to associate with the CA. Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Associate a Port with a Connectivity Association

About This Task

Perform the following procedure to associate a port with a connectivity association (CA).

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Associate a port with a CA:

```
macsec connectivity-association WORD<5-32>
```

Example

```
Switch:1(config)#interface gigabitethernet 1/2  
Switch:1(config-if)#macsec connectivity-association caname12
```

Variable Definitions

The following table defines parameters for the **macsec** command.

Variable	Value
<i>connectivity-association</i> <i>WORD<5-32></i>	Specifies the connectivity-association name as an alphanumeric ASCII string up to 32 characters long. The device uses this value for the connectivity-association key name (CKN).

Update the Connectivity Association for a port

Use the following procedure to change the Connectivity Association (CA) to which a port is associated.

Before You Begin

Ensure that the new CA is created at the global level. For information about configuring a CA, see [Configure a Connectivity Association](#) on page 2618.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable MACsec on the port:

```
no macsec enable
```

3. Change the CA to which the port is associated:

```
macsec connectivity-association WORD<5-32>
```

4. Enable MACsec on the port:

```
macsec enable
```

Example

Change the CA to which a port is associated:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabit 1/2
Switch:1(config-if)#no macsec enable
Switch:1(config-if)#macsec connectivity-association caname12
Switch:1(config-if)#macsec enable
```

Variable Definitions

The following table defines parameters for the **macsec** command.

Variable	Value
<i>connectivity-association</i> <i>WORD<5-32></i>	<p>Specifies the connectivity-association name as an alpha-numeric ASCII string up to 16 characters long. The device uses this value for the connectivity-association key name (CKN).</p> <p>Tip: Configure the CKN in multiples of 4 characters to avoid MKA interoperability issues between Fabric Engine switches and EXOS or Switch Engine switches. For example, Macsecma (8 characters) or Macsecmka123 (12 characters) are valid, but Macsec (6 characters) is not valid.</p>

The following table defines parameters for the **interface gigabitethernet** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Specifies the port to associate with the connectivity association (CA).</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View MACsec Status

Perform this procedure to view MACsec status.

About This Task

This command displays the status for the following:

- MACsec status
- MACsec encryption status
- MACsec encryption cipher suite
- The associated Connectivity Association (CA) name



Note

If you do not specify a port number, the information on all MACsec capable interfaces is displayed.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. View the MACsec status:

```
show macsec status {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]  
[slot/all][all]}
```

3. Display all MACsec related information:

```
show macsec
```

Examples



Note

The switch does not support replay protect.

View MACsec status:

```
Switch:1>show macsec status
```

MACSEC Port Status									
PortId	MACSEC Status	Encryption Status	Replay Protect	Replay Protect W'dow	Encryption Offset	Cipher Suite	CA Name	MKA-Profile Name	MKA Connect Status
1/1	enabled	enabled	enabled	--	ipv4Offset(30)	AES-128	KeyNumber1	profileMKA	secure
1/13	disabled	disabled	disabled	--	none	AES-128	NIL	--	--
1/14	disabled	disabled	disabled	--	none	AES-128	NIL	--	--
1/15	enabled	disabled	enabled	50	ipv4Offset(30)	AES-256	mkanka	extreme	pending

View MACsec status on port 1/13:

```
Switch:1>show macsec status 1/13
```

MACSEC Port Status									
PortId	MACSEC Status	Encryption Status	Replay Protect	Replay Protect W'dow	Encryption Offset	Cipher Suite	CA Name	MKA-Profile Name	MKA Connect Status
1/13	enabled	disabled	enabled	50	ipv4Offset(30)	AES-256	mkanka	extreme	pending

View the MACsec Connectivity Association Details

Perform this procedure to view the MACsec connectivity association (CA) details.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the MACsec CA details:

```
show macsec connectivity-association [WORD<5-32]
```



Note

This command displays the MACsec CA details, including the hashed value of the CA key.

Example

View the MACsec connectivity association details:

```
Switch:1>show macsec connectivity-association
```

MACSEC Connectivity Associations Info			
Connectivity Association Name	SHA-1 Connectivity Association Key Hash	AN_Mode / TxKeyParity	Port Members
caname50	ba6b005bef79e7b95f3e08181e2501ce	2AN / NA	1/49
caname51	5b41f44ecaa54f3873e781557b39230b	4AN / odd	
caname15	053f26fb96b011191f2da28849f08677	4AN / Even	1/50

Configure Static MACsec using CLI

Use the following procedures to configure static MACsec using the Command Line Interface (CLI).

Enable Static MACsec on a Port

Use the following procedure to enable MACsec with static Secure Association Keys (SAK) on a port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
- enable
- configure terminal
- interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}



Note If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable MACsec on the port:
- macsec enable

Variable Definitions

The following table defines parameters for the **macsec enable** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configure the Confidentiality Offset on a Port

Use the following procedure to configure the confidentiality offset on a port. The default is disabled.

About This Task

The confidentiality offset provides a way to start encryption after a few bytes following the Ethernet header. The confidentiality offset facilitates traffic flow inspection and classification on intermediate devices by not encrypting the Network Layer header for IPv4 or IPv6. For instance, if you configure the offset to 30, the IPv4 header and the

TCP/UDP header are not encrypted. If you configure the offset to 50, the IPv6 header and the TCP/UDP header are not encrypted.



Note

On a MACsec-enabled port with confidentiality offset configured to 50 on the 5320 Series or 5420 Series, all packets less than 67 bytes drop and discarded packets increment.

As a best practice, do not configure the confidentiality offset to 50 on the 5320 Series or 5420 Series.

On a MACsec-enabled port with data encryption enabled and confidentiality offset configured to 30 or 50 on the 5320 Series or 5420 Series, InOctetsValidated counter also increments in addition to InOctetsDecrypted counters in Macsec secure channel Inbound statistics.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure confidentiality offset on the port using one of the following commands:

- To enable confidentiality-offset, use `macsec confidentiality-offset <30-50>`
- To disable confidentiality-offset, use `no macsec confidentiality-offset`

Example

Configuring the confidentiality offset on the port:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabit 1/2  
Switch:1(config-if)#macsec confidentiality-offset 30
```

Variable Definitions

The following table defines parameters for the **macsec confidentiality-offset** command.

Variable	Value
<30-50>	Specifies the bytes after the Ethernet header from which data encryption begins. Valid values are 30 and 50.

The following table defines parameters for the **interface gigabitethernet** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Specifies the port that you want to associate with the connectivity association (CA).</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure MKA MACsec using CLI

Use the following procedures to configure MACsec Key Agreement (MKA) MACsec using the Command Line Interface (CLI).

Perform the following procedures in order to configure MKA MACsec.

- [Configure MKA Keychains](#) on page 2626
- [Configure an MKA Profile](#) on page 2628
- [Associate an MKA Profile to a Port](#) on page 2630
- [Associate an MKA Keychain to a Port](#) on page 2631
- [Enable MKA on a Port](#) on page 2632

If you apply a Connectivity Association to a port before you apply an MKA profile to the port, the port is considered to be configured with static MACsec and an error occurs.

Configure MKA Keychains

About This Task

Perform this procedure to configure MKA keychains.



Note

Creating an MKA keychain changes the CLI command mode to mka-keychain mode.

Procedure

1. Enter mka-keychain Configuration mode:

```
enable
```

```
configure terminal
```

```
macsec mka keychain WORD<5-32>
```


2. Configure keys in the MKA keychain:

```
key WORD<5-32> {cak WORD<10-64> | endtime <MMddyyyyhhmmss |
unlimited>}
```

**Note**

The CKN/CAK pair must match between devices.

3. Exit from macsec mka keychain mode:

```
exit
```

4. [Configure an MKA Profile](#) on page 2628.5. [Associate an MKA Profile to a Port](#) on page 2630.6. [Associate an MKA Keychain to a Port](#) on page 2631.

7. Exit to Global Configuration mode:

```
exit
```

8. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

9. [Enable MKA on a Port](#) on page 2632.

10. Display MACsec MKA keychain status on the port:

```
show macsec mka keychain
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#macsec mka keychain KeyChain1
2025-03-14T15:50:21.813Z 5720-24MXW-FabricEngine CP1 - 0x0034c630 - 00000000 GlobalRouter
MACSEC INFO MKA Keychain KeyChain1 created successfully
Switch:1(mka-keychain)#key KeyNumber1 cak 123456789abcdef endtime 08102026124500
Switch:1(mka-keychain)#key KeyNumber2 cak 987654321abc123 endtime unlimited
Switch:1(mka-keychain)#exit
Switch:1(config)#macsec mka profile profileMKA
Switch:1(config)#confidentiality-offset 30
Switch:1(config)#include-sci enable
Switch:1(config)#exit
Switch:1(config)#interface gigabitethernet 1/1
Switch:1(config-if)#macsec mka profile profileMKA
Switch:1(config-if)#macsec mka keychain KeyChain1
Switch:1(config-if)#macsec mka enable
```

The following example displays MACsec MKA key status.

```
Switch:1#show macsec mka keychain
=====
MACSEC MKA Keychain Info
=====
```

Connectivity	Configured	Port	Key Name (CKN)	Key	Key
Keychain name	Keys number	Members	Key SHA-1 Hash (CAK SHA-1 Hash)	Expiry (UTC)	Status
Keychain1	3	1/1	KeyNumber1 08e621a7fb61258e130fb626d766278772d82688	Unlimited	Valid

	KeyNumber2	07/18/2025 07:28:00	Expired
	62c07c431fe79e9bdf6e443caddb4b9d8a8c5fc9		
1/4-1/5	KeyNumber3	07/18/2028 12:12:12	In-use
	0e98c2db829c2cb599eedfbdec01d6c6efeacd7		
All 1 out of 1 Total Num of Macsec mka keychains displayed			

When the key reaches the configured expiry, the MKA connection automatically switches to the next key:

```
Switch CP1 - 0x0034c634 - 00000000 GlobalRouter MACSEC INFO MACSec MKA: Participant
switched to the next keychain key on port 1/1
```

Variable Definitions

The following table defines parameters for the **key** command.

Variable	Value
<i>cak</i> WORD<10-64>	Specifies the Connectivity Association key value.
<i>endtime</i> <MMddyyyyhhmmss> <i>unlimited</i>	Specifies the expiry for the MKA keychain key in Coordinated Universal Time (UTC) regardless of the switch time-zone.
WORD<5-32>	Specifies the Connectivity Association key name.
<i>keychain</i> WORD<5-32>	Specifies the name applied to the keychain.

Configure an MKA Profile

About This Task

Perform this procedure to create a MACsec Key Agreement (MKA) profile. Creating an MKA profile changes the CLI command mode to mka profile mode.

You can also use this procedure to enter mka profile CLI command mode for an existing MKA profile.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Create an MKA profile name:
macsec mka profile WORD<1-16>



Note

An MKA profile name consists only of alphanumeric characters (0-9, A-Z, and a-z). The profile name is case sensitive.

3. Configure the MKA profile using one or more of the following options:

- Enable replay protection and configure the window size: `replay-protect enable window-size <5-500>`



Note

The configuration should be the same at both ends of the link, either enabled or disabled.

- Configure a value for the confidentiality offset to specify the number of unencrypted bytes that precede MACsec encryption: `confidentiality-offset <30-50>`
- Enable SCI Tagging for a MACsec-enabled switch:



Note

If you are running VOSS 8.4.2 or earlier, disable SCI tagging on the switch.

```
include-sci enable
```



Note

The configuration must be the same at both ends of the link, either enabled or disabled.

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#macsec mka profile profileMKA
Switch:1(mka-profile)#confidentiality-offset 30
Switch:1(mka-profile)#include-sci enable
```

Variable Definitions

The following table defines parameters for the **macsec mka profile** command.

Variable	Value
<code>window-size <5-500></code>	Specifies the maximum acceptable difference in packet ID numbers between out of order packets. If a packet ID number differs from the ID number of the previously received packet by more than the specified window size, the packet is dropped.
<code>WORD<1-16></code>	Specifies the MKA profile name. An MKA profile name consists only of alphanumeric characters (0-9, A-Z, and a-z). The profile name is case sensitive.
<code><30 50></code>	Specifies the number of bytes after the Ethernet header from which data encryption begins. Possible values are 30 (IPv4 plus TCP/UDP header) and 50 (IPv6 plus TCP/UDP header). The default is no offset.

*Associate an MKA Profile to a Port***Before You Begin**

You must create an MKA profile globally before you can associate it to a port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Apply an MKA profile to the port:

```
macsec mka profile WORD<1-16>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/1
Switch:1(config-if)#macsec mka profile profileMKA
```

Variable Definitions

The following table defines parameters for the **macsec mka profile** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>WORD<1-16></code>	Specifies the MKA profile name. An MKA profile name consists only of alphanumeric characters (0-9, A-Z, and a-z). The profile name is case sensitive.

Associate an MKA Keychain to a Port

About This Task

Perform this procedure to associate an MKA keychain to a port.



Note

You can update keys end time, delete keys, and add new keys while the keychain is associated to a port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Associate an MKA keychain to a port:

```
macsec mka keychain WORD<5-32>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/1
Switch:1(config-if)#macsec mka keychain KeyChain1
```

Variable Definitions

The following table defines parameters for the **macsec mka keychain** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>WORD<5-32></code>	Specifies the name applied to the keychain.

*Enable MKA on a Port***Before You Begin**

- Apply an MKA profile to the port.
- Before you enable MKA on a port, the port must be a member of a connectivity association.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable MKA on a port:

```
macsec mka enable
```

Example

```
Switch:1#enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#macsec mka enable
```

Variable Definitions

The following table defines parameters for the **macsec mka enable** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

*Configure MKA Actor Priority***Before You Begin**

- Apply an MKA profile to the port.
- Disable MKA on the port before you configure a value for actor priority. Enable MKA on the port after you configure an actor priority value.

About This Task

The MKA participant with the highest actor priority is designated as the key server.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
- enable
- configure terminal
- interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}



Note
If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure a value for actor priority:
- macsec actor-priority <0x00-0xff>

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitethernet 1/4
Switch:1(config-if)#macsec actor-priority 0x0f
```

Variable Definitions

The following table defines parameters for the **macsec actor-priority** command.

Variable	Value
<0x00-0xff>	Specifies a hexadecimal value for actor priority, which determines key server selection. Lower values indicate a higher priority. The default is 10.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Display MKA Profiles

About This Task

Use the following procedure to display information about all MKA profiles configured on the switch. Optionally, you can view information about a specific MKA profile.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all MKA profiles on the switch:

```
show macsec mka profile
```
3. (Optional) Display a specific MKA profile:

```
show macsec mka profile WORD<1-16>
```

Example

The following example displays MACsec MKA profile information:

The following example displays MACsec MKA information for a specific profile.

```
Switch:1>show macsec mka profile profileMKA
```

MACsec MKA Profile							
Profile Name	Profile Id	Cipher Suite	Confidentiality Offset	Replay Protect	Window Size	Port	Include-SCI
profileMKA	4	gcm-aes-128	30	Disabled	0	1/1	true

Variable Definitions

The following table defines parameters for the **show macsec mka profile** command.

Variable	Value
WORD<1-16>	Specifies the MKA profile name. An MKA profile name can consist only of alphanumeric characters (0-9, A-Z, and a-z). The profile name is case sensitive.

Display MKA Keychain Associations

About This Task

Perform the following procedure to display information about all MKA keychain associations configured on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display MKA keychain associations configured on the switch:

```
show macsec mka keychain [WORD<5-32>]
```

Example

The following example displays MKA keychains configured on the switch:

```
Switch:1#show macsec mka keychain
```

MACSEC MKA Keychain Info					
Connectivity	Configured	Port	Key Name (CRN)	Key	Key
Keychain name	Keys number	Members	Key SHA-1 Hash (CAK SHA-1 Hash)	Expiry (UTC)	Status
Keychain1	3	1/1	KeyNumber1 08e621a7fb61258e130fb626d766278772d82688	Unlimited	Valid

	KeyNumber2	07/18/2025 07:28:00	Expired
	62c07c431fe79e9bdf6e443caddb4b9d8a8c5fc9		
1/4-1/5	KeyNumber3	07/18/2028 12:12:12	In-use
	0e98c2db829c2cb599eedfbdec01d6c6efeacd7		

All 1 out of 1 Total Num of Macsec mka keychains displayed

Display MKA Participants

About This Task

Use the following procedure to display information about all participants in an MKA session. You can also display MKA information for a specific port participating in an MKA session.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about all participants in an MKA session:

```
show macsec mka participant
```
3. (Optional) Display MKA information for a specific port participating in an MKA session:

```
show macsec mka participant [verbose] {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]
```

Examples

The following example displays information for all participants in an MKA session.

```
Switch:1>show macsec mka participant
```

MACsec MKA Participants				
Port Id	CA Name	MKA-Profile Name	MKA Enable	Actor Priority
1/3	CA120022	extreme030519	Enabled	A
1/4	CA121023	extreme031519	Enabled	14
2/2	CA122024	extreme032019	Enabled	1E

The following example displays information for a specific port participating in an MKA session.

```
Switch:1>show macsec mka participant 1/3
```

MACsec MKA Participant				
Port Id	CA Name	MKA-Profile Name	MKA Enable	Actor Priority
1/3	CA120022	extreme030519	Enabled	A

Variable Definitions

The following table defines parameters for the **show macsec mka participant** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
verbose	Displays detailed information for a specific port participating in an MKA session.

MACsec Configuration using EDM

Use the following procedures to configure MACsec in EDM.

Before You Begin

Product Notice: Before you enable MACsec on the 5320 Series and 5420 Series, you must enable the boot configuration flag: **Configuration > Edit > Chassis > Boot Config. EnableIpv6EgressFilterMode** and **EnableMacsec** are mutually exclusive.

Configure Connectivity Associations

About This Task

Use the following procedure to configure connectivity associations (CA) using EDM.



Note

- You can configure MACsec on physical ports *only*. However, the physical ports can belong to an MLT trunk group that includes: Split MultiLink Trunking (SMLT), distributed MultiLink Trunking (DMLT), or Link aggregate group (LAG).
- MACsec encryption algorithms follow AES-GCM-128, AES-GCM-256, AES-GCM-XPB-128, or AES-GCM-XPB-256 standard depending on the configured MACsec cipher suite. The default is the AES-GCM-128 standard.



Note

AES-GCM-XPB-128 and AES-GCM-XPB-256 standards are supported only on 7830 Series, 7520-48YE-8CE, and 5720-VIM-2CE and with MKA only.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **MACSec**.
3. Select the **Associations** tab.
4. Select **Insert**.
5. In **AssociationName**, type the connectivity-association name.
6. In **AssociationKey**, type the value of the connectivity-association key.

**Note**

The system displays the connectivity-association key as an MD5-hashed text in the MAC security table.

7. In **AssociationTxKeyParity**, select an option for Tx key parity.

**Note**

Tx key parity configuration applies only to static MACsec configurations.

8. Select **Insert**.
9. Select **Apply**.

Associations Field Descriptions

Use the data in the following table to use the **Associations** tab.

Name	Description
AssociationName	Specifies the connectivity-association name as an alpha-numeric ASCII string up to 32 characters long. The device uses this value for the connectivity-association key name (CKN).
AssociationKey	<p>Specifies the connectivity-association key (CAK) value as a 32-character (128-bit) or a 64 character (256-bit) hexadecimal string.</p> <p>Note: Always select the 128-bit CAK value for AES-GSM-128 and the 256-bit CAK value for AES-GSM_256.</p>

Name	Description
AssociationPortMembers	Specifies the set of ports for which this connectivity association is associated.
AssociationTxKeyParity	<p>Specifies Tx key parity using the following values:</p> <ul style="list-style-type: none"> • None — key parity is not specified <p>Note: The none value only applies to platforms that support 2AN mode. If you do not specify a key parity value, the system defaults to 2AN mode.</p> <ul style="list-style-type: none"> • Even — generates even-numbered keys • Odd — generates odd-numbered keys

Associate a Port with a Connectivity Association

Use the following procedure to associate a port with a connectivity association (CA) using EDM. You can optionally configure a MACsec encryption cipher suite on the port.



Note

- You can configure MACsec on physical ports *only*. However, the physical ports can belong to an MLT trunk group that includes: Split MultiLink Trunking (SMLT), distributed MultiLink Trunking (DMLT), or Link aggregate group (LAG).
- MACsec encryption algorithms follow AES-GCM-128, AES-GCM-256, AES-GCM-XPB-128, or AES-GCM-XPB-256 standard depending on the configured MACsec cipher suite. The default is the AES-GCM-128 standard.



Note

AES-GCM-XPB-128 and AES-GCM-XPB-256 standards are supported with only on 7830 Series, 7520-48YE-8CE, and 5720-VIM-2CE and with MKA only.

Procedure

1. On the **Device Physical View** tab, select on one or more ports to associate with the connectivity association.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **MACsec** tab.
5. In **CAName**, type the connectivity-association name.
6. In **OffsetValue**, select the value of confidentiality offset to be achieved.
7. Select **EncryptionEnable** to enable encryption for the frames transmitted on the port.
8. (Optional) In **CipherSuite**, select the MACsec encryption cipher suite.

9. Select **MACsec Enable** to enable MACsec on the port.
10. Select **Apply**.

MACsec Field Descriptions

Use the data in the following table to configure the **MACsec** tab.

Name	Description
CAName	Specifies the name of the connectivity association attached to the port or interface.
OffsetValue	<p>Offsets MACsec encryption in an IPv4 TCP/UDP header or IPv6 TCP/UDP header. The confidentiality offset provides a way to start encryption after a few bytes following the Ethernet header. The confidentiality offset facilitates traffic flow inspection and classification on intermediate devices by not encrypting the Network Layer header for IPv4 or IPv6. For instance, if you configure the offset to 30, the IPv4 header and the TCP/UDP header are not encrypted. If you configure the offset to 50, the IPv6 header and the TCP/UDP header is not encrypted.</p> <p>Note: On a MACsec-enabled port with confidentiality offset configured to 50 on the 5320 Series or 5420 Series, all packets less than 67 bytes drop and discarded packets increment. As a best practice, do not configure the confidentiality offset to 50 on the 5320 Series or 5420 Series.</p> <p>Note: On a MACsec-enabled port with data encryption enabled and confidentiality offset configured to 30 or 50 on the 5320 Series or 5420 Series, InOctetsValidated counter also increments in addition to InOctetsDecrypted counters in Macsec secure channel Inbound statistics.</p>
EncryptionEnable	Specifies the encryption status per port. Use this field to enable or disable encryption for each MACsec capable port.

Name	Description
CipherSuite Note: Exception: AES-GCM-XPB-128 and AES-GCM-XPB-256 ciphers are supported only on 7830 Series, 7520-48YE-8CE, and 5720-VIM-2CE and with MKA only.	Configures the cipher suite for encrypting traffic with MACsec. The following cipher suites are supported: <ul style="list-style-type: none"> • aes-gcm-128—maximum key length of 128 bits • aes-gcm-256—maximum key length of 256 bits • aes-gcm-xpn-128—maximum key length of 128 bits • aes-gcm-xpn-256—maximum key length of 256 bits The default is the AES-GCM-128 standard.
MACsec Enable	Enables or disables MACsec on the port.

Configure MKA MACsec using EDM

Use the following procedures to configure MKA MACsec using EDM.



Note

You must configure MKA MACsec by performing procedures in the following order:

1. [Configure MKA Keychains](#) on page 2640
2. [Configure an MKA Profile](#) on page 2641
3. [Apply an MKA Profile to a Port](#) on page 2642
4. [Configure an MKA Keychain Association](#) on page 2643
5. [Enable MKA on a Port](#) on page 2632

Configure MKA Keychains

About This Task

Perform this procedure to create an MKA keychain.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **MACsec**.
3. Select the **MACsec KA** tab.
4. Select **Insert**.
5. In **Id**, type an MKA keychain ID value.
6. In **Name**, type a keychain name.
7. Select **Insert**.

MACsec KA Field Definitions

Use the data in the following table to use the **MACsec KA** tab.

Name	Description
Id	Specifies a unique identification number for the MKA keychain configured on the switch. You can configure up to 512 MKA keychains.
Name	Specifies the name applied to each keychain association configured on the switch.
PortMembers	Specifies the port or ports associated with the keychain on the switch.
KeyNum	Specifies keychain association key number.
KeysCount	Specifies the number of keychain association keys configured on the switch. You can configure up to 10 MKA key associations per keychain.

Configure an MKA Profile

Procedure

1. In the navigation pane, expand **Security > Data Path** .
2. Select **MACSec**.
3. Select **MKA Profile**.
4. Select **Insert**.
5. For **Id**, type an MKA ID value.
6. For **Name**, type the profile name.
7. (Optional) Select **ReplayProtectEnable** to enable replay protect.
8. (Optional) For **ReplayProtectWindow**, type a value for the replay protect window size.
9. (Optional) From the **OffsetValue** options, select a level for confidentiality offset.
10. Select **IncludeSCIEnable** to enable the SCI field in MACsec frames.
11. Select **Insert**.

MKA Profile Field Descriptions

Use the data in the following table to use the **MKA Profile** tab.

Name	Description
Id	Specifies a unique identification number for an MKA profile.
Name	Specifies the profile name.
ReplayProtectEnable	Specifies whether replay protect is enabled. The default is disabled.

Name	Description
ReplayProtectWindow	Specifies the maximum acceptable difference in packet ID numbers between out of order packets. If a packet ID number differs from the ID number of the previously received packet by more than the specified window size, it is dropped.
OffsetValue	Specifies the number of bytes after the Ethernet header from which data encryption begins. The default is no offset.
PortMembers	Specifies the ports that are members of an MKA profile.
CipherSuite	Displays the cipher suite used to encrypt traffic with MACsec.
IncludeSCIEnable	Specifies whether SCI tagging is enabled for a MACsec-enabled switch. The default is disabled.

Apply an MKA Profile to a Port

Before You Begin

You must configure an MKA profile globally before you can apply it to a port.

About This Task

Use the following procedure to apply a MACsec Key Agreement (MKA) profile on a port.

Procedure

1. In the Device Physical View, select the port to which you want to apply an MKA profile.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **MACsec MKA** tab.
5. For **MACsecMKAProfileName**, select the name of the MKA profile you want to apply.
6. Select **Apply**.

What to Do Next

[Enable MKA on a Port](#) on page 2644

MACsec MKA Field Descriptions

Use the data in the following table to use the **MACsec MKA** tab.

Name	Description
MACsecMKAProfileName	Specifies the MKA profile name.
CAName	Specifies the name of the connectivity association attached to the port or interface.

Name	Description
OffsetValue	Specifies the number of bytes after the Ethernet header from which data encryption begins. The default is no offset.
CipherSuite	Displays the cipher suite used to encrypt traffic with MACsec.
Keychain Name	Specifies the name associated with the keychain.
MKA Enable	Enables the MKA profile on the port. The default is disabled.

Configure an MKA Keychain Association

About This Task

Perform this procedure to configure a MKA keychain association.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **MACsec**.
3. Select the **MACsec KA Key** tab.
4. Select **Insert**.
5. In **Id**, type the ID associated with the keychain.
6. In **KeyId**, type the keychain association key.
7. (Optional) In **KeyCKN**, type the connectivity association key name.
8. (Optional) In **KeyCAK**, type the Connectivity Association Key value.
9. (Optional) In **KeyExpiry**, type the expiry time for the keychain association key.
10. Select **Insert**

MACsec KA Key Field Definitions

Use the data in the following table to use the **MACsec KA Key** tab.

Name	Description
Id	Specifies the unique identification number for the MKA keychain configured on the switch.
KeyId	Specifies the unique identification number for the keychain association key configured on the switch.
KeyCKN	Specifies the connectivity association key name (CKN).
KeyCAK	Specifies the Connectivity Association Key (CAK).
KeyExpiry (UTC)	Specifies the expiry time for the keychain association key in Coordinated Universal Time (UTC) regardless of the switch time-zone.
KeyStatus	Specifies the status of the keys in the following states: <ul style="list-style-type: none"> • Valid—a key with a valid end time but no MACsec MKA enabled on the port • Expired—a key that is no longer valid • In-use—a key with a valid end time and assigned ports with MACsec MKA enabled

Enable MKA on a Port

Before You Begin

Ensure that you apply an MKA profile to the port and associate the port with a Connectivity Association (CA) before you enable MKA on the port.

About This Task

Use the following procedure to enable a MACsec Key Agreement (MKA) profile on a port.

Procedure

1. In the Device Physical View, select the port to which you want to apply an MKA profile.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **MACsec MKA** tab.
5. Select **MKA Enable** to enable MKA on the port.
6. Select **Apply**.

MACsec MKA Field Descriptions

Use the data in the following table to use the **MACsec MKA** tab.

Name	Description
MACsecMKAProfileName	Specifies the MKA profile name.
CAName	Specifies the name of the connectivity association attached to the port or interface.
OffsetValue	Specifies the number of bytes after the Ethernet header from which data encryption begins. The default is no offset.
CipherSuite	Displays the cipher suite used to encrypt traffic with MACsec.
Keychain Name	Specifies the name associated with the keychain.
MKA Enable	Enables the MKA profile on the port. The default is disabled.

Configure MKA Actor Priority

About This Task

Use the following procedure to configure an actor priority value for a port. The MKA participant with the highest actor priority is designated as the key server.



Note

Ensure that peers in an MKA connection are not configured with the same actor priority value.

Procedure

1. In the navigation pane, expand **Security > Data Path**.
2. Select **MACSec**.
3. Select the **MKA Participants** tab.
4. Select a port.
5. For **ActorsPriority**, double-click the field and type a value for **ActorsPriority**.
6. Select **Apply**.

MKA Participants Field Descriptions

Use the data in the following table to use the **MKA Participants** tab.

Name	Description
PortNumber	Specifies the port number of the MKA session participant.
CAName	Specifies the Connectivity Association (CA) name associated with the MKA session participant.
MKAProfileName	Specifies the name of the MKA profile.
MKA Enable	Specifies whether MKA is enabled for the port.
ActorsPriority	Specifies a hexadecimal value for actor priority. The default is 10.

MACsec Performance

The topics in this section provide information on MACsec statistics.

MACsec Statistics

MAC Security (MACsec) is an IEEE 802® standard that allows authorized systems in a network to transmit data confidentially and to take measures against data transmitted or modified by unauthorized devices.

The switch supports the following statistics that provide a measure of MACsec performance.

Table 202: General MACsec statistics

Statistics	Description
TxUntaggedPkts	Specifies the number of transmitted packets without the MAC security tag (SecTAG), with MACsec disabled on the interface.
TxTooLongPkts	Specifies the number of transmitted packets discarded because the packet length is greater than the Maximum Transmission Unit (MTU) of the Common Port interface.
RxUntaggedPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec not operating in strict mode.

Table 202: General MACsec statistics (continued)

Statistics	Description
RxNoTagPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec operating in strict mode.
RxBadTagPkts	Specifies the number of received packets discarded with an invalid SecTAG or with a zero value Packet Number (PN)/invalid Integrity Check Value (ICV).
RxUnknownSCIPkts	Specifies the number of packets received with an unknown Secure Channel Identifier (SCI) and with MACsec not operating in strict mode.
RxNoSCIPkts	Specifies the number of packets received with an unknown Secure Channel Identifier (SCI) and with MACsec operating in strict mode.
RxOverrunPkts	Specifies the number of packets discarded because the number of received packets exceeded the cryptographic performance capabilities.

Table 203: Secure-channel inbound MACsec statistics

Statistics	Description
UnusedSAPkts	Specifies the summation of received unencrypted packets on all SAs of this secure channel, with MACsec not in strict mode.
NoUsingSAPkts	Specifies the summation of received packets that were discarded along with either encrypted packets or packets that were received with MACsec operating in strict mode.
LatePkts	Specifies the number of packets received that have been discarded for this Secure Channel (SC) with Replay Protect enabled. Note: Replay Protect is supported only by MACsec configurations using MACsec Key Agreement (MKA) protocol.
NotValidPkts	Specifies the summation of packets that were discarded in all SAs of the SC because they were not valid with one of the following conditions: <ul style="list-style-type: none"> • MACsec was operating in strict mode • The packets received were encrypted but contained erroneous fields.
InvalidPkts	Specifies the summation of all packets received that were not valid for this SC, with MACsec operating in check mode.
DelayedPkts	Specifies the summation of packets for this SC, with the Packet Number (PN) of the packets lower than the lower bound replay protection PN. Note: Replay Protect is supported only by MACsec configurations using MKA protocol.

Table 203: Secure-channel inbound MACsec statistics (continued)

Statistics	Description
UncheckedPkts	The total number of packets for this SC that: <ul style="list-style-type: none"> • were encrypted and failed the integrity check • were not encrypted and failed the integrity check • were received when MACsec validation was not enabled
OKPkts	Specifies the total number of Integrity Check Validated (ICV) packets for all SAs of this Secure Channel. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted.
OctetsValidated	Specifies the number of octets of plain text recovered from received packets that were integrity protected but not encrypted.
OctetsDecrypted	Specifies the number of octets of plain text recovered from received packets that were integrity protected and encrypted.

Table 204: Secure-channel outbound MACsec statistics

Statistics	Description
ProtectedPkts	Specifies the number of integrity protected but not encrypted packets for this transmitting SC.
EncryptedPkts	Specifies the number of integrity protected and encrypted packets for this transmitting SC.
OctetsProtected	Specifies the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.
OctetsEncrypted	Specifies the number of plain text octets that are integrity protected and encrypted on the transmitting SC.

Table 205: MACsec Key Agreement statistics

Statistics	Description
MKPDUs Validated & Rx	Specifies the number of MACsec Key Agreement Protocol Data Units (MKPDU) validated and received.
Rx Distributed SAK	Specifies the number of Secure Association Keys (SAK) received.
MKPDUs Transmitted	Specifies the number of MKPDUs transmitted.
Tx Distributed SAK	Specifies the number of SAKs transmitted.

Interoperability Behaviors

If the static key MACsec configuration involves interoperability with VSP 4900 Series, VSP 7400 Series, 5520 Series (only front panel ports), and VIMs, you can see the following display inconsistencies for MACsec statistics:

- one side can show delayed packets on the ingress side

- one side can throw an Ingress Key expiry log message or it might never show an Ingress Key expiry log message

View MACsec Statistics using CLI

Use the following procedures to view statistics for MACsec using the Command Line Interface (CLI).

View MACsec Statistics

Perform this procedure to view the MACsec statistics.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the general MACsec statistics:

```
show macsec statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]]]
```

3. View the secure-channel inbound MACsec statistics:

```
show macsec statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]]] secure-channel inbound
```

4. View the secure-channel outbound MACsec statistics:

```
show macsec statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]]] secure-channel outbound
```

Example

Display general MACsec statistics, inbound MACsec statistics, and outbound MACsec statistics:

The switch supports MACsec on specific ports. For more information, see your hardware documentation.



Note

Display inconsistencies might occur for MACsec statistics when different ports interoperate with one another.

```
Switch:1>enable
Switch:1#show macsec statistics 1/40
```

MACSEC Port Statistics				
PortId	TxUntagged Packets	TxTooLong Packets	RxUntagged Packets	RxNoTag Packets
1/40	0	0	0	0
PortId	RxBadTag Packets	RxUnknown SCIPackets	RxNoSCI Packets	RxOverrun Packets

```

1/40      0      0      0      0

Switch:1#show macsec statistics 1/40 secure-channel inbound

=====
MACSEC Port Inbound Secure Channel Statistics
=====
PortId      UnusedSA      NoUsingSA      Late      NotValid      Invalid
Packets      Packets      Packets      Packets      Packets      Packets
-----
1/40      0      0      0      100037      0

PortId      Delayed      Unchecked      Ok      Octets      Octets
Packets      Packets      Packets      Pkts      Validated      Decrypted
-----
1/40      0      0      0      53528828      0

Switch:1#show macsec statistics 1/40 secure-channel outbound

=====
MACSEC Port Outbound Secure Channel Statistics
=====
PortId      Protected      Encrypted      Octets      Octets
Packets      Packets      Packets      Protected      Encrypted
-----
1/40      0      99946      0      53434154

```

Clear MACsec Statistics

About This Task

You have the option to clear MACsec statistics for all ports, or clear MACsec statistics for a specific port. Clearing MACsec statistics can be useful for debugging purposes.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Clear all MACsec statistics:
macsec clear-stats
3. (Optional) Clear MACsec statistics for a specific port:
macsec clear-stats port {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}

Example

Clear all MACsec statistics:

```

Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#macsec clear-stats

```

Clear MACsec statistics for a specific port:

```

Switch:1>enable
Switch:1#configure terminal

```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#macsec clear-stats port 1/3
```

Variable Definitions

Use the data in the following table to use the **clear macsec-stats** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View MKA MACsec Statistics using CLI

Use the following procedures to view statistics for MKA MACsec using the Command Line Interface (CLI).

Display MKA Statistics

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display MKA statistics for a specific port:

```
show macsec mka statistics {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}
```

Example

The following example displays MACsec MKA statistics for a port.

```
Switch:1>show macsec mka statistics 1/3
```

```
=====
                        MKPDU Statistics on interface 1/3
=====
      MKPDUs Validated & Rx      : 1630
      Rx Distributed SAK          : 2
      MKPDUs Transmitted         : 1694
      Tx Distributed SAK          : 0
-----
-----
```


Variable Definitions

Use the data in the following table to use the **show macsec mka statistics** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Clear MKA Statistics

About This Task

You have the option to clear MKA statistics for all ports, or clear MKA statistics for a specific port. Clearing MKA statistics can be useful for debugging purposes.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Clear all MKA statistics:

```
macsec mka clear-stats
```
3. (Optional) Clear MKA statistics for a specific port:

```
macsec mka clear-stats port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}
```

Example

Clear all MKA statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#macsec mka clear-stats
```

Clear MKA statistics for a specific port:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#macsec mka clear-stats port 1/3
```

Variable Definitions

Use the data in the following table to use the **macsec mka clear-stats** command.

Variable	Value
port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View MACsec Statistics using EDM

Use the following procedures to view statistics for static MACsec using EDM.



Note

For MACsec Key Agreement (MKA) MACsec, the procedures View Secure Channel Inbound Statistics and View Secure Channel Outbound Statistics display only the statistics for the current Secure Association (SA), instead of displaying the statistics for all SAs created under that Secure Channel.

View MACsec Interface Statistics

Use this procedure to view the MACsec interface statistics using EDM.

Procedure

1. In the Device Physical View tab, select one or more ports for which you need to view the MACsec interface statistics.

The switch supports MACsec on specific ports. For more information, see your hardware documentation.

2. In the navigation pane, expand **Edit > Port > General**.
3. Select the **MACsec Interface Stats** tab.



Note

Use the **Clear Stats** button to clear MACsec interface statistics. The **Clear Stats** button is available to clear single-port as well as multiple-port MACsec interface statistics.

MACsec Interface Stats Field Descriptions

The following table describes the fields in the **MACsec Interface Stats** tab.

Field	Description
TxUntaggedPkts	Specifies the number of transmitted packets without the MAC security tag (SecTAG), with MACsec disabled on the interface.
TxTooLongPkts	Specifies the number of transmitted packets discarded because the packet length is greater than the maximum transmission unit (MTU) of the common port interface.
RxUntaggedPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec <i>not</i> operating in strict mode.
RxNoTagPkts	Specifies the number of received packets without the MAC security tag (SecTAG), with MACsec operating in strict mode.
RxBadTagPkts	Specifies the number of received packets discarded with an invalid SecTAG, or with a zero value packet number (PN), or invalid Integrity Check Value (ICV).
RxUnknownSCIPkts	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec <i>not</i> operating in strict mode.
RxNoSCIPkts	Specifies the number of packets received with an unknown secure channel identifier (SCI), and with MACsec operating in strict mode.
RxOverrunPkts	Specifies the number of packets discarded because the number of received packets exceeded the cryptographic performance capabilities.

View Secure Channel Inbound Statistics

Use this procedure to view the secure channel (SC) inbound statistics using EDM.

Procedure

1. In the Device Physical View tab, select one or more ports for which to view the SC inbound statistics.
For information about which ports support MACsec, see the hardware documentation.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.

4. Select the **SC Inbound Stats** tab.



Note

Use the **Clear Stats** button to clear single-port secure channel inbound statistics. The **Clear Stats** button is not available to clear multiple-port secure channel inbound statistics.

SC Inbound Stats Field Descriptions

The following table describes the fields in the **SC Inbound Stats** tab.

Field	Description
UnusedSAPkts	Specifies the summary of received unencrypted packets on all SAs of this secure channel, with MACsec <i>not</i> in strict mode.
NoUsingSAPkts	Specifies the summary of received packets that were discarded along with either encrypted packets or packets that were received with MACsec operating in strict mode.
LatePkts	Specifies the number of packets received that have been discarded for this secure channel (SC) with Replay Protect enabled. Note: Replay Protect is supported only by MACsec configurations using MACsec Key Agreement (MKA) protocol.
NotValidPkts	Specifies the summary of packets that were discarded in all SAs of the SC because they were not valid with one of the following conditions: <ul style="list-style-type: none"> • MACsec was operating in strict mode. • The packets received were encrypted but contained erroneous fields.
InvalidPkts	Specifies the summary of all packets received that were not valid for this SC, with MACsec operating in <i>check</i> mode.
DelayedPkts	Specifies the summary of packets for this SC, with the packet number (PN) of the packets lower than the lower bound replay protection PN. Note: Replay Protect is supported only by MACsec configurations using MACsec Key Agreement (MKA) protocol.

Field	Description
UncheckedPkts	The total number of packets for this SC that: <ul style="list-style-type: none"> • Were encrypted and failed the integrity check. • Were <i>not</i> encrypted and failed the integrity check. • Were received when MACsec validation was not enabled.
AcceptedPkts	Specifies the total number of Integrity Check Validated (ICV) packets for all SAs of this Secure Channel. The number of octets of User Data recovered from received frames that were integrity protected but not encrypted.
OctetsValidated	Specifies the number of octets of plain text recovered from received packets that were integrity protected but not encrypted.
OctetsDecrypted	Specifies the number of octets of plain text recovered from received packets that were integrity protected and encrypted.

View Secure Channel Outbound Statistics

Use this procedure to view the secure channel (SC) outbound statistics using EDM.

Procedure

1. In the Device Physical View tab, select one or more ports for which you need to view the SC outbound statistics.
For information about which ports support MACsec, see the hardware documentation.
2. In the navigation pane, expand **Configuration > Edit > Port** folders.
3. Select **General**.
4. Select the **SC Outbound Stats** tab.



Note

Use the **Clear Stats** button to clear single-port secure channel outbound statistics. The **Clear Stats** button is not available to clear multiple-port secure channel outbound statistics.

SC Outbound Stats Field Descriptions

The following table describes the fields in the **SC Outbound Stats** tab.

Field	Description
ProtectedPkts	Specifies the number of integrity protected but not encrypted packets for this transmitting SC.
EncryptedPkts	Specifies the number of integrity protected and encrypted packets for this transmitting SC.
OctetsProtected	Specifies the number of plain text octets that are integrity protected but not encrypted on the transmitting SC.
OctetsEncrypted	Specifies the number of plain text octets that are integrity protected and encrypted on the transmitting SC.

View MKA MACsec Statistics using EDM

Use the following procedure to view statistics for MKA MACsec using EDM.

Display MKA Statistics for a Port

Procedure

1. In the Device Physical View, select the port or ports for which to display MKA statistics.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **MACsec MKA Stats** tab.
5. (Optional) Select **Clear Stats** to clear MKA MACsec statistics for the selected port or ports.

Clearing MKA MACsec statistics can be useful for debugging purposes.

MACsec MKA Stats Field Descriptions

Use the data in the following table to use the **MACsec MKA Stats** tab.

Name	Description
MKPDUValidatedPkts	Specifies the number of MKPDU packets received and validated.
RxDistributedSAKPkts	Specifies the number of SAK packets received.
MKPDUTransmittedPkts	Specifies the number of MKPDU packets transmitted.
TxDistributedSAKPkts	Specifies the number of SAK packets transmitted.



MultiLink Trunking and Split MultiLink Trunking

[Virtual Inter-Switch Trunk \(vIST\) on page 2657](#)

[Simplified Virtual-IST on page 2659](#)

[MultiLink Trunking on page 2660](#)

[Split MultiLink Trunking on page 2662](#)

[MLT and SMLT Configuration Requirements on page 2671](#)

[MLT and SMLT Link Aggregation Configuration Using the CLI on page 2675](#)

[MLT and SMLT Link Aggregation Configuration using EDM on page 2696](#)

[MLT Configuration Examples on page 2711](#)

[MLT Network Topology and Configuration Reference on page 2713](#)

This section provides the concepts and procedures you need to configure MultiLink Trunking (MLT) and Split MultiLink Trunking (SMLT).

Virtual Inter-Switch Trunk (vIST)

Table 206: Virtual Inter-Switch Trunk product support

Feature	Product	Release introduced
Switch cluster (multi-chassis LAG) -Virtual Inter-Switch Trunk (vIST)	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Split MultiLink Trunking provides subsecond failover when a switch fails. Virtual Inter-Switch Trunk (vIST) improves upon that Layer 2 and Layer 3 resiliency by using a virtualized IST channel through the SPBM cloud. The vIST channel carries the vIST

control traffic and data traffic during an SMLT failover. This feature dramatically improves resiliency over other methods.

Because vIST uses a virtual channel and because IS-IS runs over it, vIST eliminates the potential single point of failure with a dedicated MLT. The vIST channel is always up as long as there is SPBM connectivity between the vIST peers.

vIST interoperates between any two devices that support vIST, and the devices do not have to be of the same type.

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.

**Important**

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

**Note**

- Users may observe a momentary increase in activity when a MAC delete message is received from a peer. This is due to vIST engaging in MAC learning activities. This is a normal operational procedure.
- For proper traffic flow, if a Layer 2 VSN is created on one vIST peer, it must also be created on the other vIST peer. For more information on Layer 2 VSN, see [Layer 2 VSN Configuration](#) on page 1514.

vIST Configuration Best Practice

If you need to update the vIST VLAN IP address on vIST peers by deleting and recreating the vIST vlan IP address (for example, as part of maintenance), ensure that you update one vIST BEB at a time.

**Caution**

Always perform vIST configuration updates under no traffic. Otherwise, it results in traffic loss.

Before you begin updating a device, as a first step, isolate the device by shutting down all the links and failing over the traffic to its vIST peer. Then, delete and recreate the vIST VLAN IP on the device and save your configuration. When bringing the device back into operation, first unshut those network-to-network interface (NNI) ports that bring up vIST, followed by the SMLT configured ports, and then all the remaining ports, to prevent network loops or duplicate traffic.

For information on vIST configuration, see [Create a Virtual IST](#) on page 2683 or [Create a Virtual IST using EDM](#) on page 2705.

vIST Operations

When you enable IST and boot the chassis, the SMLT enabled trunk ports (SMLT ports) are automatically locked. A timeout mechanism automatically unlocks the SMLT ports when the IST control channel fails to establish within a reasonable amount of time. The timeout mechanism prevents the SMLT ports from being locked forever. Initially 240 seconds are allowed for the switch to determine the IST VLAN status.

The IST VLAN is considered up if at least one port is forwarding traffic and an ARP entry is populated for the IP address of the IST peers. Once the IST VLAN is up, the timeout value is reset to 60 seconds. The IST control channel must be up within the timeout period. If the timeout period is exceeded, then the SMLT ports are automatically unlocked and a message is logged stating that the SMLT ports are unlocked due to a timeout.



Note

If the IST filter is enabled before the timeout, then the IST filter is unaffected and remains enabled.

If a virtual BMAC mismatch occurs between two vIST peers, incorrect information can show in CLI regarding the ports that the client information is learned on, such as MAC addresses or IPv6 neighbor addresses. This can occur before the mismatch is detected, because the two peers receive traffic from each other and learn the information on the local ports. Each vIST peer synchronizes the information as learned locally, even when learned from vIST peer. You can resolve the mismatch by changing the virtual BMAC on one of the vIST peers, but the traffic might remain affected. After confirming the mismatch is resolved from a log report, you can resolve the traffic issue by disabling and enabling the router ISIS on both vIST peers.

Simplified Virtual-IST

Simplified Virtual-IST (vIST) is for conventional switch clustering deployments that use SMLT and not SPB. The Simplified vIST feature provides a single CLI command to enable the virtual IST for SMLT deployments.

- Simplified vIST is available for conventional multicast deployments with PIM and IGMP only when the boot flag (**spbm-config-mode**) is disabled.



Important

PIM is supported with Simplified vIST only, not with SPB vIST.

- When the **spbm-config-mode** boot flag is enabled (default setting), Simplified vIST is not available. This means that you continue to configure SPB/IS-IS for vIST as described in [Create a Virtual IST](#) on page 2683 and [Create a Virtual IST using EDM](#) on page 2705.
- Simplified VIST requires that the two vIST devices be directly connected.
- For legacy IGMP snooping and IGMP snooping over Simplified vIST, the IGMP sender information is NOT synchronized across the vIST. This means that **show ip igmp sender** displays the sender record only on the switch that the multicast stream actually ingresses.

- In Simplified vIST mode, you should not use LACP on vIST MLT.
- In Simplified vIST mode, you cannot enable PIM Infinite Threshold Policy.

**Note**

For Simplified vIST deployment, if a VLAN is part of an SMLT it must be configured on both of the IST peers.

**Important**

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of a header size increase.

For configuration information, see [Configure Simplified vIST in SMLT Topologies](#) on page 2688 or [Configure Simplified vIST in SMLT Topologies](#) on page 2707.

For information about how to configure Simplified vIST with multicast, see [Enabling Multicast on the Switch](#) on page 1696.

MultiLink Trunking

Table 207: MultiLink Trunking product support

Feature	Product	Release introduced
MultiLink Trunking (MLT) / Link Aggregation Group (LAG)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

MultiLink Trunking (MLT) is a point-to-point connection that aggregates multiple ports to logically act like a single port with aggregated bandwidth. Grouping multiple ports into a logical link provides a higher aggregate on a switch-to-switch or switch-to-server application.

To include ports as trunk group members of an MLT, you must statically configure the ports.

MLT Traffic Distribution Algorithm

You can use a multilink trunk to aggregate bandwidth between two switches. The MLT algorithm ensures that each packet in a flow does not arrive out of sequence, and that a flow always traverses the same link path.

The hashing algorithm uses the following packet fields and the incoming interface (source) port number to calculate the index to outgoing (destination) port number in an MLT:

Traffic type	Hashing algorithm
IPv4 traffic	Hash Key = [Destination IP Address (32 bits), Source IP Address (32 bits), Source TCP/UDP Port, Destination TCP/UDP port]
IPv4 traffic without TCP/UDP header	Hash Key = [Source IP Address (32 bits), Destination IP address (32 bits)]
IPv6 traffic	Hash Key = [Destination IPv6 Address (128 bits), Source IPv6 address (128 bits), Source TCP/UDP Port, Destination TCP/UDP port]
IPv6 traffic without TCP/UDP header	Hash Key = [Source IP Address (128 bits), Destination IP address (128 bits)]
Mac-In-Mac transit traffic	<ul style="list-style-type: none"> Inner Layer 2 non-IP packet: Hash Key = [Source Port (8bits), Customer VLAN(12bits), Customer Destination Mac Address (48 bits), Customer Source Mac Address (48 bits)] Inner IP packet: Hash Key = [Source Port (8bits), VLAN(12bits), Destination IP (32bits), Source IP (32bits)]
Layer 2 Non-IP traffic	Hash Key = [Destination MAC Address (48 bits), Source MAC Address(48 bits)]

MultiLink Trunking and Autonegotiation Interaction

To use MLT with the switch, you can have ports running at different speeds. After you use MLT with LACP, LACP dynamically checks for proper speed on all port members. You do not need to have similar physical connection types. After you use autonegotiation with MLT and not LACP, you need to ensure that all ports run at the same speed.

MLT Configuration Rules

Multilink trunks adhere to the following rules. Unless otherwise stated, these rules also apply to MLT with LACP.

- You cannot configure an MLT name that uses all numbers, for example, 222.
- Multilink trunk ports support mixed speed links, for example, one link can be 10Gb and another 1Gb. However, no weighting of traffic distribution occurs so if you mix links of different operational speeds, you can overload the lower speed link or underutilize a higher speed link.

This rule applies to multilink trunks only. MLT with LACP does not support different link speeds.

- All multilink trunk ports must be in the same Spanning Tree Group (STG) unless the port is tagged. Use tagging so ports can belong to multiple STGs, as well as multiple VLANs.
- After the port is made a member of MLT, it inherits the properties of the MLT and hence the STG properties are inherited from the VLAN associated with that MLT. After you remove the port from MLT or after you delete the MLT, the ports are removed from the MLT STG and added into the default STG.
- MLT is compatible with Multiple Spanning Tree Protocol (MSTP) (IEEE 802.1s) and Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w).
- Tagging (IEEE 802.1Q) is supported on a multilink trunk.

Multilink trunks have the following general features and requirements:

- Supports MLT groups with as many as 8 ports belonging to a single multilink trunk. For more information about the number of MLT groups supported for each hardware platform, see *Fabric Engine Release Notes*.
- Apply filters individually to each port in a multilink trunk.

With MSTP or RSTP enabled, ports in the same multilink trunk operate as follows:

- The designated port sends the Bridge Protocol Data Unit (BPDU).
- The multilink trunk port ID is the ID of the lowest numbered port.
- If identical BPDUs are received on all ports, the multilink trunk mode is forwarding.
- If ports do not receive BPDUs on a port or BPDU and port tagging do not match, the individual port is taken offline.
- Path cost is inversely proportional to the active multilink trunk bandwidth.

MLT with LACP LAG Rules

The Link Aggregation Group (LAG) adheres to the following rules:

- All LAG ports operate in full-duplex mode.
- All LAG ports operate at the same data rate.
- All LAG ports must belong to the same set of VLANs.
- Link aggregation is compatible with MSTP, and RSTP.
- Assign all LAG ports to the same MSTP or RSTP groups.
- You can configure a LAG with up to 24 ports, but only a maximum of 8 can be active at a time.
- After you configure a multilink trunk with LACP, you cannot add or delete ports or VLANs manually without first disabling LACP.

Split MultiLink Trunking

Split MultiLink Trunking (SMLT) is an option that improves Layer 2 and Layer 3 resiliency. The following sections discuss SMLT in more detail.

SMLT Overview

Split MultiLink Trunking is an option that improves Layer 2 (bridged) resiliency by providing for the addition of switch failure redundancy with subsecond failover, on top of all standard MLT link failure protection and flexible bandwidth scaling functionality. Use Split MultiLink Trunking to connect a device that supports some form of link aggregation, be it a switch or a server, to two distinct separate SMLT endpoints or switches. These SMLT devices form a virtualized Switch Cluster through the SPBM cloud and are referred to as a Virtual Inter-Switch Trunk (vIST) Core Switch pair.

Switch Clusters are always formed as a pair, but you can combine pairs of clusters in either a square or full-mesh fashion to increase the size and port density of the Switch Cluster. If you configure SMLT in a Layer 3 or routed topology, the configuration is referenced as Routed SMLT (RSMLT).

For information about Routed SMLT, see [RSMLT](#) on page 3285.

You can form SMLT connections through single links from the Switch Cluster to the edge connection, standard MLTs, or MLTs with LACP. Optionally, SMLT links can have VLACP enabled as well. You can mix these various link connections. Within the same Switch Cluster, you can configure both SMLT and RSMLT to allow a mixture of both Layer 2 and Layer 3 VLANs.

Split MultiLink Trunking networks do not need to use RSTP or MSTP to enable loop-free triangle topologies because SMLT inherently avoids loops due to its superior enhanced link aggregation protocol. The loop-free link is accomplished by having two aggregation switches, displays it as a single device to edge switches, which dual-home to the aggregation switches. The aggregation switches interconnect using a Virtual Inter-Switch trunk (vIST), exchanging addressing and state information (permitting rapid fault detection and forwarding path modification). Split MultiLink Trunking is designed for Layer 2 network connectivity, but you can configure in Layer 3 networks by working with VRRP or RSMLT Layer 2 edge.

SMLT Advantages

SMLT eliminates all single points of failure and creates multiple paths from all user access switches to the network core. In case of failure, SMLT recovers as quickly as possible using all capacity. SMLT provides a transparent and interoperable solution that requires no modification on the part of the majority of existing user access devices.

SMLT improves the reliability of Layer 2 networks that operate between user access switches and the network center aggregation switch by providing:

- load sharing among all links
- fast failover in case of link failure
- elimination of single points of failure
- fast recovery in case of node failure
- transparent and interoperable solutions
- removal of MSTP and RSTP convergence issues

SMLT, MSTP, and RSTP

Networks designed to have user access switches dual-home to two aggregation switches, and have VLANs spanning two or more user access switches, experience the following design constraints:

- no load sharing exists over redundant links
- network convergence is slow in case of failure

With the introduction of SMLT, all dual-home Layer 2 frame-switched network devices with dual homes are no longer dependent on the MSTP or RSTP for loop detection. A properly designed SMLT network inherently does not have logical loops.

SMLT solves the spanning tree problem by combining two aggregation switches into one logical MLT entity, thus making it transparent to all types of edge switches. In the process, it provides quick convergence, while load sharing across all available trunks.

If you use STP mode on a switch that is in an SMLT configuration, you can experience traffic loss for 30 seconds if you change the port membership of the MLT, even if the port is in a down state. The traffic loss is because the convergence time for STP is 30 seconds. Use MSTP or RSTP on all switches in SMLT configurations.

SMLT Topologies

The following are the three generic topologies in which you can deploy SMLT, depending on the resiliency and redundancy required:

- a triangle topology
- a square topology
- a full-mesh topology

SMLT and Virtual Inter-Switch Trunk (vIST)

The following illustration shows an SMLT configuration with a pair of switches as aggregation systems (E and F) and four separate switches as user access switches (A, B, C, and D).



Note

In the following figure, the pair of aggregation switches (E and F) are connected by a Virtual Inter-Switch Trunk (vIST). This means that although the system displays these two switches to be directly connected, they could be anywhere within the SPBM cloud.




Figure 181: SMLT aggregation switches and operations

You must connect SMLT aggregation switches through vIST. For example, user access switches B and C connect to the aggregation systems through multilink trunks split between the two aggregation systems. As shown in the previous figure, the implementation of SMLT only requires two SMLT-capable aggregation switches.

Aggregation switches use vIST to do the following:

- Confirm that they are alive and exchange MAC address forwarding tables.
- Send traffic between single switches attached to the aggregation switches.
- Serve as a backup if one SMLT link fails.

Because SMLT requires vIST, use multiple links on vIST to ensure reliability and high availability. Use Gigabit Ethernet links for vIST connectivity to provide enough bandwidth for potential cross traffic.

For more information about vIST, see [Virtual Inter-Switch Trunk \(vIST\)](#) on page 2657.

A vIST multilink trunk must contain at least two physical ports.

Other SMLT Aggregation Switch Connections

The previous figure includes end stations that connect to each of the switches. In this example, a, b1, b2, c1, c2, and d are clients and printers, while e and f are servers and routers.

User access switches B and C can use a method to determine which link of the multilink trunk connections to use to forward a packet, as long as the same link is used for a Source Address and Destination Address (SA/DA) pair. The packet is routed correctly regardless of whether B or C knows the DA. SMLT aggregation switches always send traffic directly to a user access switch, and only use vIST for traffic that they cannot forward in another, more direct way.

SMLT Environment Traffic Flow Rules

Traffic flow in an SMLT environment adheres to the following rules:

- If a packet is received from a vIST trunk port, it is not forwarded to an active SMLT group in order to prevent network loops.
- After a packet is received, the system performs a look-up on the forwarding database. If an entry exists, and if the entry was learned locally from the SMLT or through the vIST as a remote SMLT, it is forwarded to the local port (the packet must not be sent to the vIST for forwarding unless there is no local connection). Unknown and Broadcast packets flood out all ports that are members of this VLAN.
- For load sharing purposes in an SMLT configuration, the switch obeys the MLT traffic distribution algorithm.

For more information, see [SMLT Traffic Flow Examples](#) on page 2669 .

SMLT and LACP Support

The switch fully supports IEEE 802.3ad LACP on MLTs and on a pair of SMLT systems.

With LACP the switch provides a standardized external link aggregation interface to third-party vendor IEEE 802.3ad implementations. This protocol extension provides dynamic link aggregation mechanisms. Only dual-home devices benefit from this enhancement.

Advantages of this protocol extension include the following:

- MLT peers and SMLT client devices can be both network switches and a type of server or workstation that supports link bundling through IEEE 802.3ad.
- Single link and multilink trunk solutions support dual-home connectivity for attached devices, so that you can build dual-home server farm solutions.

Supported SMLT/LACP Scenarios

SMLT/IEEE link aggregation interaction supports all known SMLT scenarios in which an IEEE 802.3ad SMLT pair connects to SMLT clients, or in which two IEEE 802.3ad SMLT pairs connect to each other in a square or full-mesh topology.

Unsupported SMLT/LACP Scenarios

Some of the unsupported SMLT/LACP scenarios include the following factors, which lead to failure:

- Incorrect port connections.
- Mismatched MLT IDs assigned to SMLT client. SMLT switches can detect if MLT IDs are not consistent. The SMLT aggregation switch, which has the lower IP address, does not allow the SMLT port to become a member of the aggregation thereby avoiding misconfigurations.
- The SMLT client switch does not have automatic aggregation enabled (LACP disabled). SMLT aggregation switches can detect that aggregation is not enabled on the SMLT client, thus no automatic link aggregation is established until the configuration is resolved.

SMLT and the Virtual Router Redundancy Protocol

Use Virtual Router Redundancy Protocol (VRRP) to have one active primary router for each IP subnet, with all other network VRRP interfaces operating in backup mode.

The VRRP has only one active routing interface enabled. Users that access switches aggregated into two SMLT switches send their shared traffic load (based on source and destination MAC or IP addresses) on all uplinks towards the SMLT aggregation switches.



Note

A VRRP virtual IP address must not be the same as the VLAN IP address of the device.

The VRRP is less efficient if you use it with SMLT. All other interfaces are in backup (standby) mode. In this case, all traffic is forwarded over the vIST link towards the primary VRRP switch. All traffic that arrives at the VRRP backup interface is forwarded, so there is not enough bandwidth on the vIST link to carry all the aggregated riser traffic. However, an enhancement to VRRP overcomes this issue by ensuring that the vIST trunk is not used in such a case for primary data forwarding.

SMLT and VRRP BackupMaster

The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. The system directly routes all traffic to the destined subnetwork and not through Layer 2 switches to the VRRP master. This avoids a potential limitation in the available vIST bandwidth.

To avoid potential frame duplication problems, you can only use the VRRP BackupMaster feature for SMLT on interfaces that you configure for SMLT. You cannot use VRRP BackupMaster with hubs to avoid frame duplication or on brouter or VLAN interfaces.

If you use an SMLT with routing on SMLT aggregation switches, use VRRP for default gateway redundancy. In a VRRP environment, one switch is active and the other is a backup. In an SMLT environment, you can enable the VRRP BackupMaster and use an

active-active concept. The VRRP BackupMaster router routes traffic that is received on the SMLT VLAN and avoids traffic flow across the vIST. This provides true load-sharing abilities.

Follow these guidelines if you use VRRP BackupMaster with SMLT:

- The VRRP virtual IP address and the VLAN IP address cannot be the same.
- Configure the hold-down timer for VRRP to a value that is approximately one hundred and fifty percent of the Interior Gateway Protocol (IGP) convergence time to allow the IGP enough time to reconverge following a failure. For example, if OSPF takes 40 seconds to reconverge, configure the hold-down timer to 60 seconds.
- Enable hold-down times on both VRRP sides (Master and BackupMaster).

SMLT and SLPP

You can use Simple Loop Prevention Protocol (SLPP) to prevent loops in an SMLT network. SLPP focuses on SMLT networks but works with other configurations. Always use SLPP in an SMLT environment.



Note

If SLPP is used in a vIST environment, it must be enabled on both the vIST peers. Because, when an SLPP packet of a vIST peer is looped through UNI ports to the other device, that device will shut down its UNI port due to receiving SLPP packets from its peer. A device's own SLPP packets will go over a vIST connection but will not be forwarded by its vIST peer back onto its UNI ports.

For square and full-mesh configurations that use a routed core, create a separate core VLAN. Enable SLPP on the core VLAN and the square or full mesh links between switch clusters. This configuration detects loops created in the core and loops at the edge do not affect core ports. If you use RSMLT between the switch clusters, enable SLPP on the RSMLT VLAN. Because you enable SLPP only on one or two VLANs in the core, changing the RX threshold values will not be necessary.

The SLPP design is to shut down the port where the SLPP packets originate or to shut down the vIST peer switch port, after the counter reaches the threshold. A loop can still occur after ports are shut down. SLPP can shut down all SMLT ports on a triangle SMLT topology, which results in isolating the edge switch.

SLPP Guard

Table 208: SLPP Guard product support

Feature	Product	Release introduced
SLPP Guard	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Use SLPP Guard with SMLT to provide additional loop protection to protect wiring closets from erroneous connections.

SLPP Guard requires Simple Loop Prevention Protocol (SLPP) to be configured in the core of the network. SLPP detects loops in the SMLT network. Because SMLT networks disable Spanning Tree (STP), Rapid Spanning Tree (RSTP), or Multiple Spanning Tree Protocol (MSTP) for participating ports, SLPP Guard provides additional network loop protection, extending the loop detection to individual edge access ports.

SLPP Guard can be configured on MLT or LAG ports. If the edge switch with SLPP Guard enabled receives an SLPP-PDU packet on a port, SLPP Guard operationally disables the port for the configured timeout interval and appropriate log messages and SNMP traps are generated. If the disabled port does not receive any SLPP-PDU packets after the configured timeout interval expires, the port automatically reenables and generates a local log message, a syslog message, and SNMP traps, if configured.

SMLT Traffic Flow Examples

The first four traffic flow examples are based on [Figure 181](#) on page 2665.

Example 1: Traffic Flow From a to b1 or b2

Assuming a and b1/b2 are communicating through Layer 2, traffic flows from A to switch E and is forwarded over its direct link to B. Traffic coming from b1 or b2 to a is sent by B on one of its multilink trunk ports.

B can send traffic from b1 to a on the link to switch E, and traffic from b2 to a on the link to F. In the case of traffic from b1, switch E forwards the traffic directly to switch A, while traffic from b2, which arrived at F, is forwarded across the vIST to E and then to A.

Example 2: Traffic Flow From b1/b2 to c1/c2

Traffic from b1/b2 to c1/c2 is always sent by switch B through its multilink trunk to the core. No matter at which switch E or F arrives at, it is sent directly to C through the local link.

Example 3: Traffic Flow From a to d

Traffic from a to d (and d to a) is forwarded across vIST because it is the shortest path. The link is treated as a standard link; SMLT and vIST parameters are not considered.

Example 4: Traffic Flow From f to c1/c2

Traffic from f to c1/c2 is sent out directly from F. Return traffic from c1/c2 passes through one active VRRP Master for each IP subnet. The traffic is passed across vIST if switch C sends it to E.

SMLT and vIST Traffic Flow Example

In an SMLT environment, the two aggregation switches share the same forwarding database by exchanging forwarding entries using the vIST. The entry for 00:E0:7B:B3:04:00 is shown on switch C as an entry learned on MLT-1, but because SMLT Remote is true, this entry was actually learned from switch B. On B, that same entry is shown as directly learned through MLT-1 because SMLT Remote is false.

The following illustration shows the network topology.




Figure 182: Network topology for traffic flow example

After a packet arrives at switch C destined for 00:E0:7B:B3:04:00, if the SMLT Remote status is true, the switch tries to send the packet to MLT-1 rather than through vIST. Traffic rarely traverses vIST unless there is a failure. If this same packet arrives at B, it is then forwarded to MLT-1 on the local ports.

MLT and SMLT Configuration Requirements

Use the information in this section to understand the considerations and guidelines while configuring link aggregation into your network.



Note

Static MAC is not supported against SMLT.

MLT with LACP

After you configure MLT with LACP, you must enable the aggregation parameter. After you enable the aggregation parameter, the LACP aggregator maps one-to-one to the specified multilink trunk.

The following lists the steps that are involved to configure MLT with LACP:

1. Assign a numeric key to the ports you want to include.
2. Configure the LAG for aggregation.
3. Enable LACP on the port.
4. Create an MLT and assign to it the same key as in step 1.

The multilink trunk with LACP only aggregates ports whose key matches its own.

The newly created MLT with LACP adopts the VLAN membership of its member ports after the first port is attached to the aggregator associated with this LAG. After a port detaches from an aggregator, the associated LAG port deletes the member from its list.

After a multilink trunk is configured with LACP, you cannot add or delete ports manually without first disabling LACP. You can add or remove VLANs to an MLT without manually disabling LACP.

The following list identifies expected configuration behavior for all platforms when you add VLANs to an LACP-enabled MLT:

- If only one port is part of the MLT/LAG, use the **vlan members add <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}** command to add newer VLANs to the LAG.
- If two or more ports with several VLANs are active members of the MLT/LAG, you must use the **mlt <1-512> vlan <1-4059>** command to add new VLANs.
- If at least one port is active on the MLT/LAG (LACP), and you need to add a second port that is physically up and has the same key, you must use the **mlt <1-512> vlan <1-4059>** command. The port values must match otherwise the new port enters the churn condition.
- If the switch displays a persistent error that you cannot change the port VLAN membership, you must correct the configuration during a maintenance window:
 1. Shutdown the ports.
 2. Disable LACP on the ports.
 3. Add all VLANs to the port members being aggregated.

4. Enable LACP on the ports.
5. Enable the ports.

To enable tagging on ports belonging to a LAG, disable LACP on the port, and then enable tagging and LACP on the port.

If you enable Open Shortest Path First (OSPF) routing on a port, do not set the LACP periodic transmission timer to less than 1 second.

MLT with LACP and SMLT

You can configure Split MultiLink Trunking (SMLT) with MLT, or with MLT with LACP. Follow these guidelines while you configure SMLT with LACP:

- If you configure LACP for SMLT, you must configure the same LACP smlt-sys-id on both sides. After you configure the LACP system ID for SMLT, configure the same LACP smlt-sys-id on both aggregation switches to avoid loss of data. Configure the LACP smlt-sys-id to be the base MAC address of one of the aggregate switches, and include the MLT-ID. Configure the same system ID on both of the SMLT core aggregation switches.
- If you use LACP in an SMLT square configuration, the LACP ports must have the same keys for that SMLT LAG; otherwise, the aggregation can fail if a switch fails.
- If an SMLT aggregation switch has LACP enabled on some of its multilink trunks, do not change the LACP system priority.
- After you configure SMLT links, set the multicast packets per-second value to 6000 pps.
- To avoid traffic loss when an SMLT goes down during a CP switchover, distribute multiple SMLT links across different slots.
- To avoid unnecessary processing, do not enable LACP on vISTs. Use VLACP if an optical network between the SMLT core switches requires a failure detection mechanism.

Using the LACP smlt-sys-id enables you to use a third-party switch as a wiring closet switch in an SMLT configuration. This enhancement provides an option for the administrator to configure the SMLT Core Aggregation Switches to always use the system ID. In this way, the SMLT Core Aggregation Switch always uses the same LACP key regardless of the state of the SMLT Core Aggregation Switch neighbor (or the vIST link). Therefore no change in LAGs must occur on the attached device regardless of whether the device is a server or a third-party switch. This situation does not affect edge switches used in SMLT configurations. The actor system priority of LACP_DEFAULT_SYS_PRIO, the actor system ID the user configures, and an actor key equal to the MLT-ID are sent to the wiring closet switch. Configure the system ID to be the base MAC address of one of the aggregate switches along with its MLT-ID. The administrator must ensure that the same value for the system ID is configured on both of the SMLT Core Aggregation Switches.

You can configure the LACP smlt-sys-id used by SMLT core aggregation switches. After you set the LACP system ID for SMLT, configure the same LACP smlt-sys-id on both aggregation switches to avoid the loss of data.

The LACP System ID is the base MAC address of the switch, which is carried in Link Aggregation Control Protocol Data Units (LACPDU). If two links interconnect two switches that run LACP, each switch knows that both links connect to the same remote device because the LACPDUs originate from the same System ID. If you enable the links for aggregation using the same key, LACP can dynamically aggregate them into an MLT LAG.

If SMLT is used between the two switches, they act as one logical switch. Both aggregation switches must use the same LACP System ID over the SMLT links so that the edge switch sees one logical LACP peer, and can aggregate uplinks towards the SMLT aggregation switches. This process automatically occurs over the vIST connection, where the base MAC address of one of the SMLT aggregation switches is chosen and used by both SMLT aggregation switches.

However, if the switch that owns that Base MAC address restarts, the vIST goes down, and the other switch reverts to using its own Base MAC address as the LACP System ID. This action causes all edge switches that run LACP to think that their links are connected to a different switch. The edge switches stop forwarding traffic on their remaining uplinks until the aggregation can reform (which can take several seconds). Additionally, after the restarted switch comes back on line, the same actions occur, thus disrupting traffic twice.

The solution to this problem is to statically configure the same LACP smlt-sys-id MAC address on both aggregation switches.

**Note**

The SMLT ID is always the same as the MLT ID. For instance, both sides can have an MLT 10, but once SMLT is enabled on both sides it will function as an SMLT. Until SMLT is enabled on both peers however, it will function as a normal MLT.

MLT with LACP and Spanning Tree

Only the physical link state or its LACP peer status affects LACP module operation. After a link is enabled or disabled, an LACP module is notified. The MSTP or RSTP forwarding state does not affect LACP module operation. LACPDUs can be sent if the port is in an MSTP or RSTP blocking state.

Unlike legacy MultiLink trunks, configuration changes (such as speed and duplex mode) to a LAG member port are not applied to all member ports in the multilink trunks. The changed port is removed from the LAG and the corresponding aggregator, and the user is alerted that the configuration is created.

**Important**

Link Aggregation Control Protocol, as defined by IEEE, is a protocol that exists between two bridge endpoints; therefore, the LACP PDUs are terminated at the next Server Provider (SP) interface.

SMLT and VLACP

Use Virtual Link Aggregation Control Protocol (VLACP) for all SMLT access links configured as MultiLink Trunks to ensure both end devices can communicate. The switch does not support LACP and VLACP on the same links simultaneously.

VLACP for SMLT also protects against CPU failures by causing traffic to switch or reroute to the SMLT peer if the CPU fails or stops responding.

The following table provides the values for VLACP in an SMLT environment:

Table 209: VLACP values

Parameter	Value
SMLT access	
Timeout	Short
Timer	500ms
Timeout scale	5
VLACP MAC	01:80:C2:00:00:0F
SMLT core	
Timeout	Short
Timer	500ms
Timeout scale	5
VLACP MAC	01:80:C2:00:00:0F
vIST	
Timeout	Long
Timer	10000
Timeout scale	3
VLACP MAC	01:80:C2:00:00:0F

SMLT with NNI Ports

If you want to modify SMLTs that contain network-to-network interface (NNI) ports, do the modification during maintenance windows. Otherwise, if you create or delete SMLTs that contain NNI ports running MSTP, IS-IS adjacencies that connect to those ports can bounce even if the SMLT is not used.

MLT and Private VLANs

The switch supports private VLANs and E-Tree configuration. MLT and private VLANs operate as follows:

- When using static MLT, configure the private-vlan as part of the overall MLT configuration, not at the port member level. All ports in the MLT use the private VLAN type for that MLT.
- When using LACP, configure the private-vlan at the interface level.
- When the private VLAN port type is trunk, the MLT automatically becomes tagged.
- If there are other non-private VLANs using the MLT configured as isolated, the following message is displayed: `All non private VLANs using this interface will be removed once this mlt becomes isolated. Do you wish to continue (y/n) ?`
- If there are other non-private VLANs using the MLT configured as promiscuous, the following message is displayed: `All non private VLANs using this interface will be removed once this mlt becomes promiscuous. Do you wish to continue (y/n) ?`

MLT and SMLT Link Aggregation Configuration Using the CLI

This section describes how to configure and manage link aggregation using the CLI, including MultiLink Trunking (MLT), to increase the link speed and redundancy for higher availability.

Configure link aggregation to provide link level redundancy and increase load sharing. MultiLink Trunking (MLT) is a link aggregation technology that you can use to group several physical Ethernet links into one logical Ethernet link to provide fault-tolerance and high-speed links between routers, switches, and servers. Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency.

Configure an MLT

Before You Begin

See [MLT and SMLT Configuration Requirements](#) on page 2671 to understand the commands you can use to add VLANs to an LACP-enabled MLT. Ensure you use the correct command to avoid the LACP churn condition.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create an MLT:
`mlt <1-512>`
3. (Optional) Set the private VLAN type for the MLT:
`mlt <1-512> private-vlan <isolated|promiscuous|trunk>`

4. Add a VLAN to the MLT:

```
mlt <1-512> vlan <1-4059>
```

5. (Optional) Change the name of the MLT:

```
mlt <1-512> name WORD<0-64>
```

6. Enable trunking on the MLT:

```
mlt <1-512> encapsulation dot1q
```

7. Enable the MLT:

```
mlt <1-512> enable
```

8. Display the MLT configuration:

```
show mlt <1-512>
```

Examples

Create MLT 40. Configure the private VLAN type for MLT 40 to isolated. Add VLAN 10 to the MLT. Enable the MLT.

```
Switch:1(config)#mlt 40
Switch:1(config)#mlt 40 private-vlan isolated
Switch:1(config)#mlt 40 vlan 10
Switch:1(config)#mlt 40 enable
```

Display the MLT configuration:

```
Switch:1(config)#show mlt 40
```

Mlt Info							
MLTID	IFINDEX	NAME	PORT TYPE	MLT ADMIN	MLT CURRENT	PORT MEMBERS	VLAN IDS
40	6183	MLT-40	access	norm	norm		CONFIG 10
MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER			
40	6183	null	disable	down			
MLTID	NAME	WHERE CREATED	LOCAL PORT MEMBERS	REMOTE PORT MEMBERS	WHICH PORTS PROGRAMMED IN DATA PATH		
40	MLT-40	LOCAL			NONE		
MLTID	IFINDEX	ENCAP DOT1Q	LOSSLESS	PVLAN	PVLAN TYPE	VID TYPE	FLEX-UNI
40	6183	disable	disable	enable	isolated	secondary	disable

Variable Definitions

Use the data in the following table to use the **mlt** command.

Variable	Value
<i>enable</i>	Creates and enables a new MLT.
<i>encapsulation dot1q</i>	Enables trunking on the MLT.
<i>member {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	<p>Adds ports to the MLT.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i><1-512></i>	Specifies the MLT ID.
<i>name WORD<0-64></i>	Configures the name for the MLT.
<i>private-vlan {isolated promiscuous trunk}</i>	Specifies a private VLAN type for the MLT.
<i>vlan <1-4059></i>	Specifies a VLAN ID to add to the MLT.

Display MLT Names

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the names of MLTs on the switch:

```
show mlt name [<1-512>]
```

Example

Display the name for MLT 2:

```
Switch:1>show mlt name 2
=====
Mlt Name
=====
MLTID IFINDEX NAME
-----
2      6145    BN1-BN8
```

View MLT Statistics

About This Task

View MLT statistics to display MultiLink Trunking statistics for the switch or for the specified MLT ID.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View MLT statistics:

```
show mlt stats [<1-512>]
```

Example

```
Switch:1#show mlt stats
```

```
=====
```

Mlt Interface				
ID	IN-OCTETS	OUT-OCTETS	IN-UNICST	OUT-UNICST
1	256676904	183670662	1397	456
2	61737348498	61584347982	1450182	1490619
4	229256124	47472778	0	0
100	251678170	32332107	0	0

```
-----
```

ID	IN-MULTICST	OUT-MULTICST	IN-BROADCAST	OUT-BROADCAST	MT
1	2419514	2295274	41	268194	E
2	962303832	960067410	765	237	E
4	2159884	666153	0	90	E
100	2095269	504965	13	0	E

```
-----
```

ID	IN-LSM	OUT-LSM
1	0	0
2	957925732	957929399
4	0	0

```
-----
```

```
--More-- (q = quit)
```

Variable Definitions

Use the data in the following table to help you use the **show mlt stats** command.

Variable	Value
<1-512>	Specifies the MLT ID.

View MLT Port Members

View the port members in the specified MLT, or for all MLTs.

MLTID	IFINDEX	NAME	TYPE	ADMIN	CURRENT	MEMBERS	ORIGIN	IDS
1	6144	MLT-1	trunk	norm	norm	1/5	CONFIG	1000 2000
MLTID	IFINDEX	DESIGNATED PORTS	LACP ADMIN	LACP OPER				
1	6144	1/5	disable	down				
MLTID	NAME	WHERE CREATED	LOCAL PORT MEMBERS	REMOTE PORT MEMBERS	WHICH PORTS PROGRAMMED IN DATA PATH			
1	MLT-1	LOCAL	1/5		LOCAL			
MLTID	IFINDEX	ENCAP DOT1Q	LOSSLESS	PVLAN	PVLAN TYPE	VID TYPE	FLEX-UNI	
1	6144	enable	disable	disable	-	-	disable	

Variable definitions

Use the data in the following table to use the **show mlt error collision** command.

Variable	Value
<1-512>	Specifies the MLT ID. The value ranges from 1–512.

Add Ports to an MLT LAG

Perform this procedure to add ports to an MLT LAG.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Add ports to an MLT LAG:

```
mlt <1-512> member {slot/port[-slot/port]}[,...][slot/all][all]}
```

Example

Add port 1/24 to MLT 40:

```
Switch:1(config)#mlt 40 member 1/24
```

Variable Definitions

Use the data in the following table to use the **mlt** command.

Variable	Value
<i>enable</i>	Creates and enables a new MLT.
<i>encapsulation dot1q</i>	Enables trunking on the MLT.
<i>member {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	<p>Adds ports to the MLT.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i><1-512></i>	Specifies the MLT ID.
<i>name WORD<0-64></i>	Configures the name for the MLT.
<i>private-vlan {isolated promiscuous trunk}</i>	Specifies a private VLAN type for the MLT.
<i>vlan <1-4059></i>	Specifies a VLAN ID to add to the MLT.

Remove Ports from an MLT LAG

About This Task



Important

Before removing a port member from an MLT, you must first disable the port. This ensures that the other side brings its corresponding port member down. This achieves parity on both sides and avoids traffic disruptions.

The shutdown port requires that you enter the interface GigabitEthernet configuration mode. However, to remove ports from an MLT LAG, you can enter any configuration mode.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable the ports:

```
shutdown port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/  
all][all]}
```

3. Remove ports from an MLT LAG:

```
no mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

Example

Remove port 1/24 from MLT 40:

```
Switch:1(config)#interface GigabitEthernet 1/24  
Switch:1(config-mlt)#shutdown port 1/24  
Switch:1(config-mlt)#no mlt 40 member 1/24
```

Variable definitions

Use the data in the following table to use the **mlt** command.

Variable	Value
<1-512>	Specifies which MLT to add the ports to.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	<p>Specifies the ports to add to the MLT. Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Create an SMLT from an Existing MLT

Before You Begin

- Create an MLT before you create a split in the MLT.
- To create an SMLT from an existing MLT, first shut down the MLT ports. Configure the MLT as SMLT, and then enable the ports again.



Important

Failure to perform this operation can lead to ports in STP blocked state and traffic loss.

- If you are configuring SMLT with vIST, all the SMLT VLANs associated with vIST must have an I-SID.

About This Task

Create an SMLT from an existing MLT to split physical ports between two switches to improve resiliency and provide active load sharing.



Important

To remove an SMLT configuration from the MLT on a vIST switch, first shut down the SMLT MLT ports on its vIST peer.

Failure to perform this operation can lead to Layer 2 loops.

Procedure

1. Log on to the MLT Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface mlt <1-512>
```

2. Create an SMLT from an existing MLT:

```
smlt
```

Example

Create an SMLT on MLT 1:

```
Switch:1(config)#interface mlt 1  
Switch:1(config-mlt)#smlt
```

Virtual interswitch trunk (vIST)

Create a Virtual IST

Use this procedure to create a virtual IST (vIST).

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.

**Important**

- When you create a vIST, VLANs assigned to SMLT ports must have an I-SID assigned.
- If you assign a VLAN to an I-SID on one SMLT-BEB node, then you must create the same VLAN and assign it to the same I-SID on the peer SMLT-BEB node even if no devices are connected to this second node.
- The vIST VLAN must also be associated with an I-SID. You can use the same vIST VLAN in another part of your network, but the I-SID associated with the vIST VLAN must not be used anywhere else.

**Note**

Simplified Virtual-IST (vIST) is for non-SPB customers who use SMLT with legacy IST. Simplified VIST is available for legacy multicast deployments only when the boot flag (**spbm-config-mode**) is disabled.

For more information, see [Configure Simplified vIST in SMLT Topologies](#) on page 2688 .

About This Task

The following list provides an overview of the configuration steps:

- Enable SPBM and IS-IS globally.
- Configure SPBM and IS-IS.
- Configure an Layer 2 VSN by assigning an I-SID to the C-VLAN, which is used by the vIST.

For information about SPBM and IS-IS, see [SPBM and IS-IS Infrastructure Fundamentals](#) on page 1178.

**Important**

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create a vIST:
`virtual-ist peer-ip <A.B.C.D> vlan <1-4059>`
3. Verify your vIST configuration:
`show virtual-ist`

Example

The configuration example contains annotations (in parenthesis) to explain the configuration information.

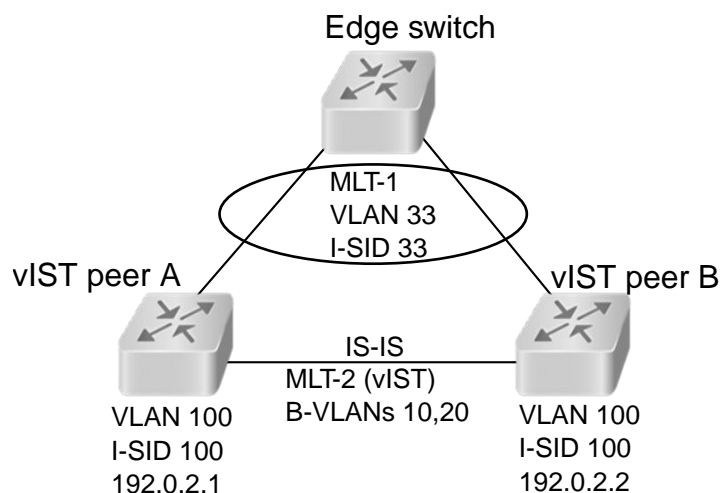


Figure 183: SMLT with vIST configuration example



Note

This configuration example is for vIST peer B.

```
#
# ISIS SPBM CONFIGURATION
#

router isis
spbm 1
spbm 1 nick-name 0.00.16
spbm 1 b-vid 10,20 primary 10
spbm 1 smlt-virtual-bmac 00:00:84:04:01:01
spbm 1 smlt-peer-system-id a051.c6eb.c865
exit

#
# MLT CONFIGURATION
#

mlt 1 enable
mlt 1 member 1/28

mlt 2 enable
mlt 2 member 1/10
mlt 2 encapsulation dot1q

#
# VLAN CONFIGURATION
#

vlan members remove 1 1/28,1/10

(The rest of this VLAN section are all prerequisites for configuring a vIST.)
vlan create 10 type spbm-bvlan (This command and the next create the B-VLANs.)
vlan create 20 type spbm-bvlan
```

```
vlan create 33 type port-mstprstp 0
vlan i-sid 33 33
mlt 1 vlan 33

vlan create 100 type port-mstprstp 0  (Creates the C-VLAN for the vIST VLAN.)
vlan i-sid 100 100  (Creates a Layer 2 VSN for the vIST VLAN.)
interface Vlan 100

ip address 192.0.2.2 255.255.255.0 0  (Assigns an IP to the vIST VLAN.)
exit

#
# VIRTUAL vIST CONFIGURATION
#

virtual-ist peer-ip 192.0.2.1 vlan 100  (vIST configuration.)

#
# MLT INTERFACE CONFIGURATION
#

interface mlt 1
smlt
exit

interface mlt 2
isis
isis spbm 1
isis enable
exit

#
# PORT CONFIGURATION - PHASE II
#

interface GigabitEthernet 1/28
no shutdown
no spanning-tree mstp force-port-state enable
exit

interface GigabitEthernet 1/10
no shutdown
exit

#

# ISIS CONFIGURATION
#

router isis
is-type 11
manual-area 11.1111
exit
router isis enable
```

Variable definitions

Use the data in the following table to use the **ist peer-ip** command.

Variable	Value
<A.B.C.D>	Specifies the peer IP address—the IP address of the vIST VLAN on the other aggregation switch.
<1-4059>	Specifies the VLAN ID for this vIST.

View vIST Statistics

View virtual IST (vIST) statistics for the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the vIST statistics:
show virtual-ist stat
3. To clear the vIST statistics:
clear virtual-ist stats

Example

```
Switch:1#show virtual-ist stat
=====
                        IST Message Statistics
=====
PROTOCOL MESSAGE      COUNT
-----
Ist Down              : 0
Hello Sent            : 0
Hello Recv           : 0
Learn MAC Address Sent : 0
Learn MAC Address Recv : 0
MAC Address AgeOut Sent : 0
MAC Address AgeOut Recv : 0
MAC Address Expired Sent : 0
MAC Address Expired Recv : 0
Delete Mac Address Sent : 0
Delete Mac Address Recv : 0
Smlt Down Sent       : 0
Smlt Down Recv       : 0
Smlt Up Sent         : 0
Smlt Up Recv         : 0
Send MAC Address Sent : 0
Send MAC Address Recv : 0
IGMP Sent            : 0
IGMP Recv            : 0
Port Down Sent       : 0
Port Down Recv       : 0
Request MAC Table Sent : 0
Request MAC Table Recv : 0
Unknown Msg Type Recv : 0
Mlt Table Sync Req Sent : 0
Mlt Table Sync Req Recv : 0
Mlt Table Sync Sent   : 0
Mlt Table Sync Recv   : 0
```

```

Port Update Sent      : 0
Port Update Recv      : 0
Entry Update Sent     : 0
Entry Update Recv     : 0
Dialect Negotiate Sent : 0
Dialect Negotiate Recv : 0
Update Response Sent   : 0
Update Response Recv   : 0
Transaction Que HiWaterM : 0
Poll Count Hi Water Mark : 0

```

Edit a vIST

If you need to change the virtual IST (vIST) **peer-ip** or **vlan**, use this procedure to delete the vIST first.



Note

- You must disable IS-IS globally before deleting a vIST, and then re-enable it after creating a new vIST.
- You do not have to set the SMLT peer system ID or the virtual BMAC to 0 before you change the vIST peer IP address or VLAN ID number.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Disable IS-IS globally:

```
no router isis enable
```
3. Delete the vIST:

```
no virtual-ist peer-ip
```
4. Create a vIST:

```
virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
```
5. Enable IS-IS globally:

```
router isis enable
```

Configure Simplified vIST in SMLT Topologies

This procedure shows how to configure Simplified vIST in an SMLT environment. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

When you enable Simplified vIST with the **virtual-ist enable** command, two VLANs are automatically created to support vIST. The VLAN IDs are: 4086 and 4087.

Before You Begin

- SPBM must not be enabled on the vIST peers.

- Disable PIM Infinite Threshold Policy.

About This Task



Important

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of a header size increase.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable the boot flag:

```
no boot config flags spbm-config-mode
```

The system responds with these messages:

Warning: Please save the configuration and reboot the switch for this to take effect.

Warning: Please carefully save your configuration file before rebooting the switch. Saving configuration file when spbm-config-mode is changed to disable, removes SPBM configurations from the configuration file.

3. Save the configuration, and then reboot the switch.



Important

Any change to the **spbm-config-mode** boot flag requires a reboot for the change to take effect.

4. Create the vIST VLAN:

```
vlan create <2-4059> type port-mstprstp <0-63>
```

```
interface vlan <1-4059>
```

```
ip address <A.B.C.D/X>
```

5. Configure the vIST peer address and VLAN:

```
virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
```

6. Configure the SMLT MLT:

```
mlt <1-512> enable
```

```
mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]
[slot/all][all]}
```

```
interface mlt <1-512>
```

```
smlt
```

7. Configure the vIST MLT:

```
mlt <1-512> enable

mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]
[slot/all][all]}

mlt <1-512> encapsulation dot1q

interface mlt <1-512>

virtual-ist enable
```

**Note**

The **virtual-ist enable** command enables Simplified vIST and is only available when the **spbm-config-mode** boot flag is disabled.

8. Create a customer VLAN and assign the SMLT MLT ID:

```
vlan create <2-4059>

vlan mlt <1-4059> <1-512>

interface vlan <1-4059>

ip address <A.B.C.D/X>
```

Example

```
enable
configure terminal
no boot config flags spbm-config-mode
```

Save the configuration and reboot the switch.

```
virtual-ist peer-ip 192.0.2.1 vlan 50

mlt 3 enable
mlt 3 member 1/35,1/36
interface mlt 3
smlt
exit
mlt 5 enable
mlt 5 member 2/1,2/2
mlt 5 encapsulation dot1q
interface mlt 5
virtual-ist enable
exit
vlan create 50 type port-mstprstp 0
interface vlan 50
ip address 192.0.2.2 255.255.255.0 1
exit
vlan create 100
vlan mlt 100 3
interface vlan 100
ip address 198.51.100.1 255.255.255.0 2
exit
```


Troubleshoot vIST Failure

About This Task

When you use Virtual Inter-Switch Trunk (vIST), all critical network traffic runs on this link. If vIST fails, network protocols such as RIP, VRRP, OSPF, and VLACP go down and eventually cause a network outage.

vIST uses an SPBM tunnel to virtually connect two nodes that can be anywhere in the SPBM cloud. Even if the two vIST nodes are directly connected by an MLT link, the vIST VLAN does not have MLT ports as members. Instead, it is configured to be an SPBM C-VLAN.



Note

For more information on vIST and a configuration example, see [MultiLink Trunking and Split MultiLink Trunking](#) on page 2657.

The vIST tunnel is up as long as there is SPBM connectivity between the IST peers. If there is a vIST failure, check the following procedure for some possible reasons:

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Verify that the vIST VLAN is configured on the vIST switch:

```
show virtual-ist
```

3. Verify that an I-SID is associated with the vIST VLAN:

```
show isis spbm i-sid discover
```



Important

The I-SID associated with the vIST VLAN should be the same on the vIST peer, and this I-SID should not be used anywhere else in the network.

4. Verify that the vIST peers are on the same subnet.
 5. If peer ARP is not resolved, enable **trace level 14** to see if ARP request/response are being sent/received.
 6. If vIST is not up, check the **mac fdb table** and verify that the peer MAC is synchronized:
- ```
show vlan mac-address-entry <1-4059>
```
7. If the vIST peer MAC is learned, check to see if the peer IP address is reachable.
    - a. Use **show virtual-ist** to obtain the vIST peer IP address.
    - b. Ping the peer IP address.
  8. If unable to **ping** the peer IP address, check to see if ARP is resolved.

```
show ip arp vlan <vid>
```

## View All Ports Configured for SMLT

View all ports for an SMLT to ensure the correct ports are configured.

### Procedure

1. Enter Privileged EXEC mode:  
enable
2. View all ports configured for SMLT:  
show smlt mlt

### Example

```
Switch:1#show smlt mlt
```

```
=====
 Mlt SMLT Info
=====
MLT ADMIN CURRENT
ID TYPE TYPE

1 smlt smlt
4 smlt smlt
```

### Variable definitions

Use the data in the following table to use the **show smlt** command.

| Variable | Value                                            |
|----------|--------------------------------------------------|
| mlt      | Displays SMLT information for the MLT interface. |

## View Information about Collision Errors

View information about collision errors to obtain information about collision errors in the specified MLT, or for all MLTs.

### Procedure

1. Enter Privileged EXEC mode:  
enable
2. View information about collision errors:  
show mlt error collision [<1-512>]

### Example

```
Switch:1#show mlt error collision 4
```

```
=====
 Mlt Collision Error
=====
MLT -----COLLISIONS-----
ID SINGLE MULTIPLE LATE EXCESSIVE

40 0 0 0 0
```

### Variable definitions

Use the data in the following table to use the **show mlt error collision** command.

| Variable | Value                                              |
|----------|----------------------------------------------------|
| <1-512>  | Specifies the MLT ID. The value ranges from 1-512. |

## View Information about Ethernet Errors

View information about Ethernet errors to display information about the types of Ethernet errors sent and received by a specific MLT or all MLTs.

### About This Task



#### Important

The IMAC columns refer to internal MAC address errors.

### Procedure

1. Enter Privileged EXEC mode:  
`enable`
2. View information about Ethernet errors:  
`show mlt error main [<1-512>]`

### Example

```
Switch:1#show mlt error main 40
```

```
=====
 Mlt Ethernet Error
=====
MLT ALIGN FCS IMAC IMAC CARRIER FRAMES SQETEST DEFER
ID ERROR ERROR TRNSMIT RECEIVE SENSE TOOLONG ERROR TRNSMSS

40 0 0 0 0 0 0 0 0
```


### Variable definitions

Use the data in the following table to use the **show mlt error main** command.

| Variable | Value                                              |
|----------|----------------------------------------------------|
| <1-512>  | Specifies the MLT ID. The value ranges from 1-512. |

SLPP Guard configuration

This section provides the procedures to configure Simple Loop Prevention Protocol (SLPP) Guard.



**Important**

Enable SLPP Guard on the edge switches of an SMLT network and SLPP on the aggregation layer switches.

Configure SLPP Guard

Configure SLPP Guard on MLT and LAG ports to provide additional network loop protection.

SLPP Guard is disabled by default.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
- ```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable SLPP Guard on the port:
- ```
slpp-guard enable
```
3. (Optional) Configure the timeout value on the port:
- ```
slpp-guard timeout {<10-65535> | 0}
```

Example

```
Switch:1(config-if)#slpp-guard enable
Switch:1(config-if)#slpp-guard timeout 120
```

Variable definitions

Use the data in the following table to use the **slpp-guard** command.

Variable	Value
timeout <0 10-65535>	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch reenables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default timeout value is 60 seconds.

Reenable an Operationally Disabled Port

Reenable a port that has been operationally disabled by SLPP Guard.



Note

You cannot reenable a disabled port if the timer count has not reached its timeout value. Either wait until it reaches the timeout or disable SLPP Guard for that port and then re-enable it.

Procedure

- 1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- 2. Reenable SLPP Guard on the port:
slpp-guard enable
- 3. (Optional) Configure the timeout value on the port:
slpp-guard timeout {<10-65535> | 0}

Example

```
Switch:1(config-if)#slpp-guard enable
Switch:1(config-if)#slpp-guard timeout 120
```

Variable definitions

Use the data in the following table to use the **slpp-guard** command.

Variable	Value
timeout <0 10-65535>	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch reenables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default timeout value is 60 seconds.

View SLPP Guard Status

View current SLPP Guard settings.

Procedure

- 1. To enter User EXEC mode, log on to the switch.

2. Display SLPP Guard status:

```
show slpp-guard {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]
[slot/all] [all]}
```

Example

```
Switch:1#>show slpp-guard
```

SLPP Guard							
SLPP-guard Ethertype: 0x8102							
Port Interface							
Port	Link	Oper	SLPP-guard State		Timeout	TimerCount	ORIGIN
1/1	Up	Down	Disabled	N/A	60	N/A	RADIUS
1/2	Up	Down	Enabled	N/A	120	N/A	RADIUS
1/3	Up	Up	Enabled	Monitoring	60	N/A	RADIUS
1/4	Up	Down	Enabled	N/A	120	N/A	RADIUS
1/5	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/6	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/7	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/8	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/9	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/10	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/11	Down	Down	Disabled	N/A	60	N/A	RADIUS
1/12	Down	Down	Disabled	N/A	60	N/A	RADIUS



Note

The TimerCount column in the preceding example output indicates the time, in seconds, that elapses after SLPP Guard disables a port. When the value TimerCount value equals the Timeout value, the switch re-enables the port.

Variable definitions

Use the data in the following table to use the **show slpp-guard** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

MLT and SMLT Link Aggregation Configuration using EDM

Configure link aggregation to provide link level redundancy and increase load sharing. MultiLink Trunking (MLT) is a link aggregation technology that you can use to group

several physical Ethernet links into one logical Ethernet link to provide fault-tolerance and high-speed links between routers, switches, and servers. Split MultiLink Trunking (SMLT) is an option that improves Layer 2 (bridged) resiliency.

Add a Multilink or LACP Trunk

Perform this procedure to add a multilink or LACP trunk.

About This Task



Important

Ensure that all ports that belong to the same MLT/LACP group use the same port speed, for example, 1 Gbit/s, even if autonegotiation is used. This requirement is not enforced by the software.

When you add a VLAN to a dynamic MLT, only the active ports of the MLT are added as port members of the VLAN. Ports configured with the same aggregation key, but not active, are not added to the VLAN. If these inactive ports become active later, the system does not automatically add them to the VLAN port member list.

You must add all inactive ports to the VLAN. If you do not add the inactive ports to the VLAN, when they become active later, hashing can result in choosing a newly active port for traffic forwarding. Because the port is not a port member of the VLAN, traffic will be dropped. When you add the VLAN to the MLT, also add the inactive aggregation ports to the VLAN. You may have to disable LACP on the inactive ports before you can add them to the VLAN. Because the ports are inactive, disabling LACP does not cause a traffic interruption.

Similarly when you remove a VLAN from a dynamic MLT, all active ports of the MLT are removed from the VLAN port member list but the inactive members are not removed. You must remove the inactive aggregation members from the VLAN.

If you later configure a port for the same aggregation, you must add this port to all VLANs that are members of the MLT.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **MLT/LACP**.
3. Select the **MultiLink/LACP Trunks** tab.
4. Select **Insert**.
5. In **Id**, type the ID number of the MLT.
6. In **PortType**, select **access** or **trunk**.
7. (Optional) In **Name**, type a name for the MLT, or accept the default.
8. In **PortMembers**, select the ellipsis (...).
9. In **Port Editor: PortMembers**, select the desired ports.
10. Select **Ok**.
11. In **VlanIdList**, select the ellipsis (...).
12. In **VlanIdList**, select the desired VLANs.

13. Select **Ok**.
14. In **MltType**, select the MLT type.
15. (Optional) In **Aggregatable**, select **enable** or **disable**.
16. (Optional) In **PrivateVlanType**, select **trunk**, **isolated**, or **promiscuous**.
17. Select **Insert**.

MultiLink/LACP Trunks *Field Descriptions*

Use the data in the following table to use the **MultiLink/LACP Trunks** tab.

Name	Description
Id	Specifies a unique value for this MLT.
PortType	Specifies the value to access or trunk port. If the aggregatable field is set to enable, this field is read-only. The default value is access.
Name	Configures the name given to the MLT.
PortMembers	Assigns ports to the MLT. All ports in an MLT must have the same settings for speed and duplex, but can have different media types. All untagged ports must belong to the same STG. Up to eight same-type ports can belong to a single MLT. If the aggregatable field is set to enable, this field becomes read-only. Important: Ensure that all ports that belong to the same MLT/LACP group use the same port speed, for example, 1 Gbps, even if autonegotiation is enabled.
VlanIdList	Indicates to which ports the VLANs belong. If the aggregation field is set to enable, this field is read-only.
MltType	Specifies the type of MLT <ul style="list-style-type: none"> • normalMLT (default) • istMLT • splitMLT
RunningType	Specifies the MLT running type.
IfIndex	Specifies the interface of the trunk.
ClearLinkAggregate	Clears the link aggregate, disabling and reenabling the trunk.
DesignatedPort	Specifies the designated port of this trunk.

Name	Description
Origin	Displays if the MLT is configured manually or dynamically.
Aggregatable	Enables or disables link aggregation. The default value is disable.
AggOperState	Specifies the aggregation state of the trunk.
AggTimeOfLastOperChange	Specifies the time value since the interface entered its current operational state.
PeerPortMembersList	Specifies the peer ports connected to the local ports of this trunk.
EntryOwner	Defines the owner of the MLT.
DatapathProgrammingState	Defines the datapath programming state of the MLT.
PrivateVlanType	Specifies the type of private VLAN for the MLT.
FlexUniEnable	Specifies whether the Flex UNI is enabled or disabled on the port. The default value is disable.
NodeAliasClear	Specifies whether the Node Alias is enabled or disabled for the MLT.

Add Ports to an MLT

Add ports to an MLT to insert MultiLink/LACP trunks.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **MLT/LACP**.
3. Select the **MultiLink/LACP Trunks** tab.
4. In the **PortMembers** column, double-click the field associated with the MLT to which you want to add ports to.
5. Select the port numbers to add, or select **All** to add all ports to the MLT.
Up to 16 same-type ports can belong to a single MLT.
6. Select **Ok**.
7. Select **Apply**.

View SMLT Statistics

View SMLT statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.

2. Select **MLT/LACP**.
3. Select the **Ist/SMLT Stats** tab.

IST/SMLT Stats Field Descriptions

The following table describes parameters on the **IST/SMLT Stats** tab.

Name	Description
SmltIstDownCnt	The number of times the session between the two peering switches has gone down since last boot.
SmltHelloTxMsgCnt	The count of transmitted hello messages.
SmltHelloRxMsgCnt	The count of received hello messages.
SmltLearnMacAddrTxMsgCnt	The count of transmitted learned MAC address messages.
SmltLearnMacAddrRxMsgCnt	The count of received learned MAC address messages.
SmltMacAddrAgeOutTxMsgCnt	The count of transmitted aging out MAC address messages.
SmltMacAddrAgeOutRxMsgCnt	The count of received aging out MAC address messages.
SmltMacAddrAgeExpTxMsgCnt	The count of transmitted MAC address age expired messages.
SmltMacAddrAgeExpRxMsgCnt	The count of received MAC address age expired messages.
SmltStgInfoTxMsgCnt	The count of transmitted STG information messages.
SmltStgInfoRxMsgCnt	The count of received STG information messages.
SmltDelMacAddrTxMsgCnt	The count of transmitted MAC address deleted messages.
SmltDelMacAddrRxMsgCnt	The count of received MAC address received messages.
SmltSmltDownTxMsgCnt	The count of transmitted SMLT down messages.
SmltSmltDownRxMsgCnt	The count of received SMLT down messages
SmltUpTxMsgCnt	The count of transmitted SMLT up messages.
SmltUpRxMsgCnt	The count of received SMLT up messages.
SmltSendMacTblTxMsgCnt	The count of sent send MAC table messages.
SmltSendMacTblRxMsgCnt	The count of received send MAC table messages.
SmltIgmpTxMsgCnt	The count of sent IGMP messages.

Name	Description
SmltIcmpRxMsgCnt	The count of received ICMP messages.
SmltPortDownTxMsgCnt	The count of sent port down messages.
SmltPortDownRxMsgCnt	The count of received port down messages.
SmltReqMacTblTxMsgCnt	The count of sent MAC table request messages.
SmltReqMacTblRxMsgCnt	The count of received MAC table request messages.
SmltRxUnknownMsgTypeCnt	The count of received unknown message type messages.
SmltPortTblSyncReqTxMsgCnt	The count of sent sync request messages.
SmltPortTblSyncReqRxMsgCnt	The count of received sync request messages.
SmltPortTblSyncTxMsgCnt	The count of sent sync messages.
SmltPortTblSyncRxMsgCnt	The count of received sync messages.
SmltPortUpdateTxMsgCnt	The count of sent update messages.
SmltPortUpdateRxMsgCnt	The count of received update messages.
SmltEntryUpdateTxMsgCnt	The count of sent entry update messages.
SmltEntryUpdateRxMsgCnt	The count of received entry update messages.
SmltDialectNegotiateTxMsgCnt	The count of sent protocol ID messages.
SmltDialectNegotiateRxMsgCnt	The count of received protocol ID messages.
SmltUpdateRespTxMsgCnt	The count of sent update response messages.
SmltUpdateRespRxMsgCnt	The count of received update response messages.
SmltTransQHighWaterMarkMsgCnt	The count of transaction queue high watermark messages.
SmltPollCountHighWaterMarkCnt	The count of poll count high watermark.

View MLT Interface Statistics

About This Task

Use MLT interface statistics tab to view interface statistics for the selected MLT.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **MLT/LACP**.
3. Select the **MultiLink/LACP Trunks** tab.

4. Select an MLT.
5. Select **Graph**.
6. Select the **Interface** tab.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
InOctets	Specifies the total number of octets received on the MLT interface, including framing characters.
OutOctets	Specifies the total number of octets transmitted out of the MLT interface, including framing characters.
InUcastPkts	Specifies the number of packets delivered by this MLT to higher level protocols that were not addressed to a multicast or broadcast address at this sublayer.
OutUcastPkts	Specifies the number of packets that higher-level protocols requested be transmitted that were not addressed to a multicast address at this MLT. This total number includes discarded or unsent packets.
InMulticastPkt	Specifies the number of packets delivered to this MLT that were addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
OutMulticast	Specifies the total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this MLT, including those that were discarded or unsent. For a MAC layer protocol, this number includes both Group and Functional addresses.
InBroadcastPkt	Specifies the number of packets delivered to this MLT that were addressed to a broadcast address at this sublayer.
OutBroadcast	Specifies the total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this MLT, including those that were discarded or not sent.
InLsmPkts	Specifies the total number of Link State Messaging (LSM) packets delivered on this MLT.
OutLsmPkts	Specifies the total number of Link State Messaging (LSM) packets transmitted on this MLT.

View MLT Ethernet Error Statistics

About This Task

Use MLT Ethernet error statistics to view the error statistics.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.

2. Select **MLT/LACP**.
3. Select the **MultiLink/LACP Trunks** tab.
4. Select an MLT, and then select **Graph**.
5. Select the **Ethernet Errors** tab.

Ethernet Errors *Field Descriptions*

Use the data in the following table to use the **Ethernet Errors** tab.

Name	Description
AlignmentErrors	Specifies the frame count frames received on a particular MLT that is not an integral number of octets in length and does not pass the FCS check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
FCSErrors	Specifies the frame count received on an MLT that is an integral number of octets in length, but does not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object increments when the FrameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
IMacTransmitError	Specifies the frame count for which transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
IMacReceiveError	Specifies the frame count for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object. The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent receive errors on a particular interface that are not otherwise counted.
CarrierSenseError	Specifies the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular MLT. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.

Name	Description
FrameTooLong	Specifies the frame count received on a particular MLT that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
SQETestError	Specifies the number of times that the SQE test error message is generated by the PLS sublayer for a particular MLT. The SQE test error message is defined in section 7.2.2.2.4 of ANSI/ IEEE 802.3-1985.
DeferredTransmiss	Specifies the frame count for which the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
SingleCollFrames	Specifies a count of successfully transmitted frames on a particular MLT for which transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the MultipleCollisionFrames object.
MultipleCollFrames	Specifies the successfully transmitted frame count on a particular MLT for which transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the SingleCollisionFrames object.
LateCollisions	Specifies the number of times that a collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512 corresponds to 51.2 microseconds on a 10 Mb/s system. A late collision included in a count represented by an instance of this object is also considered as a generic collision for purposes of other collision-related statistics.
ExcessiveCollis	Specifies the frame count for which transmission on a particular MLT fails due to excessive collisions.

View Trunks

Perform this procedure to view the MLT-based SMLT configuration.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SMLT**.

3. Select the **SMLT Info** tab.

SMLT Info Field Descriptions

Use the data in the following table to use the **SMLT Info** tab.

Name	Description
Id	Shows the MLT ID for this SMLT Read-only.
MltType	Shows the MLT type of this trunk.
RunningType	Shows the SMLT running type.

Create a Virtual IST using EDM

vIST creates a virtualized channel through the SPBM cloud, and this channel connects two SMLT devices to form a virtualized Switch Cluster. Note that the SPBM cloud can consist of as few as two nodes.



Important

- When you create a vIST, VLANs assigned to SMLT ports must have an I-SID assigned.
- If you assign a VLAN to an I-SID on one SMLT-BEB node, then you must create the same VLAN and assign it to the same I-SID on the peer SMLT-BEB node even if no devices are connected to this second node.
- The vIST VLAN must also be associated with an I-SID. You can use the same vIST VLAN in another part of your network, but the I-SID associated with the vIST VLAN must not be used anywhere else.



Note

Simplified Virtual-IST (vIST) is for non-SPB customers who use SMLT with legacy IST. Simplified VIST is available only for legacy multicast deployments when the boot configuration flag (**spbm-config-mode**) is disabled.

Before You Begin

- Enable SPBM and IS-IS globally.
- Configure SPBM and IS-IS.
- Configure a Layer 2 VSN by assigning an I-SID to the C-VLAN, which is used by the vIST.

About This Task



Important

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of the header size increase when transmitting packets over the SPBM cloud.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > MLT/LACP**.
2. Select the **Virtual IST** tab.
3. In the **PeerIp** field, type the peer IP address.
4. In the **VlanId** field, enter a VLAN ID.
5. Select **Apply**.

Virtual IST Field Descriptions

Use the data in the following table to use the **Virtual IST** tab.

Name	Description
SessionStatus	Displays the status of the vIST session.
PeerIp	Specifies the peer IP address, which is the IP address of the vIST VLAN on the other aggregation switch.
VlanId	Configures a vIST VLAN ID number.

Edit a Virtual IST

If you need to change the virtual IST **PeerIp** or **VlanId**, use this procedure to delete the vIST first.



Note

- You must disable IS-IS globally before deleting a vIST, and then re-enable it after creating a new vIST.
- You do not have to set the SMLT peer system ID or the virtual BMAC to 0 before you change the virtual IST peer IP address or VLAN ID number.

Procedure

1. Disable IS-IS globally.
 - a. In the navigation pane, expand **Configuration > Fabric > IS-IS**.
 - b. Select the **Globals** tab.
 - c. In the **AdminState** field, select **off**.
 - d. Select **Apply**.
2. Delete the vIST.
 - a. In the navigation pane, expand **Configuration > VLAN > MLT/LACP**.
 - b. Select the **Virtual IST** tab.
 - c. Select **Apply**.
3. Create a vIST:
 - a. In the navigation pane, expand **Configuration > VLAN > MLT/LACP**.
 - b. Select the **Virtual IST** tab.
 - c. In the **PeerIp** field, enter the peer IP address.

- d. In the **VlanId** field, enter a VLAN ID.
 - e. Select **Apply**.
4. Enable IS-IS globally.
 - a. In the navigation pane, expand **Configuration > Fabric > IS-IS**.
 - b. Select the **Globals** tab.
 - c. In the **AdminState** field, select **on**.
 - d. Select **Apply**.

Configure Simplified vIST in SMLT Topologies

This procedure shows how to configure Simplified vIST in an SMLT environment. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST



Important

When you enable Simplified vIST with the **virtual-ist enable** command, two VLANs are automatically created to support vIST. The VLAN IDs are: 4086 and 4087.

Before You Begin

SPBM must not be enabled on the vIST peers.

About This Task



Important

Do not change the system MTU to less than the default value of 1950 bytes. The system MTU must be 1950 or jumbo because of a header size increase.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Boot Config** tab.
3. Clear **EnableSpbmConfigMode** to disable the boot flag.

The system responds with the following messages:

```
Warning: Please save the configuration and reboot the switch for this
to take effect.
```

```
Warning: Please carefully save your configuration file before
rebooting the switch. Saving configuration file when spbm-config-
mode is changed to disable, removes SPBM configurations from the
configuration file.
```

4. Select **Apply**.

5. Save the configuration, and then reboot the switch.

**Important**

A change to the **EnableSpbmConfigMode** boot flag requires a reboot for the change to take effect.

6. Configure the SMLT MLT:
 - a. Expand **Configuration > VLAN**.
 - b. Select **MLT/LACP**.
 - c. Select the **MultiLink/LACP Trunks** tab.
 - d. Select **Insert**.
 - e. In **Id**, type the ID number of the MLT.
 - f. In **PortType**, select **access** or **trunk**.
 - g. In **Name**, type a name for the MLT, or accept the default.
 - h. In **PortMembers**, select the ellipsis (...).
 - i. In **Port Editor: PortMembers** dialog box, select the desired ports.
 - j. Select **Ok**.
 - k. In **VlanIdList**, select the ellipsis (...).
 - l. In **VlanIdList** dialog box, select the desired VLANs.
 - m. Select **Ok**.
 - n. In **MltType**, select **splitMLT**.
 - o. Select **Insert**.

The switch adds the SMLT MLT to the MultiLink/LACP Trunks tab in the MLT_LACP box.
7. Configure the vIST MLT:
 - a. Repeat steps [6.a](#) on page 2708 to [6.o](#) to configure the MLT.
 - b. Select **virtualistMLT** to enable Simplified vIST.
 - c. Select **Insert**.
8. Create the vIST VLAN:
 - a. Expand **Configuration > VLAN > VLANs**.
 - b. In **Basic**, select **Insert**.
 - c. In **Id**, type an unused VLAN ID, or use the ID provided.
 - d. In **MstpInstance**, select the down arrow, and then choose an MSTI instance from the list.
 - e. In **Type**, select **byPort**.
 - f. In **PortMembers**, select the ellipsis (...) .
 - g. In **Port Editor: PortMembers** dialog box, select the desired ports.
 - h. Select **OK**.
 - i. Select **Insert**.
 - j. Select the vIST VLAN from the list of VLANs, and then select **IP**.
 - k. Select **Insert**.
 - l. Configure the IP address for the vIST VLAN.
9. Repeat step [8](#) to create an SMLT VLAN and assign the SMLT MLT ID to it. Do not use the vIST MLT ID.

SLPP Guard Configuration

This section provide the procedures to configure Simple Loop Prevention Protocol (SLPP) Guard using EDM.



Important

Enable SLPP Guard on the edge switches of an SMLT network and SLPP on the aggregation layer switches.

Configure SLPP Guard

Configure SLPP Guard on MLT and LAG ports to provide additional network loop protection.

SLPP Guard is disabled by default.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SLPP**.
3. Select the **SLPP Guard** tab.
4. In the port row, double-click the cell in the **Enabled** column.
5. Select true from the drop-down list to enable SLPP Guard, or false to disable SLPP Guard for the port.
6. (Optional) In the port row, double-click the cell in the **Timeout** column.
7. (Optional) Type a value in the **Timeout** field.
8. Select **Apply**.
9. On the toolbar, select **Refresh** to update the work area data display.

SLPP Guard Field Descriptions

Use the data in the following table to use the **SLPP Guard** tab.

Name	Description
IfIndex	Specifies the port on which to configure SLPP Guard.
Enable	Enables (true) or disables (false) SLPP Guard for the port. The default is disabled.
Timeout	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch re-enables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default Timeout value is 60 seconds.
Status	Displays the SLPP Guard status for the port.

Name	Description
TimerCount	Indicates the time, in seconds, that elapses after SLPP Guard disables a port. When the TimerCount value equals the Timeout value, the switch re-enables the port.
GuardOrigin	Specifies the origin of SLPP Guard configuration on the port. The supported values are: <ul style="list-style-type: none">• config - Set by the user.• radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.

Reenable an Operationally Disabled Port

Reenable a port that has been operationally disabled by SLPP Guard.



Note

You cannot reenable a disabled port if the timer count has not reached its timeout value. Either wait until it reaches the timeout or disable SLPP Guard for that port and then re-enable it.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SLPP**.
3. Select the **SLPP Guard** tab.
4. In the port row, double-click the cell in the **Enabled** column.
5. Select true from the drop-down list to enable SLPP Guard, or false to disable SLPP Guard for the port.
6. (Optional) In the port row, double-click the cell in the **Timeout** column.
7. (Optional) Type a value in the **Timeout** field.
8. Select **Apply**.
9. On the toolbar, select **Refresh** to update the work area data display.

View SLPP Guard Status

View current SLPP Guard settings.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SLPP**.
3. Select the **SLPP Guard** tab.

SLPP Guard Field Descriptions

Use the data in the following table to use the **SLPP Guard** tab.

Name	Description
IfIndex	Specifies the port on which to configure SLPP Guard.
Enable	Enables (true) or disables (false) SLPP Guard for the port. The default is disabled.
Timeout	Specifies the time period, in seconds, for which SLPP Guard disables the port. After the timeout period expires, the switch re-enables the port. The timeout value can be 0 or a value ranging from 10 to 65535. With a value of 0, the port remains disabled until it is manually re-enabled. The default Timeout value is 60 seconds.
Status	Displays the SLPP Guard status for the port.
TimerCount	Indicates the time, in seconds, that elapses after SLPP Guard disables a port. When the TimerCount value equals the Timeout value, the switch re-enables the port.
GuardOrigin	Specifies the origin of SLPP Guard configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.

MLT Configuration Examples

This section contains configuration examples for configuring MultiLink Trunking (MLT) and MLT with Link Aggregation Control Protocol (LACP) using the Command Line Interface (CLI).



Note

Examples and network illustrations in this document may not include your specific platform. Unless otherwise noted, the concept illustrated applies to all supported platforms. For information about high-level feature support, see *Fabric Engine and VOSS Feature Matrix*.

MultiLink Trunking

This configuration example shows you how to create a multilink trunk and a Virtual Local Area Network (VLAN) between two switches.

The following illustration shows you an MLT within a VLAN used to carry user traffic.




Figure 184: MLT within a VLAN

Switch A and B configuration

Configure MLT 10 on VLAN 100 on each device.

```
#
# MLT CONFIGURATION
#

mlt 10 enable
interface mlt 10
exit
#
# VLAN CONFIGURATION
#
vlan create 100 name VLAN-MLT-10 type port-mstprstp 0 color 1
vlan mlt 100 10
vlan members 100 1/1-1/4 portmember
```

MultiLink Trunking with Link Aggregation Control Protocol

This configuration example shows you how to configure and enable a multilink trunk with LACP.

You must configure all aggregatable ports to use the same aggregator key so they can form an MLT.

The following illustration shows you an MLT created with LACP within a VLAN used to carry user traffic.




Figure 185: MLT within a VLAN

The following procedures show you how to configure both switches.

Switch A and B configuration

Configure the two devices to connect with MLT with LACP.

```
#
# MLT CONFIGURATION
#

mlt 10 enable
interface mlt 10
```

```
lacp enable key 10
exit
#
# VLAN CONFIGURATION
#
vlan create 100 name VLAN-MLT-10" type port-mstprstp 0 color 1
vlan mlt 100 1
vlan mlt 100 10
vlan members 100 1/1-1/4 portmember

#
# PORT CONFIGURATION - PHASE II
#

interface GigabitEthernet 1/1
lacp key 10 aggregation enable
lacp enable
exit
interface GigabitEthernet 1/2
lacp key 10 aggregation enable
lacp enable
exit
```

MLT Network Topology and Configuration Reference

The following reference information contains examples of MLT network topology and configuration. The same topologies apply to MLT with LACP.



Note

Examples and network illustrations in this document may not include your specific platform. Unless otherwise noted, the concept illustrated applies to all supported platforms. For information about high-level feature support, see *Fabric Engine and VOSS Feature Matrix*.

Example 1: Switch-to-switch MLT

The following illustration shows two multilink trunks (MLT1 and MLT2) connecting three switches.




Figure 186: Switch-to-switch multilink trunks configuration

In this example, you can configure each trunk with multiple switch ports to increase bandwidth and redundancy. If traffic between switch-to-switch connections approaches single port bandwidth limitations, you can create a multilink trunk to supply the additional bandwidth required to improve performance, as well as providing physical link layer redundancy.

Example 2: Switch-to-server MLT

In this example, File server 1 uses dual MAC addresses, with one MAC address for each Network Interface Card (NIC). No multilink trunk is configured on File server 1. File server 2 is a single MAC server (with a four-port NIC) configured as multilink trunk configuration MLT1.

The following illustration shows a typical switch-to-server multilink trunk configuration.




Figure 187: Switch-to-server multilink trunk configuration

In this example, one port on File server 1 is blocked and unused and File server 2 benefits from aggregated bandwidth on multilink trunk T1.

Example 3: Client/Server MLT

In this example, both servers are connected directly to the switch. File server 2 is connected through multilink trunk MLT 1. The switch-to-switch connections are through MLT 2, MLT 3, and MLT 4. Clients access data from the servers (File server 1 and File server 2) and receive maximized bandwidth through MLT 1, MLT 2, MLT 3, and MLT 4.

The following illustration shows how you can use multilink trunks in a client/server configuration.




Figure 188: Client/server multilink trunk configuration



Network Operating System Personalities

[Base MAC Address Assignment for Universal Hardware Switches](#) on page 2719

[Change the Network Operating System Personality](#) on page 2719

Table 210: Network operating system personality product support

Feature	Product	Release introduced
Dual-boot personalities	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Applicable
	VSP 7400 Series	Not Applicable

Extreme Networks offers universal hardware products that support more than one Network Operating System (NOS) personality. These hardware products ship with a default NOS personality but you can select a non-default personality. After you log in using the default NOS user name and password for the first time, you can respond to an initialization question to change the NOS personality.

The following table identifies hardware products capable of running Fabric Engine and indicates the default NOS personality, if multiple are supported.

Table 211: Hardware with Fabric Engine support

Product	NOS support
4220 Series	Switch Engine (default) Fabric Engine
5320 Series	Switch Engine (default) Fabric Engine

Table 211: Hardware with Fabric Engine support (continued)

Product	NOS support
5420 Series	Switch Engine (default) Fabric Engine Note: VOSS support on this platform ends with VOSS 8.5.x
5520 Series	Switch Engine (default) Fabric Engine Note: VOSS support on this platform ends with VOSS 8.5.x
5720 Series	Switch Engine (default) Fabric Engine
7520 Series	Switch Engine (default) Fabric Engine
7720 Series	Switch Engine (default) Fabric Engine
7830 Series	Fabric Engine only

Network Operating System Selection

The primary method to select a NOS personality is by using ExtremeCloud IQ or ExtremeCloud IQ Site Engine. To select a NOS personality using ExtremeCloud IQ, see *Read Me First – Universal Hardware*. If the network is not accessible, or if you do not use Extreme Networks management software, you can also change the NOS personality manually by using CLI commands in the running NOS.

The procedure in this document includes the commands you use if the personality is Fabric Engine and you need to change the NOS personality manually. To change the personality to Fabric Engine, see the documentation for the running NOS.

Fabric Engine Default Behavior

The following list provides details about the default behavior after you deploy a universal hardware product with the Fabric Engine personality:

- Port behavior:
 - All ports are isolated members of private VLAN 4040 and 4049.
 - The Segmented Management Instance Management VLAN (non-routing mode) and Management OOB interfaces are enabled.
 - All ports are administratively enabled.
 - To provide loop protection, no switching occurs between ports.

- DHCP:
 - DHCP client requests are cycled between In-Band, all ports untagged, and Out-of-Band ports.
- DNS:
 - The DHCP client issues a DNS query on the interface selected by Zero Touch Deployment to look for ExtremeCloud IQ. For information about Zero Touch Deployment, see [Zero Touch Deployment](#) on page 91.
 - You can then use ExtremeCloud IQ to upgrade or change the switch configuration. If you do not use ExtremeCloud IQ, you can use ExtremeCloud IQ Site Engine or on-switch user interfaces.

Base MAC Address Assignment for Universal Hardware Switches

When running Switch Engine, each universal hardware switch uses a base MAC address at offset 0 for both the default management port, if available, and in-band VLAN utilizing DHCP, for example, 00:c0:cc:8b:68:00. When the switch runs Fabric Engine, it uses a base MAC at offset 0x81 for the default management port (for example, 00:c0:cc:8b:68:81) and offset 256 for the in-band VLAN (for example, 00:c0:cc:8b:69:00).



Note

The address assignment for the in-band VLAN assumes that the VLAN has a mac-offset value of 0 assigned. If a different mac-offset value is assigned, the MAC address changes accordingly. For example, if mac-offset is 10, then the associated MAC address is 00:c0:cc:8b:69:0A.

When using a DHCP client on the switch, the switch sends a common DHCP client identifier equal to the base MAC address of the switch that is printed on the switch label. Because of this, assuming a standard DHCP pool configuration, the DHCP server always recognizes the switch by the same IP address, regardless of whether Switch Engine or Fabric Engine runs on the switch.

To statically assign IP addresses on the DHCP server, assign them based upon the DHCP client ID. This assignment will ensure that the bindings do not change when the switch alternates between Switch Engine and Fabric Engine. If you assign the DHCP IP addresses based on MAC addresses, configure multiple entries, one for the 0 offset and one for the 0x81 offset, to account for the different ways in which the two operating systems assign base MAC addresses.

Change the Network Operating System Personality

Use this procedure on a Fabric Engine switch to change the Network Operating System (NOS) personality.

The primary method to select a NOS personality for the switch is by using ExtremeCloud IQ. If the network is not accessible, or if you do not use Extreme Networks management software, you can change the NOS personality by using Fabric Engine CLI commands.

Before You Begin

Use FTP to transfer the NOS image to the `/intflash/` directory.

About This Task

If you change the NOS personality, the system deletes all configuration, licensing, and log files that pertain to the previous NOS personality.

If a power interruption occurs during a personality change, the system restarts the process after you restore power to the switch.

As part of the NOS personality change, you must restart the switch. If, during the restart, the system detects a failure, the NOS personality does not change. After the original NOS restarts, it logs an error message and raises a persistent alarm.



Note

This procedure does not provide information about upgrading Fabric Engine. For details about upgrading the Fabric Engine release, see [Upgrade the Software](#) on page 233.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Unpack the NOS image:

```
software add WORD<1-99> -y
```



Note

The NOS image file is removed after the **software add** command runs.

3. Install the NOS image:

```
software activate WORD<1-99>
```

4. Confirm you want to change the NOS personality by entering `y`.

5. Confirm the new image has been added and will become active after the next restart:

```
show software
```



Important

If you decide not to change the NOS personality, before you restart the switch, you can run the **software activate** command again with the primary release as the image name and the system cancels the NOS personality change.

6. Restart the switch.

Example

Change the NOS personality to Switch Engine.

```
5520-24W-FabricEngine:1>enable
5520-24W-FabricEngine:1#software add summit_arm-31.6.0.459.xos
Switch Engine image validated successfully
Adding Switch Engine software - /intflash/summit_arm-31.6.0.459.xos
Extraction of /intflash/summit_arm-31.6.0.459.xos to /intflash/release/
summit_arm-31.6.0.459.xos successful
5520-24W-FabricEngine:1#show software
=====
                        software releases in /intflash/release/
=====
summit_arm-31.6.0.459.xos
5520.8.5.0.0int028                (Primary Release)                (Signed Release)

-----
Auto Commit      : enabled
Commit Timeout   : 10 minutes
5520-24W-FabricEngine:1#software activate summit_arm-31.6.0.459.xos
WARNING: The specified image is for the Switch Engine Network Operating System and the
Fabric Engine Network Operating System is currently running.
If you continue, all data including configurations, logs and debugs will be cleared,
except for the license activation status. Switch Engine will be installed and Fabric
Engine image files will be removed. Do you want to continue (y/n) ? y
Executing software activate for version summit_arm-31.6.0.459.xos.
Validated Switch Engine software
Activating Switch Engine software

Primary Version:  summit_arm-31.6.0.459.xos
Backup Version:   5520.8.5.0.0int028.64

Changes will take effect on next reboot.
```

Variable Definitions

The following table defines parameters for the **software** command.

Variable	Value
<i>activate</i> <i>WORD</i> <1-99>	Specifies the software version to copy to the boot flash file.
<i>add</i> <i>WORD</i> <1-99>	Specifies the software version to unpack.
-y	Suppresses the confirmation message to automatically overwrite the non-primary image. If you omit this parameter, you must confirm the action to overwrite the non-primary image.



Network Time Protocol

[NTP Fundamentals](#) on page 2724

[NTP Configuration Using CLI](#) on page 2728

[NTP Configuration Using EDM](#) on page 2735

Table 212: Network Time Protocol product support

Feature	Product	Release introduced
NTPv4 client for IPv4	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
NTPv4 client for IPv6	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 212: Network Time Protocol product support (continued)

Feature	Product	Release introduced
NTPv4 master and restrict	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
NTP Authentication Key Obfuscation	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8
NTP Server configuration using fully qualified domain name (FQDN)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.10
	5420 Series	Fabric Engine 8.10
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.10
	VSP 7400 Series	VOSS 8.10

The following sections provide information on NTPv4.

**Important**

For NTPv4:

- The switch can operate as both NTPv4 client and NTPv4 server.
- The server selection algorithm can deem a server to be unfit to sync even though there is connectivity.
- It can take several iterations (intervals) for the server to sync.
- You must configure a Segmented Management Instance on applicable switches before you use NTPv4.

NTP Fundamentals

This section provides conceptual material about the Network Time Protocol (NTP). Review this content before you make changes to the NTP configuration.

NTP Overview

NTP synchronizes the internal clocks of various network devices across large, diverse networks to universal standard time. NTP protocol runs over the User Datagram Protocol (UDP), which in turn runs over IP.

The NTPv4 specification is documented in RFC 5905 and supports both IPv4 and IPv6 addresses. If you configure the server using a fully qualified domain name (FQDN), the switch resolves the name to an IP address.

Every network device relies on an internal system clock to maintain accurate time. On local devices, the internal system clock is typically manually set to within a minute or two of the actual time and is rarely reset at regular intervals. Many local clocks are battery-backed devices that use room temperature clock oscillators that can drift as much as several seconds each day. NTP adjusts the time of the devices so that they synchronize within a millisecond (ms) on LANs and up to a few tens of milliseconds on WANs relative to Coordinated Universal Time (UTC).

The NTP client on the switch supports only unicast client mode. In this mode, the NTP client sends NTP time requests to other remote time servers in an asynchronous fashion. The NTP client collects four samples of time from each remote time server. A clock selection algorithm determines the best server among the selected samples based on stratum, delay, dispersion and the last updated time of the remote server. The real time clock (RTC) is adjusted to the selected sample from the chosen server.

NTP Terms

A peer is a device that runs NTP software. However, this implementation of NTP refers to peers as remote time servers that provide time information to other time servers on the network and to the local NTP client. An NTP client refers to the local network device, the switch, that accepts time information from other remote time servers.

NTP System Implementation Model

NTP is based on a hierarchical model that consists of a local NTP client that runs on the switch and on remote time servers. The NTP client requests and receives time information from one or more remote time servers. The local NTP client reviews the time information from all available time servers and synchronizes its internal clock to the time server whose time is most accurate. The NTP client does not forward time information to other devices that run NTP.

Two types of time servers exist in the NTP model: primary time servers and secondary time servers. A primary time server is directly synchronized to a primary reference source, usually a wire or radio clock that is synchronized to a radio station that provides a standard time service. The primary time server is the authoritative time source in the hierarchy, meaning that it is the one true time source to which the other NTP devices in the subnet synchronize their internal clocks.

A secondary time server uses a primary time server or one or more secondary time servers to synchronize its time, forming a synchronization subnet. A synchronization subnet is a self-organizing, hierarchical master-backup configuration with the primary servers at the root and secondary servers of decreasing accuracy at successive levels.

The following figure shows NTP time servers forming a synchronization subnet.




Figure 189: NTP time servers forming a synchronization subnet

In the NTP model, the synchronization subnet reconfigures in a hierarchical primary-secondary configuration to produce accurate and reliable time, even if one or more primary time servers or the path between them fails. This feature applies in a case in which all the primary servers on a partitioned subnet fail, but one or more backup primary servers continue to operate. If all of the primary time servers in the subnet fail, the remaining secondary servers synchronize among themselves.

Time Distribution Within a Subnet

NTP distributes time through a hierarchy of primary and secondary servers, with each server adopting a stratum, see [NTP System Implementation Model](#) on page 2725. A stratum defines how many NTP hops away a particular secondary time server is from

an authoritative time source (primary time server) in the synchronization subnet. A stratum 1 time server is located at the top of the hierarchy and is directly attached to an external time source, typically a wire or radio clock; a stratum 2 time server receives its time through NTP from a stratum 1 time server; a stratum 3 time server receives its time through NTP from a stratum 2 time server, and so forth.

Each NTP client in the synchronization subnet chooses as its time source the server with the lowest stratum number with which it is configured to communicate through NTP. This strategy effectively builds a self-organizing tree of NTP speakers. The number of strata is limited to 15 to avoid long synchronization loops.

NTP avoids synchronizing to a remote time server with inaccurate time. NTP never synchronizes to a remote time server that is not itself synchronized. NTP compares the times reported by several remote time servers.

Synchronization

Unlike other time synchronization protocols, NTP does not attempt to synchronize the internal clocks of the remote time servers to each other. Rather, NTP synchronizes the clocks to universal standard time, using the best available time source and transmission paths to that time source.

Use the **show NTP statistics** command to verify the NTP synchronization status. NTP uses the following criteria to determine the best available time server:

- The time server with the lowest stratum.
- The time server closest in proximity to the primary time server (reduces network delays).
- The time server that offers the highest claimed precision.

NTP accesses several (at least three) servers at the lower stratum level because it can apply an agreement algorithm to detect a problem on the time source.

NTP Modes of Operation

NTP uses unicast client mode to enable time servers and NTP clients to communicate in the synchronization subnet. The switch supports only unicast client mode.

After you configure a set of remote time servers (peers), NTP creates a list that includes each time server IP address. If you configure the server address using a fully qualified domain name (FQDN), the switch resolves the FQDN to the IP address. The NTP client uses this list to determine the remote time servers to query for time information.



Note

Using FQDNs requires a DNS configuration for address resolution. If you add a new DNS server or change the operational state of a DNS server, NTP also restarts. For more information about how to configure DNS on the switch, see [Domain Name Service](#) on page 908.

After the NTP client queries the remote time servers, the servers respond with various timestamps, along with information about their clocks, such as stratum, precision, and time reference. The NTP client reviews the list of responses from all available servers and chooses one as the best available time source from which to synchronize its internal clock.

The following figure shows how NTP time servers operate in unicast mode.




Figure 190: NTP Time Servers Operating in Unicast Client Mode

NTP Master Mode

The switch can operate as both an NTPv4 client and an NTPv4 server. You can configure the NTPv4 server by enabling the master mode. When the master mode is configured, peers can synchronize themselves with the local clock when the NTPv4 server loses synchronization or if an external NTPv4 source is not reachable. For information about configuring NTPv4 server master mode, see [Configure NTP Master Mode](#) on page 2733 and [Configure NTP Globally](#) on page 2736.

NTP Restrict

The switch offers the restrict capability on the NTPv4 server. When the NTPv4 server master mode is disabled, the restrict capabilities are disabled by default. All addresses or networks except for those configured as servers are ignored. For addresses configured as servers, traffic is allowed but there are some default restriction values.

When the NTPv4 server master mode is enabled, there are no restrict rules configured, which means all connections are allowed or there are one or multiple rules configured and only those addresses or networks with the configured rules are allowed. For more information about creating NTP restrict entries, see [Create NTP Restrict Entries](#) on page 2735.

NTP Authentication

You can authenticate time synchronization to ensure that the local time server obtains its time services only from known sources. NTP authentication adds a level of security to your NTP configuration. By default, network time synchronization is not authenticated.

If you select authentication, the switch uses the Message Digest 5 (MD5) or the Secure Hash Algorithm 1 (SHA1) algorithm to produce a message digest of the key. The message digest is created using the key and the message, but the key itself is not sent. Depending on which algorithm you select, the MD5 or SHA1 algorithm verifies the integrity of the communication, authenticates the origin, and checks for timeliness.

To authenticate the message, the client authentication key must match that of the time server. Therefore, you must securely distribute the authentication key in advance (the client administrator must obtain the key from the server administrator and configure it on the client).

While you can configure a server with many keys (identified by many key IDs), it is possible to declare only a subset of these as trusted. The time server uses this feature to share keys with a client that requires authenticated time and that trusts the server, but that is not trusted by the time server.

NTP Configuration Using CLI

This section describes how to configure the Network Time Protocol (NTP) using Command Line Interface (CLI).

Before you configure NTPv4, you must create a Segmented Management Instance and configure routing for that instance.



Important

NTP server MD5 authentication or SHA1 authentication does not support passwords (keys) that start with a special character or that contain a space between characters.

Enable NTP Globally

Enable NTP globally. Default values are in effect for most parameters. You can customize NTP by modifying parameters.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. (Optional) Configure the NTP interval time (between successive NTP updates):
`ntp interval <4-17>`
3. Enable NTP globally:
`ntp`

4. Confirm the global configuration:

```
show ntp
```

Example

Specify the time interval between NTP updates, and then enable NTP globally.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ntp interval 10
Switch:1(config)#ntp
```

Confirm the configuration:

```
Switch:1(config)#show ntp

=====
                                NTP Master
=====
Version   Enabled   Stratum
-----
4         False    10

=====
                                NTP Client
=====
Version   Enabled   Interval   Last Update Time                Synchronized To
-----
4         True     10         Thu Jul 18 08:32:59 2019 EDT    192.0.2.0 (Stratum:2)
```

Variable Definitions

The following table defines parameters for the **ntp** command.

Variable	Value
<i>authentication-key</i> <1-65534> <i>WORD</i> <0-20>	Creates an authentication key for MD5 or SHA1 authentication. The default configuration is to delete the authentication key. NTP server MD5 or SHA1 authentication does not support passwords (keys) that start with a special character or contain a space between characters. <i>WORD</i> <0-20> specifies the secret key.
<i>type</i> <md5 sha1>	Specifies the type of authentication as MD5 or SHA1. The default is MD5 authentication.
<i>interval</i> <4-17>	Specifies the time interval between successive NTP updates as a power of 2, in seconds. The default for NTPv4 is 2 to the power of 8 seconds.

Add an NTP Server

Add a remote NTP server to the configuration by specifying its IP address or fully qualified domain name (FQDN). NTP adds the IP address to a list of servers, which the local NTP client uses to query remote time servers for time information. The list of qualified servers called to is a peer list.

Before You Begin

If you specify an FQDN for the server, you must configure DNS for address resolution.

About This Task

For NTPv4, this procedure adds the NTP server information to the switch that is acting as an NTP client. You can configure a maximum of 20 NTP servers.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Add an NTP server:

```
ntp server WORD<1-110>
```
3. Configure additional options for the NTP server:

```
ntp server WORD<1-110> [auth-enable] [authentication-key <0-65534>]
[enable]
```

The NTP server is automatically enabled by default.

4. Confirm the configuration:

```
show ntp server [name]
```

Examples

Configure a server using an IPv4 address:

```
Switch:>enable
Switch:1configure terminal
Switch:1(config)#ntp server 192.0.2.187 enable
Switch:1(config)#show ntp server
```

```
=====
                        NTP Server
=====
Server Ip                Enabled Auth    Key Id    Auth Type
-----
192.0.2.187              true      false    0         N/A
```

Configure a server using an FQDN:

```
Switch:>enable
Switch:1configure terminal
Switch:1(config)#ntp server host.example.com enable
Switch:1(config)#show ntp server
```

```
=====
                        NTP Server
=====
Server Ip                Enabled Auth    Key Id    Auth Type
-----
host.example.com         true      false    0         N/A
```

Variable Definitions

The following table defines parameters for the **ntp server** command.

Variable	Value
<i>auth-enable</i>	Activates MD5 or SHA1 authentication on this NTP server. Without this option, the NTP server will not have any authentication by default.
<i>authentication-key <0-65534></i>	Specifies the key ID value used to generate the MD5 or SHA1 digest for the NTP server. The default authentication key is 0, which indicates disabled authentication.
<i>enable</i>	Activates the NTP server.
<i>WORD<1-110></i>	Specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN) of the NTP server. If you use an FQDN, you must also configure the switch to use DNS.

The following table defines parameters for the **show ntp server** command.

Variable	Value
<i>name</i>	Displays all NTP server names.

View NTP Statistics

About This Task

Use the **show ntp statistics** command to view the output for each NTP stratum regardless of whether authentication is enabled.

If two host names resolve to the same IP address, only the first one displays in the output for the **show ntp statistics** command.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View NTP statistics:
show ntp statistics

Examples

```
Switch:1#show ntp statistics

      NTP Server : 192.0.2.187
-----
      Stratum : 16
      Version : NTPv4
      Broadcast : No
      Auth Enabled : Disabled
      Auth Status : Not-Auth
      Sync Status : Rejected
      Reachability : Unreachable
      Root Delay : 0.000
      Root Disp : 0.000
      Delay : 0.000
```

```

Dispersion : 15937.500
Offset : 0.000
Precision : -23
Jitter : 0.000
Last Event : Mobilize

NTP Server : 192.0.2.201
-----
Stratum : 4
Version : NTPv4
Broadcast : No
Auth Enabled : Enabled
Auth Status : Ok
Sync Status : Candidate
Reachability : Reachable
Root Delay : 18.448
Root Disp : 128.677
Delay : 18.448
Dispersion : 0.366
Offset : 0.202
Precision : -24
Jitter : 1.041
Last Event : Popcorn

```

Configure Authentication Keys

About This Task

Configure up to 10 NTP authentication keys to use MD5 or SHA1 authentication.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create an authentication key:

```
ntp authentication-key <1-65534> type <md5|sha1>
```
3. Enter the secret key:

```
WORD <1-20>
```
4. Re-enter the secret key:

```
WORD <1-20>
```
5. Enable MD5 or SHA1 authentication for the server:

```
ntp server WORD<1-110> auth-enable
```
6. Assign an authentication key to the server:

```
ntp server WORD<1-110> authentication-key <0-65534>
```



Note

If you must disable authentication on the server, you must also disable authentication on the switch for example: `no ntp server WORD<1-110> auth-enable`

7. Confirm the configuration:

```
show ntp key
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ntp authentication-key 5 type md5
Enter the NTP secret key: *****
Reenter the NTP secret key: *****
Switch:1(config)#ntp server 192.0.2.187 auth-enable
Switch:1(config)#ntp server 192.0.2.187 authentication-key 5
Switch:1(config)#show ntp key
=====
                        NTP Key
=====
Key_Id      Type
-----
5           MD5
10          SHA1
20          MD5
30          SHA1
100         MD5
```

Variable Definitions

The following table defines parameters for the **ntp** and **ntp server** commands.

Variable	Value
<i>auth-enable</i>	Activates MD5 or SHA1 authentication on this NTP server. The default is no authentication.
<i>authentication-key <0-65534></i>	Specifies the key ID value used to generate the MD5 or SHA1 digest for the NTPv4 server. The default value is 0, which indicates disabled authentication.
<i>type <md5 sha1></i>	Specifies the type of authentication, whether MD5 or SHA1. The default is MD5 authentication.
<i>WORD<1-110></i>	Specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN) of the NTP server. If you use an FQDN, you must also configure the switch to use DNS.

Configure NTP Master Mode

About This Task

Perform the following procedure to set the switch to act as the Network Time Protocol (NTP) server, which means it will run in the master mode. The default value is disabled. You can also enable NTP master mode for a specific stratum.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. (Optional) Enable NTP master mode:

```
ntp master [ <1-16>]
```
3. Verify the configuration:

```
show ntp
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ntp master
Switch:1(config)#show ntp
```

NTP Master				
Version	Enabled	Stratum		
4	True	11		
NTP Client				
Version	Enabled	Interval	Last Update Time	Synchronized To
4	False	60		

Variable Definitions

The following table defines parameters for the **ntp master** command:

Variable	Value
<1-16>	Specifies an optional stratum value. The default value is 10. Note: If there is a better stratum available, it is preferred than what is configured.

Create NTP Restrict Entries

About This Task

Perform the following procedure to configure the NTP restrict capability for a specific IPv4 or IPv6 address, which means the switch permits NTP traffic flow from the specified IP addresses only. By default, the NTP restrict capability is disabled.



Note

- You can configure a maximum of 128 NTP Restrict entries (IPv4 or IPv6 addresses).
- 0.0.0.0/0 or ::/0 NTP restrict entries are equivalent to no NTP restrict rules configured.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure a specific NTP restrict IP address:
`ntp restrict WORD<0-255>`
3. Verify the restricted IP address:
`show ntp restrict`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ntp restrict 192.0.0.0
Switch:1(config)#show ntp restrict
=====
                        NTP Restrict Information
=====
TYPE      ADDRESS                                MASK/PREFIX  LEN
-----
IPv4      192.0.0.0                                23
=====
```

Variable Definitions

The following table defines parameters for the **ntp restrict** command:

Variable	Value
<i>WORD<0-255></i>	Specifies the IPv4 or IPv6 address. Note: You can configure a maximum of 128 IPv4 and IPv6 addresses in the NTP restrict list.

NTP Configuration Using EDM

This section describes how to configure the Network Time Protocol (NTP) using EDM.

Before you configure NTPv4, you must create a Segmented Management Instance and configure routing for that instance.

**Important**

NTP server MD5 authentication or SHA1 authentication does not support passwords (keys) that start with a special character or that contain a space between characters.

Configure NTP Globally

You can use the following procedure to globally enable NTP.

You can also enable master mode on the NTP server configured on the switch. When the master mode is configured, peers can synchronize themselves with the local clock when the NTP server loses synchronization or if an external NTP source is not reachable.

Before You Begin

You must globally disable NTP before you change the version.

Procedure

1. In the navigation pane, expand **Configuration > Edit > NTP**.
2. Select **General**.
3. Select the **Globals** tab.
4. Select **Enable**.
5. Enter a time interval.
6. (Optional) To configure the switch as an NTP server, in the **NTP Master** section, select **Enable**.
7. (Optional) In the **Stratum** field, enter a value.
8. Select **Apply**.

Globals *Field Descriptions*

Use data in the following table to use the **Globals** tab.

NTP Client section:

Name	Description
Enable	Enables the client NTP server. By default, NTP is disabled.
Interval	Specifies the time interval between successive NTP updates as a power of 2 in seconds. The default value for NTPv4 is 2 to the power of 8 seconds.

NTP Master section:

Name	Description
Enable	Enables master mode for the configured NTP server. The default value is disabled.
Stratum	Specifies the stratum for the master NTP server. The default value is 10. Note: If a better stratum is available, it is preferred over what is configured.

Add an NTPv4 Server

Add a remote NTP server to the configuration by first specifying its IP address or fully qualified domain name (FQDN). NTP adds the IP address to a list of servers, which the local NTP client uses to query remote time servers for time information. The list of qualified servers called to is a peer list.

Before You Begin

If you specify an FQDN for the server, you must configure DNS for address resolution.

About This Task

For NTPv4, this procedure adds the NTP server information to the switch that is acting as an NTP client. You can configure a maximum of 20 NTP servers.

Procedure

1. In the navigation pane, expand **Configuration > Edit > NTP**.
2. Select **NTPv4**.
3. Select the **Server** tab.
4. Select **Insert**.
5. Specify if the IP address is an IPv4 or IPv6 address or select **dns** to use a Fully Qualified Domain Name (FQDN) instead.
6. Specify the IP address or FQDN of the NTP server.
7. Select **Insert**.

Server Field Descriptions

Use the data in the following table to use the **Server** tab.

Name	Description
ServerAddressType	Specifies the address type as either IPv4 or IPv6, or DNS for a Fully Qualified Domain Name (FQDN).
ServerAddress	Specifies the IP address or FQDN of the remote NTP server.

Name	Description
Enable	Activates or disables the remote NTP server. The default is enabled.
Authentication	Activates or disables MD5 or SHA1 authentication on this NTP server. MD5 or SHA1 produces a message digest of the key. MD5 or SHA1 verifies the integrity of the communication, authenticates the origin, and checks for timeliness. The default is no authentication.
KeyId	Specifies the key ID used to generate the MD5 or SHA1 digest for this NTP server. The default is 0, which indicates that authentication is disabled.
Stratum	Shows the stratum of the server.
Version	Shows the NTP version of the server.
Broadcast	Shows if broadcast is enabled or disabled
AuthEnabled	Shows if authentication is enabled or disabled
AuthStatus	Shows the authentication status.
Synchronized	Shows the status of synchronization with the server.
Reachable	Shows the NTP reachability status of the server.
RootDelay	Shows the root delay of the server.
RootDisp	Shows the root dispersion of the server.
ServerDelay	Shows the delay of the server.
Dispersion	Shows the dispersion of the server.
Offset	Shows the offset of the server.
Precision	Shows the NTP precision of the server in seconds.
Jitter	Shows the jitter of the server
LastEvent	Shows the last event of the server.

Configure Authentication Keys for NTPv4

Assign an NTP key to use MD5 or SHA1 authentication on the server.

Procedure

1. In the navigation pane, expand **Configuration > Edit > NTP**.
2. Select **NTPv4**.
3. Select the **Key** tab.
4. Select **Insert**.
5. Complete the fields.
6. Select **Insert**.

Key Field Descriptions

Use the data in the following table to use the **Key** tab.

Name	Description
KeyId	Specifies the key ID that generates the MD5 or SHA1 digest.
KeyType	Specifies the type of authentication as MD5 or SHA1. The default is MD5 authentication.

Create NTPv4 Restrict Entries

Before You Begin

You must enable NTPv4 server master mode. For more information, see [Configure NTP Globally](#) on page 2736.

About This Task

Perform the following procedure to configure the NTP restrict capability for a specific IPv4 or IPv6 address(es), which means the switch permits traffic flow from the specified IP address only. By default the NTP restrict capability is disabled.



Note

- You can configure a maximum of 128 NTP Restrict entries (IPv4 or IPv6 addresses).
- 0.0.0.0/0 or ::/0 NTP restrict entries are equivalent to no NTP restrict rules configured.

Procedure

1. In the navigation pane, expand **Configuration > Edit > NTP**.
2. Select **Restrict**.
3. Select the **Restrict Info** tab.
4. Select **Insert**.
5. In **RowIndex**, type a value.
6. Select the IP address type.
7. Type the IPv4 or IPv6 address.
8. Type the restrict mask value.
9. Select **Insert**.

Restrict Info Field Descriptions

Use data in the following table to use the **Restrict Info** tab.

Name	Description
RowIndex	Specifies the NTP Restrict entry.
AddressType	Specifies the NTP Restrict address type.

Name	Description
RestrictAddress	Specifies the NTP address to be restricted.
RestrictMask	Specifies the prefix length of the IPv4 or IPv6 NTP address to be restricted.



OpenAPI

[Configure OpenAPI Local Port Access](#) on page 2742
[Display OpenAPI Local Port Access Information](#) on page 2743
[Display OpenAPI Log File Information](#) on page 2743
[Configure OpenAPI Authentication on a Local Port](#) on page 2744
[Display the OpenAPI Authentication State](#) on page 2745
[OpenAPI Troubleshooting](#) on page 2746

Table 213: OpenAPI product support

Feature	Product	Release introduced
OpenAPI local access	4220 Series	Fabric Engine 9.3
	5320 Series	Fabric Engine 9.3
	5420 Series	Fabric Engine 9.3
	5520 Series	Fabric Engine 9.3
	5720 Series	Fabric Engine 9.3
	7720 Series	Fabric Engine 9.3
	7520 Series	Fabric Engine 9.3
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.3
	VSP 7400 Series	VOSS 9.3

You can use the Command Line Interface (CLI) to use OpenAPI through local port access for third-party management platforms and software tools.

Operational Considerations

Consider the following when you use OpenAPI for local port access:

- This feature requires a Premium, Pilot or Extreme Platform ONE Networking license.
- You must have an API authentication token to perform requests on the local port.
- You must have read-write-all (rwa) user access level to authenticate.

Configure OpenAPI Local Port Access

About This Task

Perform this procedure to configure local port access for OpenAPI.



Important

If you change the TCP port or TLS configuration, you must disable and then reenable the local port state.

Procedure

1. Enter Application Configuration mode:
`enable`

`configure terminal`

`application`
2. Enable local port access:
`openapi local-mgmt enable`
3. (Optional) Configure the local port number:
`openapi local-mgmt tcp-port <1-9443>`
4. (Optional) Enable secure communication for local port access:
`openapi local-mgmt tls enable`
5. (Optional) Configure the minimum TLS version:
`openapi local-mgmt minimum-tls {1.2 | 1.3}`
6. (Optional) Configure the OpenAPI server local port for the authentication token:
`openapi local-mgmt ttl <60-86400>`
7. Save the configuration file:
`save config`

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CTRL/Z.
Switch:1(config)#application
Switch:1(config-app)#openapi local-mgmt enable
Switch:1(config-app)#openapi local-mgmt tcp-port 1025
Switch:1(config-app)#openapi local-mgmt enable
Switch:1(config-app)#save config
```

Variable Definitions

The following table defines parameters for the **openapi local-mgmt** command.

Variable	Value
<i>enable</i>	Enables local port access.
<i>minimum-tls {1.2 1.3}</i>	Specifies the OpenAPI server local port minimum TLS version. The default is 1.2.
<i>tcp-port <1-9443></i>	Specifies the local port number. The default is 9443.
<i>ttl <60-86400></i>	Specifies the time-to-live value for the authentication token. The default is 86400.
<i>tls enable</i>	Enables the secure communication for local port access. The default is enabled.

Display OpenAPI Local Port Access Information

About This Task

Display OpenAPI local port access information on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the local port access information:

```
show application openapi
```

Example

```
Switch:1>show application openapi
=====
OpenAPI Info
=====
Server State           : Enabled
Server State Local Port : Disabled
Local Port State       : Disabled
Local TCP Port         : 9443
TLS                    : Enabled
TLS version            : 1.2
TTL(in seconds)        : 60
Server Version         : 0.1.10.3
Specification Version   :
```

Display OpenAPI Log File Information

About This Task

Use this procedure to display OpenAPI server log file information such as the type of request, the IP address of the requester, and the URL string.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display OpenAPI log file information:

```
show application openapi log [action WORD<1-7>] [reverse] [source WORD<1-64>] [url <WORD<1-128>]
```
3. Display the last few lines of the OpenAPI log:

```
show application openapi log |tail
```

Example

The following example output displays last few lines of the OpenAPI log:

```
Switch:1(config-app)#show application openapi log | tail
2025-10-07 06:05:22,176:INFO:voss-openapi:__main__.py:main:81:: [1|1] Starting openapi webserver version 0.1.11.2
2025-10-07 07:25:29,803:INFO:voss-openapi:__main__.py:main:39:: [1|1] Starting voss-openapi application
2025-10-07 07:25:32,996:INFO:voss-openapi:__main__.py:main:81:: [1|1] Starting openapi webserver version 0.1.11.2
2025-10-08 06:22:28,263:INFO:voss-openapi:__main__.py:main:39:: [1|1] Starting voss-openapi application
2025-10-08 06:22:30,943:INFO:voss-openapi:__main__.py:main:81:: [1|1] Starting openapi webserver version 0.1.11.2
2025-10-08 12:40:28,025:INFO:voss-openapi:__main__.py:main:39:: [1|1] Starting voss-openapi application
2025-10-08 12:40:31,011:INFO:voss-openapi:__main__.py:main:81:: [1|1] Starting openapi webserver version 0.1.11.2
2025-10-09 06:35:06,360:INFO:voss-openapi:__main__.py:main:39:: [1|1] Starting voss-openapi application
2025-10-09 06:35:09,039:INFO:voss-openapi:__main__.py:main:81:: [1|1] Starting openapi webserver version 0.1.11.2
```

Variable Definitions

The following table defines parameters for the **show application openapi log** command.

Variable	Value
<i>action WORD<1-7></i>	Filters the OpenAPI log based on the following action types: <ul style="list-style-type: none"> • DELETE • GET • PATCH • POST • PUT
<i>reverse</i>	Displays the OpenAPI log in chronological order.
<i>source WORD<1-64></i>	Filters the OpenAPI log based on the application that initiated the request.
<i>url WORD<1-128></i>	Filters the OpenAPI log based on the endpoint URL.
<i>tail</i>	Limits the output of a command to the last few lines. For more information, see GREP with CLI show command on page 191.

Configure OpenAPI Authentication on a Local Port

About This Task

Perform the following procedure to configure OpenAPI authentication on a local port.

This command is used for debugging purposes.

**Important**

You must disable and then reenable the local port to apply changes.

Procedure

1. Enter Application Configuration mode:
enable

configure terminal

application
2. Enable the debug mode:
dbg enable
3. Configure authentication on a local port:
openapi local-mgmt auth enable

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#application
Switch:1(config-app)#dbg enable
Entering privilege command mode
Switch:1(config-app)#no openapi local-mgmt auth enable
Warning : openapi local-port must be restarted to disable token authentication
```

Display the OpenAPI Authentication State

About This Task

Perform the following procedure to display the OpenAPI authentication state.

This command is used for debugging purposes.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Enable the debug mode:
dbg enable
3. Display the OpenAPI authentication state:
show application openapi authentication

Example

```
Switch:1>enable
Switch:1#dbg enable
Entering privilege command mode
Switch:1#show application openapi authentication
=====
OpenAPI Auth
```

```
=====
OpenAPI authentication state      : Disabled
=====
```

OpenAPI Troubleshooting

Use the information in this section to troubleshoot OpenAPI.

OpenAPI Local Port Cannot be Enabled

```
Switch:1(config-app)#openapi local-mgmt enable
Error: Operation not allowed without a proper license file installed
```

Verify there is a valid license installed. A Premier, Pilot or Extreme Platform ONE license is required to enable use of OpenAPI.

curl Receives Error Code 401

```
curl -k -X GET
-H "X-Auth-Token:.eJxTMjcyNzczlIVSZslmaWbGZoZmybrGhsZGxuZGqcYGUGyoBADhWgoI.aIpWPA._q-4KvA-
s5A7AHCihwN0JaQr8rg" https://10.101.21.12:9443/rest/openapi/v0/configuration/vlan/1
{
  "errorCode": 401,
  "errorMessage": "Authentication was enforced please add X-Auth-Token to request
header",
  "resource": "/rest/openapi/v0/configuration/vlan/1"
}
```

The authentication token is available for 24 hours if the server is not restarted in the meantime.

In this case, the authentication token is not current one. Use curl to obtain another authentication token.

```
curl -X 'POST' 'http://10.101.24.120:9443/rest/openapi/v0/operation/auth-
token/:generate' -H 'accept: application/json' -H 'Content-Type: application/json'
-d '{"username": "rwa", "password":"rwa"}
```

curl Receives "Connection refused"

```
curl -k -X GET
-H "X-Auth-Token:.eJxTMjcyNzczlIVSZslmaWbGZoZmybrGhsZGxuZGqcYGUGyoBADhWgoI.aIpWPA._q-4KvA-
s5A7AHCihwN0JaQr8rg" https://10.101.21.12:9443/rest/openapi/v0/configuration/vlan/1
curl: (7) Failed connect to 10.101.21.12:9443; Connection refused
```

Check the local port state:

```
Switch:1(config-app)#show application openapi
=====
OpenAPI Info
=====
Server State           : Disabled
Server State Local Port : Enabled
Local Port State       : Enabled
Local TCP Port         : 2
TLS                    : Disabled
```



```

Server Version          : 0.1.10.3
Specification Version   :

```

If the local port state is enabled, use the openapi log to check if the OpenAPI process has started:

```

Switch:1(config-app)#show application openapi log | tail
2025-07-30 17:27:46,089:INFO:voss-openapi:__main__.py:main:39:: [1|1] Starting voss-
openapi application
2025-07-30 17:27:48,662:INFO:voss-openapi:__main__.py:main:81:: [1|1] Starting openapi
webserver version 0.1.10.3
2025-07-30 17:27:59,600:INFO:voss-openapi:log_requests.py:log_response:30:: [1|
1] 127.0.0.1(iqagent) POST http://127.0.0.1:7777/rest/openapi/v0/operation/auth-
token/generate {'username': 'hivemanager', 'password': '*****'} 200 OK 0.0184 seconds
2025-07-30 17:27:59,616:INFO:voss-openapi:log_requests.py:log_response:30:: [1|
1] 127.0.0.1(iqagent) POST http://127.0.0.1:7777/rest/openapi/v0/operation/auth-
token/disable None 204 NO CONTENT 0.0091 seconds
2025-07-30 17:28:23,549:INFO:voss-openapi:log_requests.py:log_response:30:: [1|1]
127.0.0.1(unknown) GET http://openapi.sm02235q-40020/rest/openapi/v0/state/system None
200 OK 0.0613 seconds

```

Check if the OpenAPI process is running. openapi.sh starts the process when the iqagent is running:

```

Switch:1(debug)#exec bash
bash-5.1# ps -ef | grep openapi
root      247514      14018    5 17:27 ttyS0      00:00:19 python /opt/appfs/bin/voss_openapi.pyz
root      247515      247514    0 17:27 pts/6        00:00:00 logger -p local0 err -t openapi.sh

```

If there is no cloud connectivity, openapi-local.sh starts the OpenAPI server:

```

Switch:1(debug)#exec bash
bash-5.1# ps -ef | grep openapi
root      252766      14018    0 17:35 ttyS0      00:00:00 /bin/bash /opt/appfs/bin/openapi-
local.sh
root      252767      252766    0 17:35 pts/6        00:00:00 logger -p local0 err -t openapi-
local.sh
root      252771      252766    91 17:35 ttyS0      00:00:08 python /opt/appfs/bin/
voss_openapi.pyz --local-port=2 --local-port-tls=0 --local-auth=1

```

If the OpenAPI process is running, use netstat to check if TCP is listening on the configured port. Confirm on which port the OpenAPI server was opened and then check if there is an opened socket on that particular port:

```

bash-5.1# tail /var/log/messages
Jul 30 17:43:40 aries64 local0.err openapi-local.sh:    WARNING: This is a development
server. Do not use it in a production deployment.
Jul 30 17:43:40 aries64 local0.err openapi-local.sh:    * Running on https://127.0.0.1:9443
Jul 30 17:43:40 aries64 local0.err openapi-local.sh:    * Running on https://
10.101.21.12:9443 (Press CTRL+C to quit)

bash-5.1# netstat -ntlp | grep 9443
tcp        0      0 0.0.0.0:9443          0.0.0.0:*             LISTEN      255857/
python

```



OSPF

- [OSPF Fundamentals](#) on page 2750
- [OSPF configuration using CLI](#) on page 2783
- [OSPFv3 Configuration using CLI](#) on page 2827
- [OSPF Configuration Using EDM](#) on page 2857
- [OSPFv3 Configuration using EDM](#) on page 2896

Table 214: OSPF product support

Feature	Product	Release introduced
Open Shortest Path First (OSPF)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 214: OSPF product support (continued)

Feature	Product	Release introduced
Secure hash algorithm 1 (SHA-1) and SHA-2	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 215: OSPFv3 product support

Feature	Product	Release introduced
OSPFv3	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 215: OSPFv3 product support (continued)

Feature	Product	Release introduced
OSPFv3 support for CLIP interfaces	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.10.1
	5420 Series	Fabric Engine 8.10.1
	5520 Series	Fabric Engine 8.10.1
	5720 Series	Fabric Engine 8.10.1
	7520 Series	Fabric Engine 8.10.1
	7720 Series	Fabric Engine 8.10.1
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.10.1
	VSP 7400 Series	VOSS 8.10.1

OSPF Fundamentals

This section introduces the fundamental concepts of Open Shortest Path First (OSPF), including its operating model, key benefits, and protocol variants.

OSPF Overview

Open Shortest Path First (OSPF) is a link-state Interior Gateway Protocol (IGP) that distributes routing information within a single autonomous system. OSPF routers build and synchronize a link-state database (LSDB) and use the Shortest Path First (SPF) algorithm to compute optimal routes.

OSPF converges quickly after link or node failures, minimizes routing traffic compared to distance-vector protocols, and supports IP subnetting, Type of Service (TOS)-based routing, external route tagging, and equal-cost multipath (ECMP) forwarding.

The OSPF protocol includes two versions: OSPFv2 for IPv4 networks and OSPFv3 for IPv6 networks.

In large networks, OSPF offers the following benefits:

- Fast convergence—After network topology changes, OSPF recalculates routes quickly.
- Minimal routing protocol traffic—Unlike distance vector routing protocols, such as Routing Information Protocol (RIP), OSPF generates a minimum of routing protocol traffic.
- Load sharing—OSPF provides support for Equal Cost Multipath (ECMP) routing. If several equal-cost routes to a destination exist, ECMP distributes traffic equally among them.
- Type of service routing—OSPF can calculate separate routes for each IP TOS.

In an OSPF network, each router maintains a link-state database that describes the topology of the AS. The database contains the local state for each router in the AS, including its usable interfaces and reachable neighbors. Each router periodically checks for changes in its local state and shares detected changes by flooding link-state advertisements (LSA) throughout the AS. Routers synchronize their topological databases based on the sharing of information from LSAs.

From the topological database, each router constructs a shortest-path tree, with itself as the root. The shortest-path tree provides the optimal route to each destination in the AS. Routing information from outside the system displays the AS on the tree as leaves. OSPF routes IP traffic based on the destination IP address, subnet mask, and IP TOS.

Autonomous System and Areas

The AS subdivides into areas that group contiguous networks, routers that connect to these networks, and attached hosts. Each area has a topological database, which is invisible from outside the area. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic compared to treating the entire AS like a single link-state domain.

You can attach a router to more than one area. When you perform this action, you can maintain a separate topological database for each connected area. Two routers within the same area maintain an identical topological database for that area. Each area uses a unique area ID and the area ID 0.0.0.0 is reserved for the backbone area.

The router routes packets in the AS based on their source and destination addresses. If the source and destination of a packet reside in the same area, the router uses intra-area routing. If the source and destination of a packet reside in different areas, the router uses inter-area routing. Intra-area routing protects the area from bad routing information because it does not use routing information obtained from outside the area. Inter-area routing must pass through the backbone area. For more information about the backbone area, see [Backbone Area](#) on page 2752.

In large networks with many routers and networks, the link-state database (LSDB) and routing table can become excessively large. Large route tables and LSDBs consume memory. The processing of link-state advertisements results in additional CPU cycles to make forwarding decisions. To reduce these undesired effects, you can divide an OSPF network into subdomains called areas.

An area comprises a number of OSPF routers that have the same area identification (ID).

By dividing a network into multiple areas, the router maintains a separate LSDB, which consists of router LSAs and network LSAs, for each area. Each router within an area maintains an LSDB only for the area to which it belongs. Area router LSAs and network LSAs do not flood beyond the area borders.

The impact of a topology change is localized to the area in which it occurs. The only exception is for the area border router (ABR), which must maintain an LSDB for each area to which they belong. The area border routers advertise changes in topology to the remainder of the network by advertising summary LSAs.

A 32-bit area ID, expressed in IP address format (x.x.x.x), identifies areas. Area 0 is the backbone area and distributes routing information to all other areas.

If you use multiple areas, they must all attach to the backbone through an ABR, which connects area 0.0.0.0 to the nonbackbone areas. If you cannot physically and directly connect an area through an ABR to area 0, you must configure a virtual link to logically connect the area to the backbone area.

Backbone Area

The backbone area consists of the following network types:

- networks and attached routers that do not exist in other areas
- routers that belong to multiple areas

The backbone is usually contiguous but you can create a noncontiguous area by configuring virtual links.

You can configure virtual links between two backbone routers that have an interface to a nonbackbone area. Virtual links belong to the backbone and use intra-area routing only.

The backbone distributes routing information between areas. The topology of the backbone area is invisible to other areas, while it knows nothing of the topology of those areas.

In inter-area routing, a packet travels along three contiguous paths in a point-to-multipoint configuration:

- an intra-area path from the source to an ABR
- a backbone path between the source and destination areas
- another intra-area path to the destination

The OSPF routing algorithm finds the set of paths that has the smallest cost. The topology of the backbone dictates the backbone paths used between areas. OSPF selects inter-area paths by examining the routing table summaries for each connected ABR. The router cannot learn OSPF routes through an ABR unless it connects to the backbone or through a virtual link.

Stub Area

Configure a stub area at the edge of the OSPF routing domain. A stub area has only one ABR. A stub area does not receive LSAs for routes outside its area, which reduces the size of its link-state database. A packet destined outside the stub area is routed to the ABR, which examines it before forwarding the packet to the destination. The network behind a passive interface is treated as a stub area and does not form adjacencies. The network is advertised into the OSPF area as an internal route.

Not so Stubby Area

A not-so-stubby area (NSSA) prevents the flooding of external LSAs into the area by replacing them with a default route. An NSSA can import small stub (non-OSPF) routing domains into OSPF. Like stub areas, NSSAs are at the edge of an OSPF routing

domain. Non-OSPF routing domains attach to the NSSAs to form NSSA transit areas. Accessing the addressing scheme of small stub domains permits the NSSA border router to also perform manual aggregation.

In an OSPF NSSA, the NSSA N/P bit notifies the ABR which external routes to advertise to other areas. If the NSSA N/P bit is set (the value is 1), the ABR exports the external route. This configuration is the default. When the NSSA N/P bit is not set (the value is 0), the ABR drops the external route. You can create a route policy to manipulate the N/P bit.

Multiarea OSPF Configuration

The following figure shows five devices (R1 to R5) in a multi-area configuration.

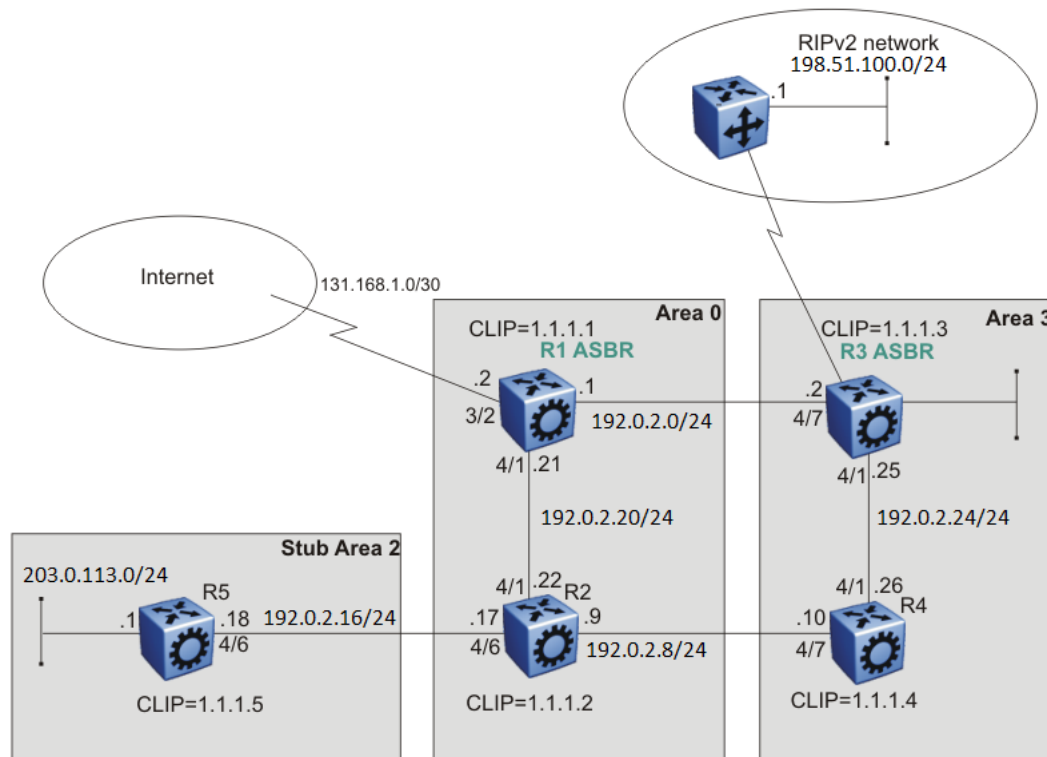


Figure 191: Multiarea configuration example

The following list explains the configuration for devices R1 through R5:

- R1 is an OSPF AS boundary router (ASBR) that is associated with OSPF Area 0 and OSPF Area 3. R1 distributes a default route for Internet traffic.
- R2 is an OSPF stub ABR for OSPF Area 2 and ABR to OSPF Area 3.
- R3 is an OSPF ASBR and distributes OSPF to RIP and RIP to OSPF.
- R4 is an OSPF internal router in Area 3.
- R5 is an internal OSPF subrouter in Area 2.
- All OSPF interfaces are brouter ports except R5.

Network 203.0.113.0/24 on R5 uses a VLAN configuration instead of a brouter port. This example uses brouter ports rather than VLANs because the spanning tree algorithm is disabled by default if you use brouter interfaces.

- All interfaces are Ethernet; therefore, the OSPF interfaces are broadcast, except the circuitless IP (CLIP) interfaces, which are passive.
- The interface priority on R5 is 0; therefore, R5 cannot become a DR.
- Configure the OSPF router priority so that R1 becomes the DR (priority 100) and R2 becomes the backup designated router (BDR) with a priority value of 50.

Use stub or NSSA areas to reduce the LSDB size by excluding external LSAs. The stub ABR advertises a default route into the stub area for all external routes.

Router types

Routers in an OSPF network can have various roles depending on how you configure them.

The following table describes the router types you can configure in an OSPF network.

Table 216: Router types in an OSPF network

Router type	Description
AS boundary router (ASBR)	A router that attaches at the edge of an OSPF network is an ASBR. An ASBR generally has one or more interfaces that run an interdomain routing protocol such as Border Gateway Protocol. In addition, a router that distributes static routes or RIP routes into OSPF is an ASBR. The ASBR forwards external routes into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area border router (ABR)	A router that attaches to two or more areas inside an OSPF network is an ABR. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information.
Internal router (IR)	A router that has interfaces only within a single area inside an OSPF network is an IR. Unlike ABRs, IRs have topological information only about the area in which they reside.
Designated router (DR)	In a broadcast or NBMA network, the routers elect a single router as the DR for that network. A DR makes sure that all routers on the network synchronize and advertises the network to the rest of the AS.
Backup designated router (BDR)	A BDR is elected in addition to the DR and, if the DR fails, can assume the DR role quickly.

To limit the amount of routing protocol traffic, the Hello protocol elects a DR and a BDR on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information, which on a large network can mean significant routing protocol traffic, all routers on the network form adjacencies with the DR and the BDR only, and send link-state information to them. The DR redistributes this information to every other adjacent router.

If the BDR operates in backup mode, it receives link-state information from all routers on the network and listens for acknowledgments. If the DR fails, the BDR can transition quickly to the role of DR because its routing tables are up-to-date.

OSPF Neighbors

In an OSPF network, two routers that have an interface to the same network are neighbors. Routers use the Hello protocol to discover their neighbors and to maintain neighbor relationships. On a broadcast or point-to-point network, the Hello protocol dynamically discovers neighbors. On an NBMA network, you must manually configure neighbors for the network.

The Hello protocol provides bidirectional communication between neighbors. Periodically, OSPF routers send hello packets over all interfaces. Included in these hello packets is the following information:

- router priority
- router hello timer and dead timer values
- list of routers that sent the router hello packet on this interface
- router choice for DR and backup designated router (BDR)

Bidirectional communication is determined after one router discovers itself listed in the hello packet of its neighbor.

NBMA interfaces whose router priority is a positive, nonzero value are eligible to become DRs for the NBMA network and are configured with a list of all attached routers. The neighbors list includes each neighbor IP address and router priority. In an NBMA network, a router with a priority other than zero is eligible to become the DR for the NBMA network. You must manually configure the IP address, mask, and router priority of neighbors on routers that are eligible to become the DR or BDR for the network.

Log messages indicate when an OSPF neighbor state change occurs. Each log message indicates the previous state and the new state of the OSPF neighbor. The log message generated for system traps also indicates the previous state and the current state of the OSPF neighbor.

Neighbors can form an adjacency to exchange routing information. After two routers form an adjacency, they perform a database exchange process to synchronize their topological databases. After the databases synchronize, the routers are fully adjacent. Adjacency conserves bandwidth because, from this point, the adjacent routers pass only routing change information.

All routers connected by a point-to-point network or a virtual link always form an adjacency. All routers on a broadcast or NBMA network form an adjacency with the DR and the BDR.

In an NBMA network, before the routers elect a DR, the router sends hello packets only to those neighbors eligible to become a DR. The NBMA DR forms adjacencies only with its configured neighbors and drops all packets from other sources. The neighbor configuration also notifies the router of the expected hello behavior for each neighbor.

If a router receives a hello packet from a neighbor with a priority different from that which is already configured for the neighbor, the router can automatically change the configured priority to match the dynamically learned priority.

OSPF Interfaces

Configure an OSPF interface, or link, on an IP interface. An IP interface can be either a single link (brouter port) or a logical interface configured on a VLAN (multiple ports). The state information associated with the interface is obtained from the underlying lower-level protocols and the routing protocol itself.



Important

To change the interface type of an enabled OSPF interface, you must first disable it, change the type, and then reenabling it. For an NBMA interface, you must first delete manually configured neighbors.

OSPF network types allow OSPF-neighboring between routers over various types of network infrastructures. You can configure each interface to support various network types.

The following table describes the supported OSPF network interface types:

Table 217: OSPF Network Types

Network interface type	Description
Broadcast Interfaces on page 2757	Broadcast interfaces automatically discover every OSPF router on the network by sending OSPF hello packets to the multicast group All OSPF Routers (224.0.0.5). Neighboring is automatic and requires no configuration.
Non-Broadcast Multiple Access Interfaces on page 2757	The NBMA network type models network environments that do not have native Layer 2 broadcast or multicast capabilities, such as Frame Relay and X.25. OSPF hello packets are unicast to manually configured neighbors.

Table 217: OSPF Network Types (continued)

Network interface type	Description
Passive Interfaces on page 2761	A passive interface is an interfacing network in OSPF that does not generate LSAs or form adjacencies. Use a passive interface on an access network or on an interface used for BGP peering. Using passive interfaces limits the amount of CPU cycles required to perform the OSPF routing algorithm.
Point-to-Point Interfaces on page 2761	A point-to-point interface is a single connection between two specific points or OSPF routers. Point-to-Point interfaces automatically discover OSPF routers (up to two) on the network by sending OSPF hello packets to the multicast group AllSPFRouters (224.0.0.5). Neighboring is automatic and requires no configuration.

Broadcast Interfaces

Broadcast interfaces support many attached routers and can address a single physical message to all attached broadcast routers (sent to AllSPFRouters and AllDRouters).

Broadcast interfaces dynamically discover neighboring routers using the OSPF Hello protocol. Each pair of routers on a broadcast network, such as an Ethernet, communicate directly.

Non-Broadcast Multiple Access Interfaces

An NBMA network interconnects multiple devices through point-to-point links. NBMA does not use broadcast and multicast data transmission.

NBMA interfaces support many routers, but cannot broadcast. NBMA networks perform the following activities:

- Statically establish OSPF neighbor relationships.
You must establish neighbor relationships because hub-and-spoke Wide Area Network (WAN) topologies do not support any-to-any broadcasting.
- Control meshed WAN connections.

In contrast to a broadcast network, where some OSPF protocol packets are multicast (sent to AllSPFRouters and AllDRouters), OSPF packets on an NBMA interface are replicated and sent in turn to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination address AllSPFRouters and AllDRouters.

The following figure shows an example of four routers attached to an NBMA subnet. The NBMA segment uses a single IP subnet and each router uses an IP address within the subnet.




Figure 192: NBMA Subnet

NBMA Interface Operations and Parameters

OSPF treats an NBMA network much like it treats a broadcast network. Because many routers attach to the network, the Hello protocol elects a designated router (DR) to generate the network link-state advertisements.

Because the NBMA network does not broadcast, you must manually configure neighbors for each router eligible to become DR (those networks with a positive, nonzero router priority value). You must also configure a poll interval for the network.

NBMA interfaces with a positive, nonzero router priority can become DR for the NBMA network and contain a list of all attached routers, or neighbors. This neighbors list includes each neighbor IP address and router priority.

The router uses neighbor information both during and after the DR election process. After an interface to a nonbroadcast network with a nonzero priority initializes, and before the Hello protocol elects a DR, the router sends hello packets only to those neighbors eligible to become DR. After the Hello protocol elects a DR, it forms adjacencies only with its configured neighbors and drops all packets from other sources. This neighbor configuration also notifies the router of the expected hello behavior of each neighbor.

If a router eligible to become the DR receives a hello packet from a neighbor that shows a different priority from that which is already configured for this neighbor, the DR changes the configured priority to match the dynamically learned priority.

Configure an NBMA interface with a poll interval. The poll interval designates the interval at which the router sends hello packets to inactive neighboring routers. The router typically sends hello packets at the Hello interval, for example, every 10 seconds. If a neighboring router becomes inactive, or if the router does not receive hello packets for the established RouterDeadInterval period, the router sends hello packets at the specified poll interval, for example, every 120 seconds.

You must configure a neighbors list for the DR to allow an NBMA network to send hello packets. If the router is eligible to become a DR, it periodically sends hello packets to all neighbors that are also eligible. The effect of this action is that two eligible routers always exchange hello packets, which is necessary for the correct DR election. You can minimize the number of hello packets by minimizing the number of eligible routers on a nonbroadcast network.

After the Hello protocol elects a DR, it sends hello packets to all manually configured neighbors to synchronize their link-state databases, establish itself as the DR, and identify the backup designated router (BDR).

If a router is not eligible to become DR, it periodically sends hello packets to both the DR and the BDR. The router also sends a hello packet in reply to a hello packet received from an eligible neighbor (other than the current DR and BDR). This process establishes an initial bidirectional relationship with a potential DR.

When a router sends hello packets to a neighbor, the neighbor state determines the interval between hello packets. If the neighbor is in the down state, the router sends hello packets at the designated poll interval, for example, every 120 seconds. Otherwise, the router sends hello packets at the designated hello interval, for example, every 10 seconds.

OSPF and NBMA Example: Adjacency Formation

In an NBMA network, as in a broadcast network, all routers become adjacent to the DR and the BDR. The adjacencies form after you assign the router priorities, configure the neighbors, and the Hello protocol elects the network DR.

The following figure shows an NBMA subnet with router priorities and manually configured neighbors.




Figure 193: NBMA Subnet Configuration Example

Because R1 and R2 have a router priority of 0, they are not eligible to become the DR. Also, R1 and R2 do not require configuration of a neighbors list; R1 and R2 discover neighbors dynamically through the Hello protocol.

R3 and R4 both have a positive, nonzero priority and are eligible to become the DR. Manually configure neighbor lists on R3 and R4.

To create this NBMA network, configure the following parameters:

1. On each router: NBMA interface type, poll interval, router priority
2. On R3: R1, R2, and R4 as neighbors
3. On R4: R1, R2, and R3 as neighbors

If all routers start at the same time, the routers perform the following steps:

1. R3 and R4 send each other a hello packet to elect a DR.
2. The Hello protocol elects R3 as the DR, and R4 as the BDR.
3. R3 (DR) and R4 (BDR) send hello packets to all other routers on the NBMA subnet to synchronize their link-state databases and establish themselves as DR and BDR.

4. R1 and R2 reply to R3 and R4.
5. R3 and R4 each form three adjacencies (one with each router on the NBMA subnet).
6. R1 and R2 each form two adjacencies (one with the DR and one with the BDR).

Passive Interfaces

Use a passive interface to enable an interface to advertise into an OSPF domain while limiting its adjacencies.

After you change the interface type to passive, the router advertises the interface into the OSPF domain as an internal stub network with the following behaviors:

- Does not send hello packets to the OSPF domain.
- Does not receive hello packets from the OSPF domain.
- Does not form adjacencies in the OSPF domain.

If you configure an interface as passive, the router advertises it as an OSPF internal route. If the interface is not a passive interface, to advertise a network into OSPF and not form OSPF adjacencies, you must configure the interface as nonOSPF, and the router must redistribute the local network as an autonomous system external (ASE) LSA.

Point-to-Point Interfaces

A point-to-point interface supports a single OSPF router on either end of the link and can address a single physical message between the two routers (sent to AllSPFRouters). A point-to-point link does not elect a designated router (DR) or a backup designated router (BDR).



Note

Unnumbered point-to-point interfaces are not supported.

OSPF Packets

All OSPF packets start with a 24-octet header that contains information about the OSPF version, the packet type and length, the ID of the router that transmits the packet, and the ID of the OSPF area that sends the packet. An OSPF packet is one of the following types:

- The router transmitted hello packets between neighbors and never forwards them. The Hello protocol requires routers to send hello packets to neighbors at predefined hello intervals. A neighbor router that does not receive a hello packet declares the other router dead.
- The router exchanges DD packets after neighboring routers establish a link, which synchronizes their LSDBs.
- Link-state request packets describe one or more link-state advertisements that a router requests from its neighbor. Routers send link-state requests if the information received in DD packets from a neighbor is not consistent with its own link-state database.

- Link-state update packets contain one or more LSAs and the router sends them following a change in network conditions.
- The router sends link-state acknowledgement packets to acknowledge receipt of link-state updates. Link-state acknowledgement packets contain the headers of the received LSAs.

Intra-area Link-state Advertisements

OSPF does not require each router to send its entire routing table to its neighbors. Instead, each OSPF router floods only link-state change information in the form of Link-state advertisements (LSA) throughout the area or Autonomot System (AS). LSAs in OSPF are one of the following five types:

- A router links advertisement is flooded only within the area and contains information about neighbor routers and the LANs to which the router attaches. A backbone router can flood router link advertisements within the backbone area.
- A DR on a LAN generates network links advertisement to list all routers on that LAN, and floods network links advertisements only within the area. A backbone DR can flood network links advertisements within the backbone area.
- An ABR floods a network summary link advertisement into an area and describes networks that are reachable outside the area. An ABR attached to two areas generates a different network summary link advertisement for each area. ABRs also generate area summary link advertisements that contain information about destinations within an area that are flooded to the backbone area.
- An ASBR summary link advertisement describes the cost of the path to an ASBR from the router that generates the advertisement.
- An ASBR sends an ASE link advertisement to describe the cost of the path to a destination outside the AS from the ASBR that generates the advertisement. This information is flooded to all routers in the AS.

OSPF and IP

OSPF runs over IP, which means that an OSPF packet transmits with an IP data packet header. The protocol field in the IP header is 89, which identifies it as an OSPF packet and distinguishes it from other packets that use an IP header.

An OSPF route advertisement expresses a destination as an IP address and a variable-length mask. Together, the address and the mask indicate the range of destinations to which the advertisement applies.

Because OSPF can specify a range of networks, it can send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 192.0.2.0 with a mask of 255.255.0.0 describes a single route to destinations 192.0.2.0 to 192.0.2.255.

ASE routes

OSPF considers the following routes as ASE routes:

- a route to a destination outside the AS
- a static route
- a default route
- a route derived by RIP
- a directly connected network that does not run OSPF

Dijkstras algorithm

A separate copy of the OSPF routing algorithm (Dijkstra's algorithm) runs in each area. Routers that connect to multiple areas run multiple copies of the algorithm. The sequence of processes governed by the routing algorithm is as follows:

1. After a router starts, it initializes the OSPF data structures, and then waits for indications from lower-level protocols that the router interfaces are functional.
2. A router then uses the Hello protocol to discover neighbors. On point-to-point and broadcast networks the router dynamically detects neighbors by sending hello packets to the multicast address AllSPFRouters. On Non-Broadcast Multiple Access (NBMA) networks, you must provide some configuration information to discover neighbors.
3. On all multiaccess networks (broadcast or nonbroadcast), the Hello protocol elects a designated router (DR) for the network.
4. The router attempts to form adjacencies with some of its neighbors. On multiaccess networks, the DR determines which routers become adjacent. This behavior does not occur if you configure a router as a passive interface because passive interfaces do not form adjacencies.
5. Adjacent neighbors synchronize their topological databases.
6. The router periodically advertises its link state, and does so after its local state changes. LSAs include information about adjacencies, enabling quick detection of dead routers on the network.
7. LSAs flood throughout the area to ensure that all routers in an area have an identical topological database.
8. From this database each router calculates a shortest-path tree, with itself as the root. This shortest-path tree in turn yields a routing table for the protocol.

OSPF Metrics

For OSPF, the best path to a destination is the path that offers the least-cost metric (least-cost delay). You can configure OSPF cost metrics to specify preferred paths. You can configure metric speed globally or for specific ports and interfaces on the network. In addition, you can control redistribution options between non OSPF interfaces and OSPF interfaces.

Assign default metric speeds for different port types, such as 10 Mb/s or 1 Mb/s ports. You can specify a new metric speed for an IP interface. An IP interface can be a brouter port or a VLAN.

OSPF virtual links

On an OSPF network, a switch that acts as an ABR must connect directly to the backbone. If no physical connection is available, you can automatically or manually configure a virtual link.

An automatic virtual link can provide redundancy support for critical network connections. Automatic virtual linking creates virtual paths for vital traffic paths in your OSPF network. If a connection fails on the network, such as after an interface cable that provides connection to the backbone (either directly or indirectly) disconnects from the switch, the virtual link is available to maintain connectivity.

Use automatic virtual linking to ensure that a link is created to another router. If automatic virtual linking uses more resources than you want to expend, creating a manual virtual link can be the better solution. Use this approach to conserve resources and control virtual links in the OSPF configuration.

On the switch, OSPF behavior follows OSPF standards; the router cannot learn OSPF routes through an ABR unless the ABR connects to the backbone or through a virtual link.

The following figure shows how to configure a virtual link between the ABR in area 2.2.2.2 and the ABR in area 0.0.0.0.




Figure 194: Virtual link between ABRs through a transit area

To configure a virtual link between the ABRs in area 1 and area 3, define area 2 as the transit area between the other two areas, and identify R2 as the neighbor router through which R2 must send information to reach the backbone through R1.

OSPF ASBRs

ASBRs advertise nonOSPF routes into OSPF domains so that they can pass through the OSPF routing domain. A router can function as an ASBR if one or more interfaces connects to a nonOSPF network, for example, RIP, BGP, or Exterior Gateway Protocol (EGP).

An ASBR imports external routes into the OSPF domain by using ASE LSAs (LSA type 5) originated by the ASBR.

ASE LSAs flood across area borders. When an ASBR imports external routes, it imports OSPF route information using external type 1 or type 2 metrics. The result is a four-level routing hierarchy, as shown in the following table, according to routing preference.

Table 218: ASBR routing hierarchy

Level	Description
1	Intra-area routing
2	Inter-area routing
3	External type 1 metrics
4	External type 2 metrics

The use of these metrics results in a routing preference from most preferred to least preferred of

- routing within an OSPF area
- routing within the OSPF domain
- routing within the OSPF domain and external routes with external type 1 metrics
- routing within the OSPF domain and external routes with external type 2 metrics

For example, an ASBR can import RIP routes into OSPF with external type 1 metrics. Another ASBR can import Internet routes and advertise a default route with an external type 2 metric. This results in RIP-imported routes that have a higher preference than the Internet-imported default routes. In reality, BGP Internet routes must use external type 2 metrics, whereas RIP imported routes must use external type 1 metrics.

Routes imported into OSPF as external type 1 are from IGP's whose external metric is comparable to OSPF metrics. With external type 1 metrics, OSPF adds the internal cost of the ASBR to the external metric. EGP's, whose metric is not comparable to OSPF metrics, use external type 2 metrics. External type 2 metrics use only the internal OSPF cost to the ASBR in the routing decision.

To conserve resources, you can limit the number of ASBRs in your network or specifically control which routers perform as ASBRs to control traffic flow.

Area link-state advertisements

The following table explains the seven LSA types exchanged between areas. LSAs share link-state information among routers. LSAs typically contain information about the

router and its neighbors. OSPF generates LSAs periodically to ensure connectivity or after a change in state of a router or link (that is, up or down).

Table 219: OSPF LSA types

LSA type	Description	Area of distribution
1	A router originates type 1 LSAs (router LSAs) to describe its set of active interfaces and neighbors.	Passed only within the same area
2	Type 2 LSAs (network LSAs) describe a network segment such as broadcast or NBMA. In a broadcast network, the DR originates network LSAs.	Passed only within the same area
3	The ABR originates type 3 LSAs (network-summary LSAs) to describe the networks within an area.	Passed between areas
4	Type 4 LSAs (ASBR-summary LSAs) advertise the location of the ASBRs from area to area.	Passed between areas
5	Type 5 LSAs (ASE LSAs) describe networks outside of the OSPF domain. The ASBR originates type 5 LSAs. In stub and NSSA areas, a single default route replaces type 5 LSA routes.	Passed between areas
6	Type 6 LSAs (group membership LSAs) identify the location of multicast group members in multicast OSPF.	Passed between areas
7	Type 7 LSAs import external routes in OSPF NSSAs.	Translated between areas

OSPF Security Mechanisms

The switch implementation of OSPF includes security mechanisms to prevent unauthorized routers from attacking the OSPF routing domain. These security mechanisms prevent a malicious person from joining an OSPF domain and advertising false information in the OSPF LSAs. Likewise, security prevents a misconfigured router from joining an OSPF domain.

Simple password

The simple password security mechanism is a simple-text password; only routers that contain the same authentication ID in their LSA headers can communicate with each other.

Do not use this security mechanism because the system stores the password in plain text. A user or system can read the password from the configuration file or from the LSA packet.

Message Digest 5

Message Digest 5 (MD5) for OSPF security provides standards-based (RFC1321) authentication using 128-bit encryption, usually expressed as a 32-digit hexadecimal

number. When you use MD5 for OSPF security, it is almost impossible for a malicious user to compute or extrapolate the decrypting codes from the OSPF packets.

If you use MD5, each OSPF packet has a message digest appended to it. The digest must match between the sending and receiving routers. Both the sending and receiving routers calculate the message digest based on the MD5 key and padding, and then compare the results. If the message digest computed at the sender and receiver does not match, the receiver rejects the packet.

Secure hash algorithm 1

The secure hash algorithm 1 (SHA-1) is a cryptographic hash function that uses 160-bit encryption, usually given in a 40 digit hexadecimal number. SHA-1 is one of the most widely used of the existing SHA hash functions and is more secure than MD5.

SHA-1 takes a variable length input message and SHA-1 creates a fixed length output message referred to as the hash, or message digest, of the original message. If you use SHA-1 with OSPF, each OSPF packet has a message digest appended to it.

The message digest or hash must match between the sending and receiving routers. If the message digest computed at the sender and receiver does not match, the receiver rejects the packet. The hash functions produce a type of checksum or summary of the input.

It is almost impossible to determine the original input message based on the output hash message.

A cryptographic hash function is fully defined and uses no secret key.

Secure hash algorithm 2

Secure hash algorithm 2 (SHA-2) is also a cryptographic hash function. SHA-2 updates SHA-1 and offers six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits message digest size values. Output size depends on the hash function, so, for instance, SHA-256 is 256 bits.

SHA-2 is more secure than SHA-1 and MD5.

SHA-2 works similarly to SHA-1, in that SHA-2 takes a variable length input message and creates a fixed length output message referred to as the hash, or message digest, of the original message. If you use SHA-2 with OSPF, each OSPF packet has a message digest appended to it. Among the differences in SHA-2 from SHA-1 are an increased bit encryption length.

Similarly with other hash functions, for SHA-2, the message digest or hash must match between the sending and receiving routers. If the message digest computed at the sender and receiver does not match, the receiver rejects the packet. The hash functions produce a type of checksum or summary of the input.

OSPFv3

OSPFv3, defined in RFC 2740 and updated by RFC 5340, extends OSPF to support IPv6 and introduces several protocol changes for IPv6 networks. In OSPFv3, IPv6 replaces the IPv4 concepts of subnets and networks with links; a link can carry multiple prefixes and allows nodes with different prefixes to communicate directly. OSPFv3 operates on each link rather than each subnet and uses IPv6 link-local addresses for most control traffic. OSPFv3 packets do not carry IPv6 addresses in their headers; instead, LSAs include all addressing information, while area IDs, link-state IDs, and router IDs remain 32-bit values. Routers identify OSPFv3 neighbors by Router ID and use additional fields such as interface ID and instance ID to support multiple OSPFv3 instances on the same link.

Flooding Scope

LSA flooding scope is generalized in OSPFv3 and coded in the LS type field of the LSA. The following three flooding scopes are available for LSAs:

- Link scope: The LSA is not flooded beyond the local link.
- Area scope: The LSA is flooded in a single OSPF area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix-LSAs.
- AS scope: The LSA is flooded through the routing domain. AS scope is used for ASexternal-LSAs.

Link-local Addresses

IPv6 uses link-local addresses on a single link. Link-local addresses facilitate features such as neighbor discovery and autoconfiguration. Datagrams with link-local sources are not forwarded. Instead, routers assign link-local unicast addresses from the IPv6 address range.

OSPF for IPv6 does not assign link-local unicast addresses to physical segments attached to a router; it assumes that each router already has link-local unicast addresses assigned. The source for all OSPF packets sent on OSPF physical interfaces is the associated link-local unicast address. Routers learn link-local addresses for all other nodes on links. The nexthop information during packet forwarding includes the learned addresses.

OSPFv3 packets always use link-local addresses as the source and destination, except on a virtual link. All OSPFv3 packets sent over a virtual link use global addresses.

Link LSA is the only OSPF LSA type that includes link-local addresses. Link-local addresses must not be advertised in other LSA types.

Authentication

OSPFv3 for IPv6 requires the IP authentication header and the IP encapsulating security payload for authentication and security. OSPFv3 does not support the authentication feature from OSPFv2.

Packet Format

OSPFv3 runs directly over IPv6. All other addressing information is absent in OSPF packet headers. OSPFv3 is network-protocol-independent. LSA types contain addressing information.

OSPFv3 implements the following packet changes from OSPFv2:

- The hello packet and database description packet operations fields are expanded to 24 bits.
- The packet header does not include Authentication and AuType fields.
- The interface ID replaces the address information in the hello packet. The Interface ID becomes the network LSA link-state ID, if the router becomes the designated router on the link.
- Router-bit (R-bit) and V6-bit in the options field process router LSAs during Shortest Path First (SPF) calculation. R-bits and V6-bits determine participation in topology distribution. The V6-bit specializes the R-bit. If the V6-bit is clear, the OSPF speaker can participate in the OSPF topology distribution without forwarding IPv6 datagrams. If the R-bit is set and the V6-bit is clear, the OSPF speaker still does not forward IPv6 datagrams, but it can forward IPv4 datagrams.
- The packet header includes the instance ID, which allows multiple OSPF protocol instances on the same link.

R-bit

Unlike OSPF for IPv4, OSPFv3 for IPv6 supports the R-bit. The R-bit indicates whether the originating node is an active router. If the R-bit is cleared, routes that transit the advertising node cannot be calculated.

For example, if a multi-homed host participates in routing without forwarding non-locally-addressed packets, the R-bit is cleared.

An IPv6-enabled switch can continue to operate as an OSPFv3 neighbor even if you disable IPv6 forwarding on the switch. This behavior differs from IPv4 OSPF, in which the switch drops a neighbor, if IP forwarding on the neighbor is disabled.

If an OSPFv3 neighbor does not provide the R-bit in the Network Discovery (ND) packet, the system enables R-bit for every OSPFv3 neighbor with dependent routes to avoid neighbor deletion from inactivity. This IPv6 enhancement is to ensure an OSPFv3 neighbor without R-bit that experiences a timeout can trigger the Network Unreachability Detection (NUD) instead of being deleted.

Link LSAs

The link LSA uses link flooding scope, not flooded beyond the associated link.

Link LSAs have three purposes:

- to provide the link-local address of the router to all other nodes on the link
- to provide the list of IPv6 prefixes associated with the link
- to allow the router to associate options bits with the network LSA for the link

Intra-Area-Prefix LSA

The Intra-Area-Prefix-LSA carries all IPv6 prefix information. In IPv4, this information is in router LSAs and network LSAs.

Unknown LSA types

In OSPFv3, unknown LSA types are either stored and flooded as though understood or given link flooding scope. Specific behavior is coded in the LS type field of the header.

Link LSA Suppression

To decrease unnecessary link LSA generation and flooding for non-broadcast and non-NBMA interfaces, the Link LSA Suppression interface configuration parameter has been added in RFC 5340. If Link LSA Suppression is configured for an interface, and the interface type is not broadcast or NBMA, the originated link LSA may be suppressed. Link LSA suppression is disabled, by default. For more information on configuration see, [Configure OSPF on a Port or VLAN](#) on page 2837.

Stub Area

OSPFv3 retains the concept of stub areas, which minimize link-state databases and routing table sizes.

IPv6 stub areas carry only router LSAs, network LSAs, Inter-Area-Prefix-LSAs, link LSAs, and Intra-Area-Prefix-LSAs.

Unlike IPv4, IPv6 can store LSAs with unrecognized link-state (LS) types or flood them as though they are understood. Rules applied to the stub area prevent the excessive growth of the link-state database.

Deprecation of MOSPF for IPV6

OSPFv3 in RFC 5340 deprecates Multicast Extensions to OSPF (MOSPF) support, and its attendant protocol fields.

NSSA Specification

RFC 2740 partially specifies this protocol feature, the level of specification was insufficient to implement it. However, RFC 5340 includes an NSSA specification unique to OSPFv3. This specification coupled with NSSA provide sufficient specification for implementation. Current Infinity IPv6 OSPF has full support for NSSA feature and is consistent with the additional specifications in RFC 5340.

Stub Area Unknown LSA Flooding Restriction Deprecated

In RFC 2740, flooding of unknown LSA was restricted within stub and NSSA areas. Following were the restrictions:

- Unlike IPv4, in IPv6 you can label unrecognized LS types as "Store and flood the LSA, as if type understood". Uncontrolled introduction of such LSAs could cause a stub area's link-state database to grow larger than its component router's capacities
- To guard against this situation, the following rule regarding stub areas has been established:

An LSA whose LS type is unrecognized can be flooded only into a stub area, if both the LSAs have area or link-local flooding scope, and the LSA has U-bit set to 0

Now these restrictions have been deprecated. OSPFv3 routers flood link and area scope LSAs whose LS type is unrecognized and U-bit is set to 1 throughout stub and NSSA areas. The only backward compatibility issue is that the OSPFv3 routers still supporting the restrictions may not propagate newly defined LSA types.

LSA Options and Prefix Options Updates

The LSA Options and Prefix Options fields have been updated to reflect recent protocol additions. Specifically, bits related to MOSPF have been deprecated, Options field bits common with OSPFv2 have been reserved, and the DN-bit has been added to the prefix- options.

IPv6 Site-Local Address

All references to IPv6 site-local addresses have been removed in RFC 5340. Infinity IPv6 OSPF does not contain any reference to IPv6 site-local addresses and is already compliant with RFC 5340 for this.

IPsec Support with OSPFv3

You can use Internet Protocol Security (IPsec) with OSPFv3 virtual link for the security protection of communication between the end points. You can also use IPsec with OSPFv3 on a brouter port or VLAN interface, for example, if you want to encrypt OSPFv3 control traffic on a broadcast network.

OSPF virtual link provides connectivity to the OSPF backbone area for redundancy or to provide a virtual link if a physical connection is not possible.

Because the device does not know the IPv6 addresses of the OSPFv3 virtual link end points at the time of configuration, you cannot manually configure the security policy ahead of time. The system must self-manage its security policy dynamically. The device also dynamically manages the IPsec enable flag, which the virtual link uses on a Layer 2 interface, either a VLAN or brouter port interface.

The following events can trigger an IPsec policy activation:

1. An OSPFv3 routing module detects the establishment of a virtual link.
2. IPsec is enabled on the already established virtual link.

On the other hand, the following two events can dynamically trigger an IPsec policy deactivation:

1. The virtual link is turn down.
2. IPsec is disabled on the virtual link.

IPsec policies can also change dynamically if a neighbor address or a local address changes.

You can enable IPsec support for IPv6 OSPF virtual link at the system level through CLI. You must disable IPsec before you can perform virtual link policy configuration changes.

Until you enable IPsec on both sides of the virtual links, the links cannot exchange OSPFv3 control messages, and the system drops OSPFv3 exchange packets.

You can configure the direction you want IPsec to protect, either, ingress, egress, or both. In addition, you can permit or drop communication for the OSPF virtual link.

You can also use IPsec with OSPFv3 on a router port or VLAN interface. For a full configuration example, see [OSPFv3 IPsec configuration example](#) on page 2085 and [OSPFv3 virtual link IPsec configuration example](#) on page 2092.

OSPF and Route Redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if OSPF routes exist in a router and they must travel through a BGP network, then configure redistribution of OSPF routes through BGP. This configuration sends OSPF routes to a router that uses BGP.

You can redistribute routes

- on an interface basis
- on a global basis between protocols on a single VRF instance (intraVRF)
- between the same or different protocols on different VRF instances (interVRF)

To configure interface-based redistribution, configure a route policy, and then apply it to the interface. Configure the match parameter to the protocol from which to learn routes.

Use the **ip ospf redistribute** command to accomplish the (intraVRF) redistribution of routes through OSPF, so that OSPF redistribution occurs globally on all OSPF-

enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

**Note**

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

OSPF Route Redistribution and DvR

DvR Controllers redistribute routes (direct routes, static routes and the default route) into the DvR domain. You can configure redistribution of DvR host routes into OSPF.

Tagged Routes

You can enable route filtering to ensure tagged IS-IS routes are not redistributed into OSPF. For more information, see [IS-IS Route Tagging](#) on page 1200.

ECMP with OSPFv3

The ECMP feature supports and complements OSPFv3 protocol.

With Equal Cost Multipath (ECMP), you can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.

**Note**

To add OSPFv3 equal cost paths in the routing table, you must first enable IPv6 ECMP feature globally.

For scaling information on the ECMP paths supported per destination prefix, see *Fabric Engine Release Notes*.

OSPFv3 and Route Redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if OSPFv3 routes exist in a router and they must travel through a BGP network, then configure

redistribution of OSPFv3 routes through BGP. This configuration sends OSPFv3 routes to a router that uses BGP.

You can redistribute routes on a global basis between protocols on a single VRF instance (intraVRF).

Use the **ipv6 ospf redistribute** command to accomplish the (intraVRF) redistribution of routes through OSPF, so that OSPF redistribution occurs globally on all OSPF-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For a redistribute policy (OSPFv3) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

OSPFv3 Support for CLIP Interfaces

You can configure the Open Shortest Path First Version 3 (OSPFv3) instance on circuitless IP (CLIP) interfaces for the Global Router or a specific Virtual Router Forwarding (VRF) instance. The switch supports a maximum of 64 OSPFv3 CLIP interfaces.

The CLIP interfaces are virtual interfaces and are not associated with the physical ports on the switch. The system considers the OSPFv3 CLIP instances as passive and the system copies the IPv6 address associated with the OSPFv3 instance into the intra-area-prefix-LSA. You must follow a two-step configuration to create an OSPFv3 instance on the CLIP interface and enable the instance, for the system to send the intra-area-prefix-LSA of the OSPFv3 instance to the neighbors. For more information, see [Configure OSPFv3 on the CLIP Interface for the Global Router](#) on page 2843 and [Configure OSPFv3 on the CLIP Interface for a VRF Instance](#) on page 2844 for CLI or [Configure OSPFv3 on a CLIP Interface](#) on page 2906. for EDM.

OSPF Graceful Restart

In many OSPF networks, OSPF routers remove a restarting OSPF router from the network topology, if the router is restarted. This action causes all OSPF routers to re-converge and route around the restarting router. The OSPF Graceful Restart feature is an OSPF enhancement to allow an OSPF router to stay on the forwarding path when the software is restarting.

This feature is documented under RFC 3623 for OSPFv2 (IPv4) and RFC 5187 for OSPFv3 (IPv6). The switch software supports only helper mode for both OSPFv2 and OSPFv3 protocols.

Helper Mode

Helper mode is a part of the OSPF Graceful restart feature. Helper mode uses the OSPF routers to help other OSPF routers on the network stay on the forwarding path while the software is restarting. The OSPF router sends a type of LSA called a GRACE-LSA to inform the other OSPF routers that it is restarting the software. When an OSPF router receives a GRACE-LSA from a neighbor OSPF Router, it enters the Helper mode for that neighbor on that network. An OSPF router supports Helper mode by default.

Operations of Helper Mode

The following section describes the operations in the Helper mode:

- Entering Helper mode — An OSPF router enters the Helper mode provided the following conditions are true:
 - The router is fully adjacent with the neighbor already.
 - No changes have been made in the LSDB since the neighbor router started.
 - The grace period has not expired.
 - Local policy configured parameters allow it to help the neighbor.
 - The router is not in the process of restarting itself.

The OSPF router does not help the neighbor if the preceding conditions are not met.

If the OSPF router is already helping a neighbor, and receives another GRACE-LSA from the neighbor, it accepts the latest GRACE-LSA, and updates the grace period accordingly. The OSPF router in Helper mode continues to advertise its LSAs like the neighbor it is helping is still full, until any changes are made on the network during the grace period.

- Exiting Helper mode — An OSPF router exits the Helper mode, under the following conditions:
 - The GRACE-LSA is flushed. It means graceful restart has successfully terminated.
 - The GRACE-LSA's grace period expires.
 - There is a network topology change.

When an OSPF router exits Helper mode, the following actions occur:

- It recalculates the DR for the network.
- It re-originates its router LSA.
- If it is the DR, it re-originates the network LSA for the network.
- If it is a virtual link, it re-originates the router LSA for the virtual link transit area.

OSPF Configuration Considerations

This section describes considerations to keep in mind as you configure OSPF.

OSPF host route advertisements and nonbackbone areas

The switch does not associate a host route with a specific area. Therefore, if you create a host route in a nonbackbone area, nonbackbone (nonOSPF core) areas do not advertise it.

For example, in an OSPF network with multiple areas, including areas not adjacent to the core, which use virtual links, a host route on a router that belongs to a nonOSPF core area is not advertised on noncore routers.

To ensure host route advertisement, disable and enable OSPF on the noncore routers.

OSPF with switch clustering

If the network loses the DR, the BDR immediately becomes the new DR on the broadcast segment. After OSPF elects the new DR, all routers perform an SPF run and issue new LSAs for the segment. The new DR generates a new network LSA for the segment and every router on the segment must refresh the router LSA.

Each router performs the SPF run as soon as it detects a new DR. Depending on the speed of the router, the router can perform the SPF run before it receives the new LSAs for the segments, which requires a second SPF run to update and continue routing across the segment. The OSPF hold-down timer does not permit two consecutive SPF runs within the value of the timer. This limitation can lead to traffic interruption of up to 10 seconds.

In a classical OSPF routed design, this situation never causes a problem because OSPF runs over multiple segments so even if a segment is not usable, routes are recalculated over alternative segments. Typical Routed Split MultiLink Trunking (RSMLT) designs only deploy a single OSPF routed vlan, which constitutes a single segment.

You can use RSMLT in a configuration with dual core VLANs to minimize traffic interruption when the network loses the DR. This configuration creates a second OSPF core VLAN, forcing different nodes to become the DR for each VLAN. Each OSPF core VLAN has a DR (priority of 100) and no BDRs. This configuration does not require a BDR because the two VLANs provide backup for each other from a routing perspective. See the following figure for a network example.




Figure 195: RSMLT with dual core VLANs

Open Shortest Path First guidelines

Use OSPF to ensure that the switch can communicate with other OSPF routers. This section describes some general design considerations and presents a number of design scenarios for OSPF.

OSPF LSA limits

To determine OSPF link-state advertisement (LSA) limits:

1. Use the command **show ip ospf area** to determine the LSA_CNT and to obtain the number of LSAs for a given area.
2. Use the following formula to determine the number of areas. Ensure the total is less than 16,000 (16K):

$$\sum Adj_N * LSA_CNT_N < 16k$$

$N = 1$ to the number of areas for each switch

Adj_N = number of adjacencies for each Area N

LSA_CNT_N = number of LSAs for each Area N

For example, assume that a switch has a configuration of three areas with a total of 18 adjacencies and 1000 routes. This includes:

- 3 adjacencies with an LSA_CNT of 500 (Area 1)
- 10 adjacencies with an LSA_CNT of 1000 (Area 2)
- 5 adjacencies with an LSA_CNT of 200 (Area 3)

Calculate the number as follows:

$$3*500+10*1000+5*200=12.5K < 16K$$

This configuration ensures that the switch operates within accepted scalability limits.

OSPF design guidelines

Follow these additional OSPF guidelines:

- OSPF timers must be consistent across the entire network.
- Use OSPF area summarization to reduce routing table sizes.
- Use OSPF passive interfaces to reduce the number of active neighbor adjacencies.
- Use OSPF active interfaces only on intended route paths.

Configure wiring-closet subnets as OSPF passive interfaces unless they form a legitimate routing path for other routes.

- Minimize the number of OSPF areas for each switch to avoid excessive shortest-path calculations.

The switch executes the Dijkstra algorithm for each area separately.

- Ensure that the OSPF dead interval is at least four times the OSPF hello-interval.
- Use MD5 authentication on untrusted OSPF links.
- Use stub or NSSAs as much as possible to reduce CPU overhead.

OSPF and CPU utilization

After you create an OSPF area route summary on an area border router, the summary route can attract traffic to the area border router for which the router does not have a specific destination route. Enabling ICMP unreachable-message generation on the switch can result in a high CPU utilization rate.

To avoid high CPU utilization, use a discard static route configuration. The discard static route is a route (equal to the OSPF summary route) with a next hop of 255.255.255.255. This configuration ensures that all traffic that does not have a specific next-hop destination route is dropped.

OSPF network design examples

You can use OSPF routing in the core of a network.

The following figure describes a simple implementation of an OSPF network: enabling OSPF on two switches (S1 and S2) that are in the same subnet in one OSPF area.




Figure 196: Example 1: OSPF on one subnet in one area

The routers in the preceding figure use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1, and the OSPF port uses an IP address of 192.0.2.1.
- S2 has an OSPF router ID of 1.1.1.2, and the OSPF port uses an IP address of 192.0.2.2.

The general method to configure OSPF on each routing switch is:

1. Enable OSPF globally.
2. Enable IP forwarding on the switch.
3. Configure the IP address, subnet mask, and VLAN ID for the port.
4. Disable RIP on the port, if you do not need it.
5. Enable OSPF for the port.

After you configure S2, the two switches elect a designated router and a backup designated router. They exchange hello packets to synchronize their link state databases.

The following figure shows a configuration in which OSPF operates on three switches. OSPF performs routing on two subnets in one OSPF area. In this example, S1 directly connects to S2, and S3 directly connects to S2, but traffic between S1 and S3 is indirect, and passes through S2.

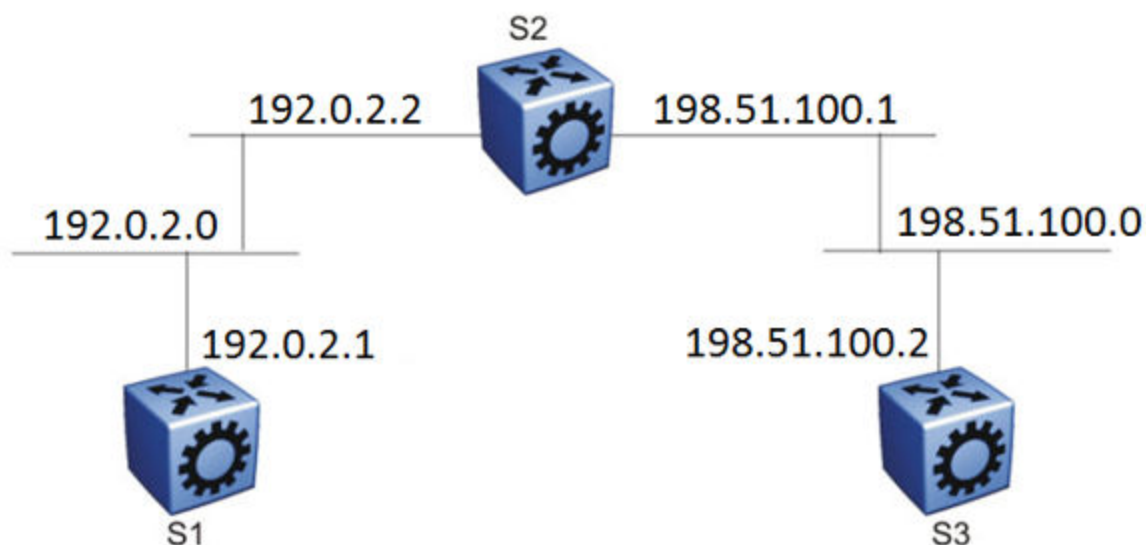


Figure 197: Example 2: OSPF on two subnets in one area

The routers in example 2 use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1, and the OSPF port uses an IP address of 192.0.2.1.
- S2 has an OSPF router ID of 1.1.1.2, and two OSPF ports use IP addresses of 192.0.2.2 and 198.51.100.1.
- S3 has an OSPF router ID of 1.1.1.3, and the OSPF port uses an IP address of 198.51.100.2.

The general method to configure OSPF on each routing switch is:

1. Enable OSPF globally.
2. Insert IP addresses, subnet masks, and VLAN IDs for the OSPF ports on S1 and S3, and for the two OSPF ports on S2. The two ports on S2 enable routing and establish the IP addresses related to the two networks.
3. Enable OSPF for each OSPF port allocated with an IP address.

After you configure all three switches for OSPF, they elect a designated router and a backup designated router for each subnet and exchange hello packets to synchronize their link-state databases.

The following figure shows an example where OSPF operates on two subnets in two OSPF areas. S2 becomes the area border router for both networks.




Figure 198: Example 3: OSPF on two subnets in two areas

The routers in scenario 3 use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1. The OSPF port uses an IP address of 192.0.2.1, which is in OSPF area 1.
- S2 has an OSPF router ID of 1.1.1.2. One port uses an IP address of 192.0.2.2, which is in OSPF area 1. The second OSPF port on S2 uses an IP address of 198.51.100.1, which is in OSPF area 2.
- S3 has an OSPF router ID of 1.1.1.3. The OSPF port uses an IP address of 198.51.100.2, which is in OSPF area 2.

The general method to configure OSPF for this three-switch network is:

1. On all three switches, enable OSPF globally.
2. Configure OSPF on one network.

On S1, insert the IP address, subnet mask, and VLAN ID for the OSPF port. Enable OSPF on the port. On S2, insert the IP address, subnet mask, and VLAN ID for the OSPF port in area 1, and enable OSPF on the port. Both routable ports belong to the same network. Therefore, by default, both ports are in the same area.

3. Configure three OSPF areas for the network.
4. Configure OSPF on two additional ports in a second subnet.

Configure additional ports and verify that IP forwarding is enabled for each switch to ensure that routing can occur. On S2, insert the IP address, subnet mask, and VLAN ID for the OSPF port in area 2, and enable OSPF on the port. On S3, insert the IP address, subnet mask, and VLAN ID for the OSPF port, and enable OSPF on the port.

The three switches exchange hello packets.

In an environment with a mix of switches and routers from different vendors, you may need to manually modify the OSPF parameter `RtrDeadInterval` to 40 seconds.

OSPFv3 Configuration Example

This section shows an example of OSPFv3 configuration. The following figure shows the network.




Figure 199: OSPFv3 configuration

To complete the configuration, you must perform the following actions:

- Configure an IPv6 VLAN (VLAN 2) with port member 1/1.
- Configure an IPv6 brouter port (1/2).
- Use IPv6 address 2001:100:102::/64.

Configure VLAN 2 and add port members.

```
vlan create 2 type port-mstprstp 0
vlan mlt 2 4
vlan members 2 1/1 portmember
interface vlan 2
ipv6 interface
ipv6 interface enable
ipv6 interface address 2001:100:102:202:0:0:0:1/64
exit
```

Enable OSPFv3 on VLAN 2.

```
# IPV6 OSPF VLAN CONFIGURATION
interface vlan 2
ipv6 ospf area 0.0.0.0
ipv6 ospf poll-interval 0
ipv6 ospf enable
exit
```

Create brouter port 1/2 with IPv6 and OSPFv3.

```
interface gigabitethernet 1/2
ipv6 interface vlan 3999
ipv6 interface enable
ipv6 interface address 2001:100:102:201:0:0:0:1/64
ipv6 ospf area 0.0.0.0
ipv6 ospf enable
exit
```

Verification

The following example shows the Global Router example for OSPFv3 Area configuration:

```
Switch:1#show ipv6 ospf area
```

```
=====
OSPF Area - GlobalRouter
```

```

=====
AREA_ID          STUB_AREA  NSSA   IMPORT_SUM TRANS_ROLE TRANS_STATE
-----
0.0.0.0          false     false true      always   disabled
STUB_METRIC STUB_METRIC_TYPE SPF_RUNS BDR_RTR_CNT ASBDR_RTR_CNT LSA_CNT LSACK_SUM
-----
10              ospfv3Metric  0      0          0          0          0

```

The following example shows the VRF example for OSPFv3 Area configuration:

```

Switch:1#show ipv6 ospf area vrf vrf1

=====
                        OSPF Area - VRF vrf1
=====
AREA_ID          STUB_AREA  NSSA   IMPORT_SUM TRANS_ROLE TRANS_STATE
-----
0.0.0.0          false     false true      candidate disabled
1.1.1.1          false     false true      candidate disabled
STUB_METRIC STUB_METRIC_TYPE SPF_RUNS BDR_RTR_CNT ASBDR_RTR_CNT LSA_CNT LSACK_SUM
-----
10              ospfv3Metric  3      0          0          0          0
10              ospfv3Metric  3      0          1          0          0

Switch:1#show ipv6 interface vlan 2

=====
                        Vlan Ipv6 Interface
=====

IF  VLAN PHYSICAL ADMIN OPER TYPE MTU HOP REACH RETRAN MCAST  IPSEC RPC  RPC
INDX                                ABLE  SMIT                                MODE
INDX      ADDRESS  STATE STATE          LMT TIME  TIME    STATUS

-----
2070 2 00:24:7f: enable up  ETHER 1500 64  30000 1000  disable disable exist
      al:7a:06                                     only

=====
                        Vlan Ipv6 Address
=====
IPV6 ADDRESS                                VLAN-ID    TYPE    ORIGIN    STATUS
-----
2001:100:102:202:0:0:0:1                    V-2        UNICAST MANUAL    PREFERRED
fe80:0:0:0:224:7fff:feal:7a06                V-2        UNICAST LINKLAYER PREFERRED

1 out of 2 Total Num of Interface Entries displayed.
2 out of 5 Total Num of Address Entries displayed.

```

OSPF configuration using CLI

Configure Open Shortest Path First (OSPF) so that the switch can use OSPF routing to communicate with other OSPF routers and to participate in OSPF routing.

Configure OSPF Globally

Configure OSPF parameters on the switch so that you can control OSPF behavior on the system. The switch uses global parameters to communicate with other OSPF routers. Globally configure OSPF before you configure OSPF for an interface, port, or VLAN.

Before You Begin

- Ensure that the switch has an IP interface.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf** to commands. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable

configure terminal

router ospf
```

2. Configure the OSPF router ID:

```
router-id {A.B.C.D}
```

3. Configure the router as an autonomous system boundary router (ASBR):

```
as-boundary-router enable
```



Note

Configure the following steps as and when needed.

4. Enable the automatic creation of OSPF virtual links:

```
auto-vlink
```

5. Configure the OSPF default metrics:

```
default-cost {ethernet|fast-ethernet|forty-gig-ethernet|gig-ethernet|
hundred-gig-ethernet|ten-gig-ethernet|twentyfive-gig-ethernet}
<1-65535>]

default-cost vlan <1-65535>]
```

6. Configure the OSPF hold-down timer value:

```
timers basic holddown <3-60>
```

7. Enable the RFC1583 compatibility mode:

```
rfc1583-compatibility enable
```

8. Enable the router to issue OSPF traps:

```
trap enable
```

9. Verify the OSPF configuration:

```
show ip ospf [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

10. Exit OSPF Router Configuration mode:

```
exit
```

You return to Global Configuration mode.

11. Enable OSPF for the switch:

```
router ospf enable
```

Example

Configure the OSPF router ID to 192.0.2.2, enable the automatic creation of OSPF virtual links, and enable traps. Configure the default cost metric for Ethernet to 101,

for fast Ethernet to 110, and for gig-Ethernet, ten-gig-Ethernet, twentyfive-gig-ethernet, forty-gig-Ethernet, and hundred-gig-ethernet to 20, and vlan to 1. Configure the basic holddown to 10. Enable OSPF for the switch, and review the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#router-id 192.0.2.2
Switch:1(config-ospf)#auto-vlink
Switch:1(config-ospf)#default-cost ethernet 101
Switch:1(config-ospf)#default-cost fast-ethernet 110
Switch:1(config-ospf)#default-cost gig-ethernet 20
Switch:1(config-ospf)#default-cost ten-gig-ethernet 20
Switch:1(config-ospf)#default-cost twentyfive-gig-ethernet 20
Switch:1(config-ospf)#default-cost Forty-gig-ethernet 20
Switch:1(config-ospf)#default-cost hundred-gig-ethernet 20
Switch:1(config-ospf)#default-cost vlan 2
Switch:1(config-ospf)#timers basic holddown 10
Switch:1(config-ospf)#trap enable
Switch:1(config-ospf)#exit
Switch:1(config)#router ospf enable
Switch:1(config)#show ip ospf
```

```
=====
                        OSPF General - GlobalRouter
=====

RouterId: 192.0.2.2
AdminStat: disabled
VersionNumber: 2
AreaBdrRtrStatus: false
ASBdrRtrStatus: true
Bad-Lsa-Ignore: false
ExternLsaCount: 0
ExternLsaCksumSum: 0(0x0)
TOSSupport: 0
OriginateNewLsas: 0
RxNewLsas: 0
TrapEnable: false
AutoVirtLinkEnable: false
SpfHoldDownTime: 10
Rfc1583Compatibility: disable
Helper mode: enabled

default-metric :
                ethernet - 101
                fast-ethernet - 110
                gig-ethernet - 20
                ten-gig-ethernet - 20
                twentyfive-gig-ethernet - 20
                forty-gig-ethernet - 20
                hundred-gig-ethernet - 20
                Vlan - 1
```

Variable definitions

The following table defines parameters for the **router-id** command.

Variable	Value
<A.B.C.D>	Configures the OSPF router ID IP address, where A.B.C.D is the IP address.

The following table defines parameters for the **default-cost** command.

Variable	Value
<i>ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>ethernet</i> is for 10 Mb/s Ethernet (default is 100).
<i>fast-ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>fast-ethernet</i> is for 100 Mb/s (Fast) Ethernet (default is 10).
<i>forty-gig-ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>forty-gig-ethernet</i> is for 40 Gigabit Ethernet (default is 1).
<i>gig-ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>gig-ethernet</i> is for Gigabit Ethernet (default is 1).
<i>hundred-gig-ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>hundred-gig-ethernet</i> is for 100 Gigabit Ethernet (default is 1).
<i>ten-gig-ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>ten-gig-ethernet</i> is for 10 Gigabit Ethernet (default is 1).
<i>twentyfive-gig-ethernet</i> <1-65535>	Configures the OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. On a channelized 100 Gbps port, the default-cost for each 25 Gbps channel is 1.
<i>vlan</i> <1-65535>	Configures the OSPF default metrics. <i>vlan</i> is for Vlan interfaces (default is 10).

The following table defines parameters for the **timers basic holddown** command.

Variable	Value
<3-60>	Configures the OSPF hold-down timer value in seconds. The default is 10.

The following table defines parameters for the **show ip ospf** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

Configure OSPF for a Port or VLAN

Configure OSPF parameters on a port or VLAN so you can control OSPF behavior on the port or VLAN.

Before You Begin

- Enable OSPF globally.
- Ensure IP interfaces exist.

About This Task

To configure OSPF on a VRF instance for a port or VLAN, you configure OSPF on the port or VLAN, and then associate the port or VLAN with the VRF.



Important

When you enable OSPF on a VLAN or a port, the switch automatically creates an area 0.0.0.0, and advertises it on the specific VLAN or port, by default. To avoid this behavior, you must manually configure the VLAN or port into a properly configured area on the switch.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the OSPF interface area ID:

```
ip ospf area {A.B.C.D}
```

3. Enable OSPF routing:

```
ip ospf enable
```

4. Choose the OSPF update authentication method:

```
ip ospf authentication-type <message-digest|none|sha-1|sha-2|simple>
```

Both sides of an OSPF connection must use the same authentication type and key.

5. If you choose simple, you must configure the password.

```
ip ospf authentication-key WORD<0-8>
```

6. If you choose an authentication key other than simple such as MD5, Sha-1 or Sha-2, you must configure the digest key first and then assign it to the authentication type.

- a. Create the digest-key:

```
ip ospf digest-key <1-255> key WORD<0-16>
```

- b. Assign the newly created digest key to the authentication type:

```
ip ospf authentication-type <message-digest|none|sha-1|sha-2|simple>
primary-digest-key <1-255>
```

7. Specify the interface type:

```
ip ospf network <broadcast|nbma|passive|p2p>
```

8. Configure the remaining parameters as required, or accept their default values. View the following variable definitions table for more information.

Example

Configure the OSPF interface area ID to 192.0.2.2, enable OSPF routing, choose the OSPF update authentication method as message-digest, and specify the interface type as broadcast.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 1
Switch:1(config-if)#ip ospf area 192.0.2.2
Switch:1(config-if)#ip ospf enable
Switch:1(config-if)#ip ospf authentication-type message-digest
Switch:1(config-if)#ip ospf network broadcast
```

Variable Definitions

The following table defines parameters for the **ip ospf** commands.

Variable	Value
<i>advertise-when-down</i> <i>enable</i>	Enables or disables AdvertiseWhenDown. If enabled, OSPF advertises the network on this interface as up, even if the port is down. The default is disabled. After you configure a port with no link and enable advertise-when-down, OSPF does not advertise the route until the port is active. OSPF advertises the route even when the link is down. To disable advertising based on link status, you must disable this parameter.
<i>area {A.B.C.D}</i>	Configures the OSPF identification number for the area, typically formatted as an IP address.
<i>authentication-key</i> <i>WORD<0-8></i>	Configures the eight-character simple password authentication key for the port or VLAN.

Variable	Value
<i>authentication-type</i> <message-digest none sha-1 sha-2 simple>	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha-1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: SHA-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
<i>bfd</i>	Enable Bidirectional Forwarding Detection (BFD) at the OSPF application level. The default is disabled.
<i>cost</i> <0-65535>	Configures the OSPF cost associated with this interface and advertised in router link advertisements. The default is 0.
<i>dead-interval</i> <0-2147483647>	Configures the router OSPF dead interval, which is the number of seconds the OSPF neighbors of a switch must wait before they assume the OSPF router is down. The default is 40. The value must be at least four times the hello interval.
<i>enable</i>	Enables OSPF on the port or VLAN.
<i>hello-interval</i> <1-65535>	Configures the OSPF hello interval, which is the number of seconds between hello packets sent on this interface. The default is 10.

Variable	Value
<i>message-digest-key</i> <1-255> md5 WORD<0-16>	<p>Configures the MD5 key. You can configure a maximum of two MD5 keys for an interface.</p> <p>If you configure two keys, the interface uses only the first key. To transition to the second key, configure a <i>primary-md5-key</i> to use the ID of the second configured key, and then delete the first key.</p> <p>Important: Use the correct key id when two keys are configured. The key id and md5 password must match with the other OSPF routers, to form the OSPF adjacencies.</p> <p><1-255> is the ID for the MD5 key WORD<0-16> is an alphanumeric password of up to 16 bytes {string length 0-16}</p>
<i>primary-digest-key</i> <1-255>	<p>Use this parameter to transition to a new MD5 key. The new MD5 key changes the primary key used to encrypt outgoing packets.</p> <p><1-255> is the ID for the new MD5 key.</p>
<i>mtu-ignore enable</i>	Enables maximum transmission unit (MTU) ignore. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable mtu-ignore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
<i>network</i> <broadcast nbma passive p2p>	Specifies the type of OSPF interface.
<i>poll-interval</i> <0-2147483647>	Configures the OSPF poll interval in seconds. The default is 120.
<i>priority</i> <0-255>	Configures the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you configure the priority to 0, the port cannot become either the designated router or a backup designated router. The default is 1.
<i>retransmit-interval</i> <0-3600>	Configures the retransmit interval for the virtual interface, which is the number of seconds between link-state advertisement retransmissions.

Variable	Value
<code>transit-delay <0-3600></code>	Configures the transit delay for the virtual interface, which is the estimated number of seconds required to transmit a link-state update over the interface.
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This variable applies only to VLAN interfaces, not to ports.

View OSPF Errors on a Port

Check OSPF errors for administrative and troubleshooting purposes.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display extended information about OSPF errors for the specified port or for all ports:
show ip ospf port-error [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]] [vrf WORD<1-16>] [vrfids WORD<0-512>]

Variable definitions

The following table defines parameters for the **show ip ospf port-error** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code>vrf WORD<1-16></code>	Specifies the VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRFs by ID number.

Configure OSPF Areas on the Router

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or not-so-stubby areas (NSSA).

Before You Begin

- Ensure that the VLAN exists if you configure OSPF on a VLAN.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

About This Task

Place stubby or NSSAs at the edge of an OSPF routing domain. Ensure that you configure all routers in a stubby or NSSA as stubby or NSSA, respectively.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable

configure terminal

router ospf
```
2. Create an OSPF area:

```
area {A.B.C.D}
```
3. Specify the area type:

```
area {A.B.C.D} import <external|noexternal|nssa>
```
4. Configure other OSPF area parameters as required.
5. Ensure that the configuration is correct:

```
show ip ospf area [vrf WORD<1-16>] [vrfsids WORD<0-255>]
```

Example

Create the OSPF area 192.0.2.10, and specify the area type as NSSA. Configure the area support to import summary advertisements into a stub area and configure the import external option for this area as stub. Ensure the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#area 192.0.2.10
Switch:1(config-ospf)#area 192.0.2.10 import nssa
Switch:1(config-ospf)#area 192.0.2.10 import stub
Switch:1(config-ospf)#area 192.0.2.10 import-summaries enable
Switch:1(config-ospf)#show ip ospf area
```

OSPF Area - GlobalRouter								
AREA_ID	STUB_AREA	NSSA	IMPORT_SUM	ACTIVE_IFCNT	NSSA_TRANSLATOR	NSSA_TRANSLATOR	NSSA_TRANSLATOR	NSSA_TRANSLATOR
TRANSLATOR					ROLE	STATE	STABILITY_INTERVAL	EVENTS

0.0.0.0	false	false	true	0	candidate	disabled	40	0
STUB_COST	INTRA_AREA_SPF_RUNS	BDR_RTR_CNT	ASBDR_RTR_CNT	LSA_CNT	LSACK_SUM			
-								
0	0	0	0	0	0			

Variable Definitions

The following table defines parameters for the **area {A.B.C.D}** command.

Variable	Value
<i>default-cost</i> <0-16777215>	Specifies the stub area default metric for this stub area, which is the cost from 0-16777215. This metric value applies at the indicated type of service.
<i>import</i> <external noexternal nssa>	Specifies the type of area: <ul style="list-style-type: none"> external—stub and NSSA are both false noexternal—configures the area as stub area. nssa—configures the area as NSSA.
<i>import-summaries enable</i>	Configures the area support to import summary advertisements into a stub area. Use this variable only if the area is a stub area.
<i>stub</i>	Configures the import external option for this area as stub. A stub area has only one exit point (router interface) from the area.

The following table defines parameters for the **show ip ospf area** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies a VRF.
<i>vrfids</i> WORD<0-512>	Specifies a range of VRF IDs.

View the OSPF Area Information

View the OSPF area information.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF area information:
show ip ospf area [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

View the OSPF area information:

```
Switch:1>enable
Switch:1#show ip ospf area
```

OSPF Area - GlobalRouter									
AREA_ID	STUB_AREA	NSSA	IMPORT_SUM	ACTIVE_IFCNT	NSSA TRANSLATOR ROLE	NSSA TRANSLATOR STATE	NSSA TRANSLATOR STABILITY INTERVAL	NSSA TRANSLATOR EVENTS	

0.0.0.0	false	false	true	0	candidate	disabled	40	0
STUB_COST	INTRA_AREA_SPF_RUNS	BDR_RTR_CNT	ASBDR_RTR_CNT	LSA_CNT	LSACK_SUM			
0	0	0	0	0	0			

Variable definitions

The following table defines parameters for the **show ip ospf area** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

Configure OSPF Aggregate Area Ranges on the Router

Use aggregate area ranges to reduce the number of link-state advertisements required within the area. You can also control advertisements.

Before You Begin

- Enable OSPF globally.
- Ensure that an area exists.
- You configure OSPF area ranges on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:


```
enable

configure terminal

router ospf
```
2. Configure an OSPF area range:


```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
```
3. Configure the advertised metric cost:


```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
advertise-metric <0-65535>
```
4. Configure the advertisement mode:


```
area range {A.B.C.D} {A.B.C.D/X} <summary-link|nssa-extlink>
advertise-mode <summarize|suppress|no-summarize>
```
5. Ensure that the configuration is correct:


```
show ip ospf area-range [vrf WORD<1-16>] [vrfids WORD<0-512>]
```


Example

Configure an OSPF area range to 192.0.2.2, configure the advertised metric cost to 10, and the advertisement mode to summarize.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)# area range 192.0.2.2 255.255.255.0/32 summary-link
Switch:1(config-ospf)# area range 192.0.2.2 255.255.255.0/32 summary-link advertise-
metric 10
Switch:1(config-ospf)# area range 192.0.2.2 255.255.255.0/32 summary-link advertise-mode
summarize
```

Variable definitions

The following table defines parameters for the **area range** command.

Variable	Value
<i>{A.B.C.D}</i> <i>{A.B.C.D/X}</i>	<i>{A.B.C.D}</i> identifies an OSPF area and <i>{A.B.C.D/X}</i> is the IP address and subnet mask of the range, respectively.
<i>advertise-metric</i> <i><0-65535></i>	Changes the advertised metric cost of the OSPF area range.
<i>advertise-mode</i> <i><summarize suppress no-summarize></i>	Changes the advertisement mode of the range.
<i><summary-link nssa-extlink></i>	Specifies the link-state advertisement (LSA) type. If you configure the range as type nssa-extlink, you cannot configure the advertise-metric.

The following table defines parameters for the **show ip ospf area-range** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies a VRF by name.
<i>vrfids WORD<0-512></i>	Specifies a range of VRF IDs.

View the OSPF Area Range Information

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF area range information:
show ip ospf area-range [vrf WORD<1-16>] [vrfids WORD<0-512>]

Variable definitions

The following table defines parameters for the **ip ospf area-range** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

Enable Automatic Virtual Links

Use automatic virtual links to provide an automatic, dynamic backup link for vital OSPF traffic. Automatic virtual links require more system resources than manually configured virtual links.

Before You Begin

- You configure automatic virtual links on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure

- Enter OSPF Router Configuration mode:

```
enable  
  
configure terminal  
  
router ospf
```
- Enable the automatic virtual links feature for the router:

```
auto-vlink
```

Configure an OSPF Area Virtual Interface

Use manual virtual interfaces to provide a backup link for vital OSPF traffic with a minimum of resource use.

Before You Begin

- Enable OSPF globally.
- You configure an OSPF area virtual interface on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

About This Task

Both sides of the OSPF connection must use the same authentication type and key.

You cannot configure a virtual link using a stub area or an NSSA.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create an OSPF area virtual interface:

```
area virtual-link {A.B.C.D} {A.B.C.D}
```

3. Choose the OSPF update authentication method:

```
area virtual-link {A.B.C.D} {A.B.C.D} authentication-type <message-  
digest|none|sha-1|sha-2|simple>
```

Both sides of an OSPF connection must use the same authentication type and key.

4. If required, configure an MD5 key for the virtual interface:

```
area virtual-link message-digest-key {A.B.C.D} {A.B.C.D} <1-255> md5-  
key WORD<1-16>
```

5. Configure optional parameters, as required.

6. Ensure that the configuration is correct:

```
show ip ospf virtual-link {A.B.C.D} {A.B.C.D} [vrf WORD<1-16>] [vrfrids  
WORD<0-512>]
```

Example

Create an OSPF area virtual interface with an area ID of 192.0.2.12 and the virtual interface ID of 198.51.100.2, choose the OSPF update authentication method to simple, and the hello-interval to 100.

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#router ospf  
Switch:1(config-ospf)#area virtual-link 192.0.2.12 198.51.100.2 198.51.100.2  
Switch:1(config-ospf)#area virtual-link 192.0.2.12 198.51.100.2 198.51.100.2  
authentication-type simple  
Switch:1(config-ospf)#area virtual-link 192.0.2.12 198.51.100.2 198.51.100.2 hello-  
interval 100
```

Variable Definitions

The following table defines parameters for the **area virtual-link** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Specifies the area ID and the virtual interface ID.
authentication-key WORD<0-8>	Configures the authentication key of up to eight characters.

Variable	Value
<i>authentication-type</i> <message-digest none sha-1 sha-2 simple>	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha-1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
<i>dead-interval</i> <0-2147483647>	Configures the number of seconds between router hello packets before neighbors declare the router down. This value must be at least four times the hello interval value. The default is 60.
<i>hello-interval</i> <1-65535>	Configures the hello interval, in seconds, on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. The default is 10.
<i>primary-digest-key</i> <1-255>	<p>Use this parameter to transition to a new MD5 key; It changes the primary key used to encrypt outgoing packets.</p> <p><1-255> is the ID for the MD5 key.</p>
<i>retransmit-interval</i> <0-3600>	<p>Configures the retransmit interval for the virtual interface, the number of seconds between LSA retransmissions.</p> <p>The range is from 1–3600.</p>
<i>transit-delay</i> <0-3600>	<p>Configures the transit delay for the virtual interface, the estimated number of seconds required to transmit a link-state update over the interface.</p> <p>The range is from 1–3600.</p>

The following table defines parameters for the **area virtual-link message-digest-key** command.

Variable	Value
<i>{A.B.C.D} {A.B.C.D}</i>	Specifies the area ID and the virtual interface ID.
<i><1-255></i>	Specifies the ID for the message digest key
<i>md5-key WORD<1-16></i>	<p>Configures the MD5 key, you can configure a maximum of two MD5 keys for an interface.</p> <p>If you configure two keys, the interface uses only the first key. To transition to the second key, configure a primary-md5-key to use the ID of the second configured key, and then delete the first key.</p> <p>Important: Use the correct key id when two keys are configured. The key id and md5 password must match with the other OSPF routers, to form the OSPF adjacencies.</p> <p><i>WORD<1-16></i> is an alphanumeric password of up to 16 characters.</p>

The following table defines parameters for the **show ip ospf virtual-link** command.

Variable	Value
<i><A.B.C.D> <A.B.C.D></i>	Specifies the area ID and the virtual interface ID.
<i>vrf WORD<1-16></i>	Specifies a VRF.
<i>vrfids WORD<0-512></i>	Specifies a range of VRF IDs.

Configure an OSPF area on a VLAN or port

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or NSSA. Place stubby or NSSAs at the edge of an OSPF routing domain.

Before You Begin

- Enable OSPF globally.
- Ensure that the VLAN exists.

About This Task

Ensure that you configure all routers in a stubby or NSSA as stubby or NSSA, respectively.

To configure OSPF areas on a VRF instance for a port or VLAN, you configure OSPF on the port or VLAN, and then associate the port or VLAN with the VRF.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create an OSPF area on the VLAN or port:

```
ip ospf area {A.B.C.D}
```

3. Specify the type of network:

```
ip ospf network <broadcast|nbma|passive|p2p>
```

4. Configure other OSPF area parameters as required.

Example

Create an OSPF area 192.0.2.2 on VLAN 1, and specify the type of network as broadcast.

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface vlan 1  
Switch:1(config-if)#ip ospf area 192.0.2.2  
Switch:1(config-if)#ip ospf network broadcast
```

Variable Definitions

The following table defines parameters for the **ip ospf** command.

Variable	Value
{A.B.C.D}	Specifies the area ID.
authentication-key WORD<0-8>	Configures the eight-character simple password authentication key for the port or VLAN.

Variable	Value
<i>authentication-type</i> <message-digest none sha-1 sha-2 simple>	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha-1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
<i>cost</i> <0-65535>	Configures the OSPF cost associated with this interface and advertised in router link advertisements. The default is 0.
<i>dead-interval</i> <0-2147483647>	Configures the number of seconds between router hello packets before neighbors declare the router down. This value must be at least four times the hello interval value. The default is 60.
<i>hello-interval</i> <1-65535>	Configures the hello interval, in seconds, on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. The default is 10.
<i>mtu-ignore enable</i>	Enables MTU ignore. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable mtu-ignore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
<i>network</i> <broadcast nbma passive p2p>	Specifies the type of OSPF interface.
<i>poll-interval</i> <0-2147483647>	Configures the OSPF poll interval in seconds. The default is 120.
<i>primary-digest-key</i> <1-255>	<p>Use this parameter to transition to a new MD5 key; it changes the primary key used to encrypt outgoing packets.</p> <p><1-255> is the ID for the message digest key.</p>

Variable	Value
<code>priority <0-255></code>	Configures the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you set the priority to 0, the port cannot become either the designated router or a backup designated router. The default is 1.
<code>retransmit-interval <0-3600></code>	Configures the retransmit interval: the number of seconds between LSA retransmissions.
<code>transit-delay <0-3600></code>	Configures the transit delay: the estimated number of seconds it takes to transmit a link-state update over the interface.

Configure an OSPF Host Route

Configure host routes when the switch resides in a network that uses routing protocols other than OSPF. A host route is a more-specific route and is used even if it is higher cost than a network route.

Before You Begin

- Globally enable OSPF.
- You configure an OSPF host route on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

About This Task

Use a host route to create a custom route to a specific host to control network traffic.

You can specify which hosts directly attach to the router, and the metrics and types of service to advertise for the hosts.

Procedure

1. Enter OSPF Router Configuration mode:


```
enable
configure terminal
router ospf
```
2. Create a host route:


```
host-route {A.B.C.D} [metric <0-65535>]
```
3. Ensure that the configuration is correct:


```
show ip ospf host-route [vrf WORD<1-16>] [vrfs WORD<0-512>]
```


Example

Create a host route on IP address 192.0.2.20 with a metric of 20.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#host-route 192.0.2.20 metric 20
Switch:1(config-ospf)#show ip ospf host-route
Switch:1(config-ospf)#show ip ospf host-route
```

```
=====
                        OSPF Host Route - GlobalRouter
=====
HOSTIPADDR      TOS  METRIC
-----
192.0.2.20      -    20
```

Variable definitions

The following table defines parameters for the **host-route** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the IP address of the host router in a.b.c.d format.
<i>metric <0-65535></i>	Configures the metric (cost) for the host route.

The following table defines parameters for the **show ip ospf host-route** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies a VRF by name.
<i>vrfids WORD<0-512></i>	Specifies a range of VRF IDs.

Configure OSPF NBMA Neighbors

Configure NBMA neighbors so that the interface can participate in designated router election. All OSPF neighbors that you manually configure are NBMA neighbors.

Before You Begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Ensure that the interface is NBMA.
- You configure OSPF NBMA neighbors on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip ospf**. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```
2. Create an NBMA OSPF neighbor:

```
neighbor {A.B.C.D} priority <0-255>
```
3. Ensure that the configuration is correct:

```
show ip ospf neighbor [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Create an NBMA OSPF neighbor.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#neighbor 198.51.100.2 priority 10
```

Variable definitions

The following table defines parameters for the **neighbor** command.

Variable	Value
<i>{A.B.C.D}</i>	Identifies an OSPF area in IP address format a.b.c.d.
<i>priority <0-255></i>	Changes the priority level of the neighbor.

The following table defines parameters for the **show ip ospf neighbors** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies a VRF by name.
<i>vrfids WORD<0-512></i>	Specifies a range of VRF IDs.

Disable Helper Mode for OSPFv2

About This Task

By default, OSPF Helper mode is enabled when OSPF is configured.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Disable Helper mode:
`helper-mode-disable`

Apply OSPF Route Acceptance Policies

Use a route policy to define how the switch redistributes external routes from a specified source into an OSPF domain. The policy defines which route types the switch accepts and redistributes.

Before You Begin

- Enable OSPF globally.
- Ensure that a route policy exists.
- Ensure that the area exists.
- You apply OSPF route acceptance policies on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix `ip ospf`. The VRF must have an RP Trigger of OSPF. Not all parameters are configurable on non0 VRFs.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter OSPF Router Configuration mode:
`enable`

`configure terminal`

`router ospf`
2. Create an acceptance policy instance:
`accept adv-rtr {A.B.C.D}`
3. Configure the type of metric to accept:
`accept adv-rtr {A.B.C.D} metric-type <type1|type2|any>`
4. Indicate the route policy:
`accept adv-rtr {A.B.C.D} route-map WORD<0-64>`

5. Enable a configured OSPF route acceptance instance:

```
accept adv-rtr {A.B.C.D} enable
```

6. Ensure that the configuration is correct:

```
show ip ospf accept [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Create an acceptance policy instance, configure the type of metric to accept, indicate the route policy and enable the OSPF route acceptance instance. Ensure the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#router ospf
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11 metric-type type1
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11 route-map test1
Switch:1(config-ospf)#accept adv-rtr 192.0.2.11 enable
Switch:1>show ip ospf accept
=====
Ospf Accept - GlobalRouter
=====
ADV_RTR      MET_TYPE ENABLE POLICY
-----
192.0.2.11   type1    true  test1
```

Variable definitions

The following table defines parameters for the **accept adv-rtr** command.

Variable	Value
<A.B.C.D>	Specifies the IP address.
<i>enable</i>	Enables an OSPF acceptance policy.
<i>metric-type</i> <type1 type2 any>	Configures the metric type as type 1, type 2, or any.
<i>route-map</i> WORD<0-64>	Configures the route policy by name.

The following table defines parameters for the **ip ospf accept** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies a VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies a range of VRF IDs.

View the OSPF Configuration Information

View the OSPF configuration information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF configuration information:

```
show ip ospf accept [vrf WORD<1-16>] [vrfrids WORD<0-512>]
```

Example

```
Switch:1>show ip ospf accept
=====
Ospf Accept - GlobalRouter
=====
ADV_RTR      MET_TYPE ENABLE POLICY
-----
192.0.2.11   type1    true    test1
```

Variable definitions

The following table defines parameters for the **show ip ospf accept** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies a VRF by name.
<i>vrfrids WORD<0-512></i>	Specifies a range of VRF IDs.

View the OSPF Link-State Database

View the area advertisements and other information in the Link-State Database (LSDB) to ensure correct OSPF operations.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF link-state database:

```
show ip ospf lsdb [adv_rtr {A.B.C.D}] [area {A.B.C.D}>] [lsa-type
<1-11] [lsid {A.B.C.D}] [vrf WORD<1-16>] [vrfrids WORD<0-512>] [detail]
```

Example

```
Switch(config-ospf)#show ip ospf lsdb
=====
OSPF LSDB - GlobalRouter
=====

Router Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID      ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
Router      192.0.2.0         192.0.2.0        617  0x80000031  0xeafd
Router      198.51.100.0      198.51.100.0     1033 0x80000030  0xa5f2

Network Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID      ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
Network     100.1.1.2         192.0.2.0        617  0x8000002f  0xd038

Summary Lsas in Area 0.0.0.0

LSTYPE      LINKSTATEID      ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
```

AsSummary Lsas in Area 0.0.0.0							
LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM		

NSSA Lsas in Area 0.0.0.0							
LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM		

=====							
AsExternal Lsas							
=====							
LSTYPE	LINKSTATEID	ADV_ROUTER	ETYPE	METRIC	ASE_FWD_ADDR	AGE	SEQ_NBR
CSUM							

Variable definitions

The following table defines parameters for the **show ip ospf lsdb** command.

Variable	Value
<i>adv_rtr</i> {A.B.C.D}	Specifies the advertising router.
<i>area</i> {A.B.C.D}	Specifies the OSPF area.
<i>detail</i>	Provides detailed output.
<i>lsa-type</i> <1-11>	Specifies the link-state advertisement type in the range of 1-11.
<i>lsid</i> {A.B.C.D}	Specifies the link-state ID.
<i>vrf</i> WORD<1-16>	Specifies a VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies a range of VRF IDs.

View the OSPF External Link-State Database

View the LSDB to determine externally learned routing information. The system displays the information for all metric types or for the type you specify.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF autonomous system external (ASE) link-state advertisements:
show ip ospf ase [metric-type <1-2>] [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

```
Switch:1#show ip ospf ase

=====
                        OSPF AsExternal Lsas - GlobalRouter
=====

LSTYPE      LINKSTATEID      ADV_ROUTER      ETYPE  METRIC  ASE_FWD_ADDR      AGE  SEQ_NBR
CSUM
=====
```

Variable definitions

The following table defines parameters for the **show ip ospf ase** command.

Variable	Value
<i>metric-type</i> <1-2>	Specifies the metric type.
<i>vrf</i> WORD<1-16>	Identifies the VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies a VRF by ID.

Configure Route Redistribution to OSPF

Before You Begin

- Enable OSPF globally.
- Ensure that a route policy exists.
- Ensure that you configure OSPF as the boundary router.

About This Task

Configure a redistribute entry to announce certain routes into the OSPF domain, including DvR host routes, static routes, direct routes, Routing Information Protocol (RIP) routes, OSPF routes, IS-IS routes or Border Gateway Protocol (BGP) routes. Optionally, use a route policy to control the redistribution of routes.



Note

Route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create the redistribution instance:

```
redistribute <bgp|direct|isis|ospf|rip|static|dvr> [vrf-src  
WORD<1-16>]
```

3. Apply a route policy if required:

```
redistribute <bgp|direct|isis|ospf|rip|static|dvr> route-map  
WORD<0-64> [vrf-src WORD<1-16>]
```

4. (Optional) Enable route filtering using tags:

```
redistribute isis auto-filter-tag
```

5. Configure other parameters, as required.

6. Enable the redistribution.

```
redistribute <bgp|direct|isis|ospf|rip|static|dvr> enable [vrf-src  
WORD<1-16>]
```

7. Ensure that the configuration is correct:

```
show ip ospf redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

8. Exit to Global Configuration mode:

```
exit
```

9. Apply the redistribution.

```
ip ospf apply redistribute <bgp|direct|isis|ospf|rip|static|dvr> [vrf  
WORD<1-16>] [vrf-src WORD<1-16>]
```

Changes do not take effect until you apply them.

10. View all routes (including DvR host routes) that are redistributed into OSPF:

- a. View the routes that are redistributed from the GRT to OSPF:

```
show ip ospf lsdb
```

- b. View the routes that are redistributed to OSPF for a specific VRF instance:

```
show ip ospf lsdb [vrf WORD<1-64>] [vrfids WORD<0-512>]
```

Examples

Redistribute static routes from the GRT to OSPF. Create the redistribution instance, apply a route policy, enable redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#redistribute static

WARNING: Routes will not be injected until apply command is issued after enable command
Switch:1(config-ospf)#redistribute static route-map policy1
Switch:1(config-ospf)#redistribute static enable
Switch:1(config-ospf)#exit
Switch:1(config)#ip ospf apply redistribute static
```



```
Switch:1(config)#show ip ospf redistribute
```

```
=====
                        OSPF Redistribute List - GlobalRouter
=====
```

SRC-VRF	SRC	MET	MTYPE	SUBNET	ENABLE	RPOLICY	AUTO-FILTER
GlobalRouter	STAT	0	type2	allow	TRUE	policy1	
GlobalRouter	DVR	0	type2	allow	FALSE		

Redistribute DvR host routes from the GRT to OSPF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#redistribute dvr
Switch:1(config-ospf)#redistribute dvr enable
Switch:1(config-ospf)#exit
Switch:1(config)#ip ospf apply redistribute dvr
Switch:1(config)#show ip ospf redistribute
```

```
=====
                        OSPF Redistribute List - GlobalRouter
=====
```

SRC-VRF	SRC	MET	MTYPE	SUBNET	ENABLE	RPOLICY	AUTO-FILTER
GlobalRouter	DVR	0	type2	allow	TRUE	test1	

Redistribute DvR host routes to OSPF for a specific VRF vrf1:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrf1
Switch:1(router-vrf)#ip ospf redistribute dvr
Switch:1(router-vrf)#ip ospf redistribute dvr enable
Switch:1(router-vrf)#exit
Switch:1(config)#ip ospf apply redistribute dvr vrf vrf1
```

View the DvR host routes that are distributed into OSPF:

```
Switch:1(config)#show ip ospf lsdb vrf vrf1
```

```
=====
                        OSPF LSDB - VRF vrf1
=====
```

```
Router Lsas in Area 0.0.0.0
```

LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM
Router	192.0.2.1	192.0.2.1	1603	0x8000002d	0xf226
Router	203.0.113.1	203.0.113.1	1608	0x8000002a	0x4104

```
Network Lsas in Area 0.0.0.0
```

LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM
Network	14.1.1.11	192.0.2.1	1635	0x80000003	0x4909

Summary Lsas in Area 0.0.0.0								
LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM			

AsSummary Lsas in Area 0.0.0.0								
LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM			

NSSA Lsas in Area 0.0.0.0								
LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM			

Opaque-Loc Lsas in Area 0.0.0.0								
LSTYPE	LINKSTATEID	ADV_ROUTER	AGE	SEQ_NBR	CSUM			

=====								
AsExternal Lsas								
=====								
LSTYPE	LINKSTATEID	ADV_ROUTER	ETYPE	METRIC	ASE_FWD_ADDR	AGE	SEQ_NBR	CSUM

AsExternal	101.1.1.3	203.0.113.1	2	1	0.0.0.0	1563	0x80000003	0xe7a7
AsExternal	101.1.1.4	203.0.113.1	2	1	0.0.0.0	1477	0x80000003	0xddb0
AsExternal	102.1.1.3	203.0.113.1	2	1	0.0.0.0	1528	0x80000003	0xdab3
AsExternal	102.1.1.4	203.0.113.1	2	1	0.0.0.0	1480	0x80000003	0xd0bc
AsExternal	103.1.1.3	203.0.113.1	2	1	0.0.0.0	1531	0x80000003	0xcdbf
...								
...								
...								

Variable Definitions

The following table defines parameters for the **redistribute** command.

Variable	Value
<i>auto-filter-tag</i>	Enables route filtering to ensure tagged routes are not redistributed into OSPF. The default is disabled.
<i>enable</i>	Enables the OSPF route redistribution instance.
<i>metric</i> <0-65535>	Configures the metric to apply to redistributed routes.
<i>metric-type</i> <type1 type2>	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.

Variable	Value
<code>route-map WORD<0-64></code>	Configures the route policy to apply to redistributed routes.
<code>subnets <allow suppress></code>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
<code>vrf-src WORD<1-16></code>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
<code><bgp direct isis ospf rip static dvr></code>	Specifies the type of routes to redistribute (the protocol source).

The following table defines parameters for the **ip ospf apply redistribute** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the VRF instance.
<code>vrf-src WORD<1-16></code>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<code>WORD<0-32></code>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

View the OSPF Redistribution Configuration Information

Displays the OSPF redistribution configuration information.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter Privileged EXEC mode:
enable

2. View the OSPF redistribution configuration information:

```
show ip ospf redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1#show ip ospf redistribute
```

```
=====
                        OSPF Redistribute List - GlobalRouter
=====

SRC-VRF          SRC  MET  MTYPE  SUBNET  ENABLE  RPOLICY  AUTO-FILTER
-----
GlobalRouter     STAT  0   type2  allow   TRUE
```

Variable definitions

The following table defines parameters for the **show ip ospf redistribute** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

Configure InterVRF Route Redistribution for OSPF

Use route redistribution so that a VRF interface can announce routes learned by other protocols, for example, OSPF or BGP. The switch supports interVRF route redistribution. Use a route policy to control the redistribution of routes.

You can also redistribute inter-VRF DvR routes to OSPF.

Before You Begin

- Enable OSPF globally.
- Ensure that a route policy exists.

- Ensure that the VRFs exist.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the redistribution instance:

```
ip ospf redistribute <bgp|direct|isis|ospf|rip|static|dvr>
```

3. Apply a route policy if required:

```
ip ospf redistribute <bgp|direct|isis|ospf|rip|static|dvr> route-map  
WORD<0-64> [vrf-src WORD<1-16>]
```

4. Configure other parameters, as required.

5. Enable the redistribution:

```
ip ospf redistribute <bgp|direct|isis|ospf|rip|static|dvr> enable  
[vrf-src WORD<1-16>]
```

6. Ensure that the configuration is correct:

```
show ip ospf redistribute [vrf WORD<1-16>] [vrffids WORD<0-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution:

```
ip ospf apply redistribute <bgp|direct|isis|ospf|rip|static|dvr> [vrf  
WORD<1-16>] [vrf-src WORD<1-16>]
```

Example

Example 1:

Create the redistribution instance, apply a route policy, enable the redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf red
Switch:1(router-vrf)#ip ospf redistribute isis
Switch:1(router-vrf)#ip ospf redistribute isis route-map test2
Switch:1(router-vrf)#ip ospf redistribute isis enable
Switch:1(router-vrf)#exit
Switch:1(config)#ip ospf apply redistribute isis
```

Example 2:

This example demonstrates redistribution of inter-VRF routes (both direct and DvR routes) to OSPF, with a route policy configured.

Redistribute inter-VRF DvR routes between VRFs (with VRF IDs 10 and 30), to OSPF.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf 10
Switch:1(router-vrf)#ip prefix-list "test10" 192.0.2.0/24 ge 25 le 32
Switch:1(router-vrf)#route-map "test10" 1
Switch:1(router-vrf)#permit
Switch:1(router-vrf)#enable
Switch:1(router-vrf)#match network "test10"
Switch:1(router-vrf)#set metric 99
Switch:1(router-vrf)#exit

Switch:1(config)#router vrf 30
Switch:1(router-vrf)#ip ospf redistribute direct vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute direct enable vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute dvr vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute dvr route-map "test10" vrf-src 10
Switch:1(router-vrf)#ip ospf redistribute dvr enable vrf-src 10
Switch:1(router-vrf)#exit

Switch:1(config)#ip ospf apply redistribute direct vrf 30 vrf-src 10
Switch:1(config)#ip ospf apply redistribute dvr vrf 30 vrf-src 10
```

Variable definitions

The following table defines parameters for the **ip ospf redistribute** command.

Variable	Value
<i>enable</i>	Enables the OSPF route redistribution instance.
<i>metric <0-65535></i>	Configures the metric to apply to redistributed routes.
<i>metric-type <type1 type2></i>	Specifies a metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.

Variable	Value
<code>route-map WORD<0-64></code>	Configures the route policy to apply to redistributed routes.
<code>subnets <allow suppress></code>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
<code>vrf-src WORD<1-16></code>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
<code><bgp direct isis ospf rip static dvr></code>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, static or dvr.

The following table defines parameters for the **ip ospf apply redistribute** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the VRF instance.
<code>vrf-src WORD<1-16></code>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<code>WORD<0-32></code>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Force Shortest-Path Calculation Updates

Force the switch to update its shortest-path calculations so that the switch uses the latest OSPF routing information. Manually initiate a shortest path first (SPF) run, or calculation, to immediately update the OSPF LSDB. This action is useful in the following circumstances:

- when you need to immediately restore a deleted OSPF-learned route
- when the routing table entries and the LSDB do not synchronize

Before You Begin

- You can perform this procedure in one of the following CLI modes: User EXEC, Privileged EXEC, or Global Configuration.

About This Task

This process is computationally intensive. Use this command only if required.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Force the router to update its shortest-path calculations:

```
ip ospf spf-run [vrf WORD<1-16>]
```

Example

Force the router to update its shortest-path calculations:

```
Switch:1>ip ospf spf-run
```

Variable definitions

The following table defines parameters for the **ip ospf spf-run** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF instance by name.

View the OSPF Default Cost Information

View the OSPF default cost information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF cost information:
show ip ospf default-cost [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

View the OSPF cost information on the switch:

```
Switch:1#show ip ospf default-cost vrf 3
```

```
=====
                        OSPF Default Metric - VRF 3
=====
10MbpsPortDefaultMetric: 100
100MbpsPortDefaultMetric: 10
1000MbpsPortDefaultMetric: 1
10000MbpsPortDefaultMetric: 1
25000MbpsPortDefaultMetric: 1
40000MbpsPortDefaultMetric: 1
100000MbpsPortDefaultMetric: 1
```

Variable definitions

The following table defines parameters for the **show ip ospf default-cost** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

View the OSPF Timer Information

Display OSPF timers information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF timers information:
show ip ospf int-timers [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

```
Switch:1#show ip ospf int-timers
```

```
=====
                        OSPF Interface Timer - GlobalRouter
=====
```

INTERFACE	AREAID	TRANSIT DELAY	RETRANS INTERVAL	HELLO INTERVAL	DEAD INTERVAL	POLL INTERVAL
192.0.2.1	0.0.0.0	1	5	10	40	120
192.0.2.11	0.0.0.0	1	5	10	40	120
192.0.2.3	0.0.0.0	1	5	10	40	120

```
=====
                        Ospf Virtual Interface Timer
=====
```

AREAID	NBRIPADDR	TRANSIT DELAY	RETRANS INTERVAL	HELLO INTERVAL	DEAD INTERVAL
--------	-----------	------------------	---------------------	-------------------	------------------

Variable definitions

The following table defines parameters for the **show ip ospf int-timers** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.

View the OSPF Neighbor Information

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF neighbor information:
show ip ospf neighbor [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

View OSPF neighbor information:

```
Switch:1#show ip ospf neighbor
```

```
=====
                        OSPF Neighbors - GlobalRouter
=====
```

INTERFACE	NBRROUTERID	NBRIPADDR	PRI	STATE	RTXQLEN	PERM	TTL	HELPER STATUS	HELPER AGE	HELPER EXIT	ACCESS REASON
192.0.2.5	192.0.2.1	192.0.2.6	1	Full	0	Dyn	31		0	None	GlobalRouter
198.51.100.7	198.51.100.1	198.51.100.8	128	Restart	0	Dyn	147	H	0	None	GlobalRouter

Total ospf neighbors: 2

Variable definitions

The following table defines parameters for the **show ip ospf neighbor** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

View the OSPF Authentication Information

Display OSPF authentication information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF authentication information:
show ip ospf int-auth [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

View the OSPF authentication information:

```
Switch:1#show ip ospf int-auth

=====
                        OSPF Interface AuthKey - GlobalRouter
=====
INTERFACE      AUTHTYPE  AUTHKEY
-----
192.0.2.1       none
192.0.2.11      none
192.0.2.3       none
```

Variable definitions

The following table defines parameters for the **show ip ospf int-auth** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a VRF by name.
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.

View the OSPF Interface Statistics

Use statistics to help you monitor OSPF performance.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF interface statistics:

```
show ip ospf ifstats [detail vrf WORD<1-16> vrfids WORD<0-512>]
[mismatch vrf WORD<1-16> vrfids WORD<0-512>] [vlan <1-4059>] [vrf
WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View the OSPF interface statistics:

```
Switch:1>show ip ospf ifstats
=====
                        OSPF Interface Statistics - GlobalRouter
=====
---HELLOS--- ---DBS--- -LS REQ-- --LS UPD--- --LS ACK---
INTERFACE      RX      TX      RX      TX      RX      TX      RX      TX      RX      TX
-----
192.0.2.3      76035  76355  33     32     4       9     2483  2551  2525  1247
192.0.2.8      76038  76349  0       0       0       0       0     0     0     0
```

Variable definitions

The following table defines parameters for the **show ip ospf ifstats** command.

Variable	Value
<i><1-4059></i>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>detail</i>	Displays the details of the OSPF.
<i>mismatch</i>	Specifies the number of times the area ID is not matched.
<i>vrf WORD<1-16></i>	Specifies a VRF by name.
<i>vrfids WORD<0-512></i>	Specifies a range of VRF IDs.

View OSPF Range Statistics

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures. OSPF range statistics include area ID, range network address, range subnet mask, range flag, and LSDB type.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display the OSPF range statistics:

```
show ip ospf stats [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>show ip ospf stats
```

```
=====
                        OSPF Statistics - GlobalRouter
=====
      NumBufAlloc: 1138
      NumBufFree: 1138
NumBufAllocFail: 0
      NumBufFreeFail: 0
          NumTxPkt: 1144
          NumRxPkt: 2287
      NumTxDropPkt: 0
      NumRxDropPkt: 0
      NumRxBadPkt: 0
          NumSpfRun: 19
      LastSpfRun: 0 day(s), 00:26:15
      LsdbTblSize: 7
      NumAllocBdDDP: 5
      NumFreeBdDDP: 5
      NumBadLsReq: 0
      NumSeqMismatch: 0
      NumOspfRoutes: 7
      NumOspfAreas: 0
NumOspfAdjacencies: 3
          NumOspfNbrs: 3
NumEnabledOspfAreas:0
```

Variable Definitions

Use the data in the following table to use the **show ip ospf stats** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies a VRF instance by VRF name.
<i>vrfids WORD<0-16></i>	Specifies a VRF or range of VRFs by ID.

Clear IP OSPF Statistics

Use the following procedure to clear all IPv4 OSPF statistics.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear IPv4 OSPF statistics:

```
clear ip ospf stats [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Variable Definitions

Use the data in the following table to use the **clear ip ospf stats** command.

Variable	Value
vrf WORD<1-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-512>	Specifies the ID of the VRF.

View OSPF Error Statistics on a Port

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display extended information about OSPF errors for the specified port or for all ports:
show interfaces GigabitEthernet error ospf [{slot/port[-slot/port]}
[,...][slot/all][all]]

Variable Definitions

The following table defines parameters for the **show interfaces GigabitEthernet error ospf** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View Basic OSPF Statistics for a Brouter Port

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

About This Task

This command provides values for brouter port implementations. For IP protocol-based VLAN interface statistics, use [View the OSPF Interface Statistics](#) on page 2820.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View basic OSPF statistics:

```
show ports statistics ospf main [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

Example

View basic OSPF statistics:

```
Switch:1>enable
Switch:1#show ports statistics ospf main
```

```
=====
                        Port Stats Ospf
=====
```

PORT_NUM	RX_HELLO	TX_HELLO	RXDB_DESCR	TXDB_DESCR	RXLS_UPDATE	TXLS_UPDATE
1/3	0	0	0	0	0	0

Variable Definitions

Use the data in the following table to use the **show ports statistics ospf main** command.

Variable	Value
<code>{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View Extended OSPF Statistics for a Brouter Port

Use statistics to help you monitor Open Shortest Path First (OSPF) performance. You can also use statistics in troubleshooting procedures.

About This Task

This command provides values for brouter port implementations. For IP protocol-based VLAN interface statistics, use [View the OSPF Interface Statistics](#) on page 2820.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display extended OSPF information about the specified port or for all ports:

```
show ports statistics ospf extended [{slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]]
```

Example

Display extended OSPF information:

```
Switch:1>enable
Switch:1#show ports statistics ospf extended

=====
                        Port Stats Ospf Extended
=====
PORT_NUM  RXLS_REQS  TXLS_REQS  RXLS_ACKS  TXLS_ACKS
-----
1/3        0            0          0          0
```

Variable Definitions

Use the data in the following table to use the **show ports statistics ospf extended** command.

Variable	Value
<code>{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View the OSPF Virtual Link Information

Displays the OSPF virtual link information to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the OSPF virtual link information:

```
show ip ospf virtual-link {A.B.C.D} {A.B.C.D} [vrf WORD<1-16>] [vrfs
WORD<0-512>]
```

Example

View the OSPF virtual link information:

```
Switch:1#show ip ospf virtual-link

=====
                        OSPF Interface AuthKey - GlobalRouter
=====
INTERFACE          AUTHTYPE AUTHKEY
-----
192.0.2.11          none
```

Variable definitions

The following table defines parameters for the **show ip ospf virtual-link** command.

Variable	Value
<i>{A.B.C.D}</i> <i>{A.B.C.D}</i> <i>vrf WORD<1-16></i>	Specifies the area ID and the virtual interface ID. The first IP address specifies the area ID and the second specifies the virtual interface ID.
<i>{A.B.C.D}</i> <i>{A.B.C.D}</i> <i>vrfids WORD<0-512></i>	Displays OSPF configuration for a particular VRF. Specifies a VRF by name.
<i>{A.B.C.D}</i> <i>{A.B.C.D}</i>	Specifies a range of VRF IDs.

View the OSPF Configuration for a VRF

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the configuration by VRF name:
show ip ospf vrf *WORD<1-16>*
3. View the configuration by VRF ID:
show ip ospf vrfids *WORD<0-512>*

Example

View the configuration for VRF named virtualrandf1:

```
Switch:1#show ip ospf vrf virtualrandf1

=====
                        OSPF General - VRF virtualrandf1
=====

RouterId: 192.0.2.1
AdminStat: enabled
VersionNumber: 2
AreaBdrRtrStatus: false
ASBdrRtrStatus: true
Bad-Lsa-Ignore: false
```



```

    ExternLsaCount: 0
    ExternLsaChecksumSum: 0 (0x0)
    TOSupport: 0
    OriginateNewLsas: 0
    RxNewLsas: 0
    TrapEnable: false
    AutoVirtLinkEnable: false
    SpfHoldDownTime: 10
    Rfc1583Compatibility: disable
--More-- (q = quit)

```

Variable Definitions

The following table defines parameters for the **show ip ospf vrf** command.

Variable	Value
<code>vrfids WORD<0-512></code>	Specifies a range of VRF IDs.
<code>vrf WORD<1-16></code>	Specifies a VRF by name.

OSPFv3 Configuration using CLI

Use the procedures in this section to configure OSPFv3 using CLI.

Configure OSPF Globally

Configure OSPFv3 globally to enable it on the system and to configure the router ID.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable OSPFv3 for IPv6:

```
router ospf ipv6-enable
```

The default is disabled.

3. Log on to OSPF Router Configuration mode:

```
router ospf
```

4. Specify the router ID:

```
ipv6 router-id {A.B.C.D}
```

- Optionally, make the router an autonomous system (AS) boundary router (BR):

```
ipv6 as-boundary-router enable
```

Enable the ASBR if the router attaches at the edge of the OSPF network, and has one or more interfaces that run an interdomain routing protocol. The default is disabled.

Example

Enable OSPFv3 for IPv6:

```
Switch:1(config)#router ospf ipv6-enable
```

Log on to OSPF Router Configuration mode:

```
Switch:1(config)#router ospf
```

Specify the router ID:

```
Switch:1(config-ospf)#ipv6 router-id 1.1.1.1
```

Variable definitions

Use the data in the following table to use the **ipv6 router-id** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies a 32-bit integer that identifies the router in the autonomous system. This value must be unique. The default value will be one of the IPv4 interface addresses.

Create an OSPF Area

Create an area to subdivide the autonomous system (AS) into areas that group contiguous networks, routers that connect to these networks, and attached hosts. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic compared to treating the entire AS like a single link-state domain.

Before You Begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace *ipv6* with *ipv6 ospf*.

About This Task

A stub area does not receive advertisements for external routes, which reduces the size of the link-state database (LSDB). A stub area uses only one area border router (ABR). Any packets destined for outside the area are routed to the area border exit point, examined by the ABR, and forwarded to a destination.

A not so stubby area (NSSA) prevents the flooding of AS-External link-state advertisements into the area by replacing them with a default route. NSSAs also import small stub (non- OSPF) routing domains into OSPF.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable

configure terminal

router ospf
```

2. Specify the area ID:

```
ipv6 area {A.B.C.D}
```

3. Configure optional area parameters if the default values do not meet your requirements:

- a. Configure the area type if you need a stub or NSSA area:

```
ipv6 area {A.B.C.D} type <nssa|stub>
```

By default, the area is a normal area; neither a stub nor NSSA area.

- b. Configure the default cost:

```
ipv6 area {A.B.C.D} default-cost <0-16777215>
```

You do not need to configure this parameter if the area is a normal area.

- c. Configure the area support for importing advertisements:

```
ipv6 area {A.B.C.D} import <external|noexternal|nssa>
```

The default is external.

- d. Disable the importation of summary advertisements into a stub area:

```
no ipv6 area {A.B.C.D} import-summaries enable
```

The default is enabled.

- e. Configure translation of Type 7 LSAs into Type 5 LSAs:

```
ipv6 area {A.B.C.D} translator-role <1-2>
```

The default value is 2-candidate.

Example

Specify the area ID:

```
Switch:1(config-ospf)#ipv6 area 0.0.0.1
```

Variable definitions

Use the data in the following table to use the **ipv6 area** command.

Variable	Value
{A.B.C.D}	Specifies a 32-bit integer to uniquely identify an area. Use 0.0.0.0 for the OSPFv3 backbone.
default-cost <0-16777215>	Configures the metric value advertised for the default route to stub and NSSA areas.

Variable	Value
<code>import <external noexternal nssa></code>	Configures area support for importing AS-external LSAs: <ul style="list-style-type: none"> external—normal area noexternal—stub area nssa—not-so-stubby area AS-scope LSAs are not imported into stub areas or NSSAs. NSSAs import AS-External data at Type 7 LSAs, which use area scope. The default is external.
<code>import-summaries enable</code>	Controls the import of inter-area LSAs into a stub area. If you disable this parameter, the router does not originate nor propagate inter-area LSAs into the stub area. If you enable this parameter (the default), the router both summarizes and propagates inter-area LSAs.
<code><nssa stub></code>	Configures the type of area. By default, the area is neither a stub area or an NSSA.
<code>translator-role <1-2></code>	Indicates if the NSSA border router can perform NSSA translation of Type 7 LSAs to Type 6 LSAs. The possible values are always (1) or candidate (2). The default is candidate (2).

Create OSPF Area Ranges

Create an area address range on the OSPF router to reduce the number of area border router (ABR) advertisements into other OSPF areas. An area address range is an implied contiguous range of area network addresses for which the ABR advertises a single summary route.

Before You Begin

- You must create the OSPF area.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace `ipv6` with `ipv6 ospf`.

About This Task

If you create two ranges, and one range is a subset of the other, the router uses the most specific match.

Procedure

- Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

- Create an area range:

```
ipv6 area range {A.B.C.D} WORD<0-255> [inter-area-prefix-link|nssa-extlink] advertise-mode <advertise|not-advertise> [advertise-metric <0-65535>]
```

Example

Create an area range:

```
Switch:1(config-ospf)#ipv6 area range 0.0.0.1 3000::0/16 advertise-mode advertise
```

Variable definitions

Use the data in the following table to use the **ipv6 area range** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the area in which the address aggregate exists. Use dotted decimal notation to specify the area name.
<i>advertise-metric</i> <i><0-65535></i>	Specifies a cost value to advertise for the OSPF area range. This value applies to summary LSAs (Type 3). If the value is 0, OSPF uses the cost to the farthest point in the network that is summarized.
<i>advertise-mode</i> <i><advertise not-advertise></i>	Specifies the advertisement mode for prefixes in the range. <i>advertise</i> advertises the aggregate summary LSA with the same link-state ID. <i>not-advertise</i> does not advertise networks that fall within the range. The default is <i>advertise</i> .
<i><inter-area-prefix-link nssa-extlink></i>	Specifies the area LSDB type to which the address aggregate applies. <i>inter-area-prefix-link</i> generates an aggregated summary. <i>nssa-extlink</i> generates an NSSA link summary.
<i>WORD<0-255></i>	Specifies the IPv6 address and prefix.

Create an OSPF Virtual Link

Create a virtual link if the switch does not connect directly to the backbone. The switch can create automatic virtual links or you can perform this procedure to create virtual links manually. Manual virtual links conserve resources and provide specific control over virtual link placement in your OSPF configuration.

Before You Begin

- The router must be an ABR to create a virtual router interface.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace *ipv6* with *ipv6 ospf*.

About This Task

Virtual linking is similar to backup redundancy. The switch creates a virtual link for vital traffic paths in your OSPF configuration if traffic is interrupted, such as when an interface cable that provides a connection to the backbone (either directly or indirectly) is disconnected from the switch. Automatic virtual linking ensures that a link is created by using another switch.

OSPF routes cannot be learned through an ABR unless it connects to the backbone directly or through a virtual link.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create a virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D}
```

3. Configure optional parameters for the virtual link if the default values do not meet your requirements:

- a. Configure the router dead interval:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} dead-interval <1-65535>
```

The default is 60 seconds.

- b. Configure the hello interval:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} hello-interval <1-65535>
```

The default is 10 seconds.

- c. Configure the retransmit interval:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} retransmit-interval  
<1-1800>
```

The default is 5 seconds.

- d. Configure the transit delay:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} transit-delay <1-1800>
```

The default is 1 second.

Example

Create a virtual link:

```
Switch:1(config-ospf)#ipv6 area virtual-link 0.0.0.1 2.2.2.2
```

Configure optional parameters for a virtual link:

```
Switch:1(config-ospf)#ipv6 area virtual-link 0.0.0.1 4.4.4.4 dead-interval 90 retransmit-  
interval 10
```

Variable definitions

Use the data in the following table to use the **ipv6 area virtual-link** command.

Variable	Value
<i>{A.B.C.D}</i> <i>{A.B.C.D}</i>	Specifies the ID for the transit area that the virtual link traverses and the router ID of the virtual neighbor. Do not use 0.0.0.0 for the transit area.
<i>dead-interval</i> <i><1-65535></i>	Specifies the number of seconds after which the neighbor declares the router down if it does not receive hello packets. Configure this value as a multiple of the hello interval. You must configure the same value on the virtual neighbor. The default is 60 seconds.
<i>hello-interval</i> <i><1-65535></i>	Specifies the number of seconds between hello packets that the router sends on this interface. Configure the same value on the virtual neighbor. The default is 10 seconds.
<i>retransmit-interval</i> <i><1-1800></i>	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface. This value also applies to the retransmissions of database description and link-state request packets. The default is 5 seconds.
<i>transit-delay</i> <i><1-1800></i>	Specifies the estimated number of seconds to transmit a link-state update packet over this interface. The default is 1 second.

Configure IPsec for the OSPF Virtual Link

Use the following procedure to configure and enable IPsec for the OSPF virtual link.

IPsec is disabled by default.

Before You Begin

- Configure the OSPF virtual link.
- Create the IPsec security association.

About This Task

Until you enable IPsec on both sides of the virtual links, the links cannot exchange OSPFv3 control messages, and the system drops OSPFv3 exchange packets.

You must disable IPsec before you can perform virtual link policy configuration changes.

For configuration examples of IPsec used with OSPFv3 virtual link, see [OSPFv3 virtual link IPsec configuration example](#) on page 2092.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create the IPsec policy under the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec
```

3. Configure the action of the IPsec policy under the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec action <drop|permit>
```

4. Configure the direction of the IPsec policy under the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec direction <both|in|out>
```

5. Link the security association to the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec security-association WORD<0-32>
```

6. Enable the IPsec policy created under the OSPF virtual link:

```
ipv6 area virtual-link {A.B.C.D} {A.B.C.D} ipsec enable
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#(config)router ospf
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec action permit
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec direction both
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec security-association test1
Switch:1(config-ospf)#ipv6 area virtual-link 1.1.1.1 2.2.2.2 ipsec enable
```

Variable definitions

Use the data in the following table to use the **ipv6 area virtual link {A.B.C.D} {A.B.C.D} ipsec** command.

Variable	Value
<code>{A.B.C.D} {A.B.C.D}</code>	The first IP address specifies the area IP address, and the second IP address specifies the virtual-link IP address.
<code>action <drop permit></code>	Configures the action of the IPsec policy under the OSPF virtual tunnel to one of the following: <ul style="list-style-type: none"> • drop—Drops the IP packets. • permit—Permits the IP packets. The default is permit.

Variable	Value
<code>direction <both in out></code>	Specifies the direction you want to protect with IPsec: <ul style="list-style-type: none"> in—Specifies ingress traffic. out—Specifies egress traffic. both—Specifies both ingress and egress traffic. The default is both.
<code>enable</code>	Enables the IPsec policy under the OSPF virtual link.
<code>security-association WORD<0-32></code>	Links the security association to the OSPF virtual link.

Configure OSPF Default Metrics

Before You Begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace `ipv6` with `ipv6 ospf`.

About This Task

Use the following procedure to configure global OSPF default metrics.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Configure OSPF default-cost:

```
ipv6 default-
```

```
cost {ethernet|fast-ethernet|forty-gig-ethernet|hundred-gig-ethernet|gig-ethernet|ten-gig-ethernet|twentyfive-gig-ethernet|vlan} <1-65535>
```



Note

Different hardware platforms support different port speeds. For more information, see your hardware documentation.

Example

Configure IPv6 default cost metric for Ethernet to 100, for fast Ethernet to 20, for gig-ethernet, twentyfive-gig-ethernet, forty-gig-Ethernet, and hundred-gig-ethernet to 2, and VLAN to 1.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#ipv6 default-cost ethernet 100
```

```

Switch:1(config-ospf)#ipv6 default-cost fast-ethernet 20
Switch:1(config-ospf)#ipv6 default-cost gig-ethernet 2
Switch:1(config-ospf)#ipv6 default-cost ten-gig-ethernet 2
Switch:1(config-ospf)#ipv6 default-cost Forty-gig-ethernet 2
Switch:1(config-ospf)#ipv6 default-cost twentyfive-gig-ethernet 2
Switch:1(config-ospf)#ipv6 default-cost hundred-gig-ethernet 2
Switch:1(config-ospf)#ipv6 default-cost vlan 1

```

Variable definitions

Use the data in the following table to use the **ipv6 default-cost** command.



Note

Different hardware platforms support different port speeds. For more information, see your hardware documentation.

Variable	Value
<i>ethernet</i> <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>ethernet</i> is for 10 Mb/s Ethernet (default is 100).
<i>fast-ethernet</i> <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>fast-ethernet</i> is for 100 Mb/s Fast-Ethernet (default is 10).
<i>forty-gig-ethernet</i> <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>forty-gig-ethernet</i> is for 10 Mb/s Forty-Gigabit-Ethernet (default is 1).
<i>gigabit-ethernet</i> <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>gigabit-ethernet</i> is for 10 Mb/s Gigabit-Ethernet (default is 1).
<i>hundred-gig-ethernet</i> <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>hundred-gig-ethernet</i> is for 100 Gigabit Ethernet (default is 1).
<i>ten-gig-ethernet</i> <1-65535>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <i>ten-gig-ethernet</i> is for 10 Mb/s Ten-Gigabit-Ethernet (default is 1).

Variable	Value
<code>twentyfive-gig-ethernet <1-65535></code>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. On a channelized 100 Gbps port, the default-cost for each 25 Gbps channel is 1.
<code>vlan <1-65535></code>	Configures the IPv6 OSPF default metrics. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet. <code>vlan</code> is for Vlan interfaces (default is 10).

Configure OSPF on a Port or VLAN

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface.

Before You Begin

- The IPv6 interface must exist.

Procedure

- Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- Create an OSPF area on the interface:

```
ipv6 ospf area {A.B.C.D}
```

- Enable OSPFv3 on the interface:

```
ipv6 ospf enable
```

The default is disabled.

4. Configure optional parameters to meet your requirements:

a. Configure the interface metric:

```
ipv6 ospf cost <0-65535>
```

The default for a brouter port or VLAN is 1.

**Note**

If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.

b. Configure the router dead interval:

```
ipv6 ospf dead-interval <1-65535>
```

The default is 40 seconds.

c. Configure the hello interval:

```
ipv6 ospf hello-interval <1-65535>
```

The default is 10 seconds.

d. Configure the link LSA suppression:

```
ipv6 ospf link-lsa-suppression
```

**Note**

Before configuring Link LSA suppression for OSPF, configure Link LSA suppression for OSPF area for point to point (p2p) or point to multipoint interfaces (p2mp), otherwise it defaults to a broadcast interface type where you cannot use Link LSA suppression.

e. Configure the poll interval:

```
ipv6 ospf poll-interval <0-65535>
```

The default is 120 seconds.

f. Configure the interface priority:

```
ipv6 ospf priority <0-255>
```

The default is 1.

g. Configure the retransmit interval:

```
ipv6 ospf retransmit-interval <1-1800>
```

The default is 5 seconds.

h. Configure the transit delay:

```
ipv6 ospf transit-delay <1-1800>
```

The default is 1 second.

Example

Create an OSPF area on the interface:

```
Switch:1(config-if)#ipv6 ospf area 0.0.0.0
```

Enable OSPFv3 on the interface:

```
Switch:1(config-if)#ipv6 ospf enable
```

Variable Definitions

Use the data in the following table to use the **ipv6 ospf** command.

Variable	Value
<i>area {A.B.C.D}</i>	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
<i>cost <0-65535></i>	Specifies the cost for the interface. The default for a brouter port or VLAN is 1.
<i>dead-interval <1-65535></i>	Specifies the number of seconds after which the neighbor declares the router down, if it does not receive hello packets. Configure this value as a multiple of the hello interval. You must configure the same value on the virtual neighbor. The default is 40 seconds.
<i>enable</i>	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is disabled.
<i>hello-interval <1-65535></i>	Specifies the number of seconds between hello packets that the router sends on this interface. Configure the same value on the virtual neighbor. The default is 10 seconds.
<i>link-lsa-suppression</i>	Configures link LSA suppression on the specified port or VLAN. It is only used for point to point or point to multipoint interfaces. By default, it is disabled.
<i>poll-interval <0-65535></i>	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
<i>priority <0-255></i>	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.

Variable	Value
<i>retransmit-interval</i> <1-1800>	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface. This value also applies to the retransmissions of database description and link-state request packets. The default is 5 seconds.
<i>transit-delay</i> <1-1800>	Specifies the estimated number of seconds to transmit a link-state update packet over this interface. The default is 1 second.

Use the data in the following table to use the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure OSPF on a Tunnel

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface.

Before You Begin

- The IPv6 interface must exist.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable

configure terminal

router ospf
```
2. Create an OSPF area on the interface:

```
ipv6 tunnel <1-2000> area {A.B.C.D}
```
3. Enable OSPFv3 on the interface:

```
ipv6 tunnel <1-2000> enable
```

4. Configure optional parameters to meet your requirements:

a. Configure the router dead interval:

```
ipv6 tunnel <1-2000> dead-interval <1-65535>
```

The default is 40 seconds.

b. Configure the hello interval:

```
ipv6 tunnel <1-2000> hello-interval <1-65535>
```

The default is 10 seconds.

c. Configure the interface metric:

```
ipv6 tunnel <1-2000> metric <0-65535>
```

d. Configure the poll interval:

```
ipv6 tunnel <1-2000> poll-interval <0-65535>
```

The default is 120 seconds.

e. Configure the interface priority:

```
ipv6 tunnel <1-2000> priority <0-255>
```

The default is 1.

f. Configure the retransmit interval:

```
ipv6 tunnel <1-2000> retransmit-interval <1-1800>
```

The default is 5 seconds.

g. Configure the transit delay:

```
ipv6 tunnel <1-2000> transit-delay <1-1800>
```

The default is 1 second.

Example

Create an OSPF area on the interface:

```
Switch:1(config-if)#ipv6 tunnel 4 area 0.0.0.0
```

Enable OSPFv3 on the interface:

```
Switch:1(config-if)#ipv6 tunnel 4 enable
```

Variable definitions

Use the data in the following table to use the **ipv6 tunnel** command.

Variable	Value
<1-2000>	Specifies the tunnel ID.
area {A.B.C.D}	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.

Variable	Value
<i>dead-interval</i> <1-65535>	<p>Specifies the number of seconds after which the neighbor declares the router down if it does not receive hello packets. Configure this value as a multiple of the hello interval.</p> <p>Tip: You must configure the same value on the virtual neighbor.</p> <p>The default is 40 seconds.</p>
<i>enable</i>	<p>Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas.</p> <p>If you disable the status, the interface is external to OSPFv3. The default is enabled.</p>
<i>hello-interval</i> <1-65535>	<p>Specifies the number of seconds between hello packets that the router sends on this interface.</p> <p>Tip: You must configure the same value on the virtual neighbor.</p> <p>The default is 10 seconds.</p>
<i>metric</i> <0-65535>	<p>Specifies the cost for the interface.</p> <p>The default for a tunnel is 100.</p>
<i>poll-interval</i> <0-65535>	<p>Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor.</p> <p>The default is 120.</p>
<i>priority</i> <0-255>	<p>Specifies the priority of this interface.</p> <p>Multiaccess networks use the priority in the designated router election.</p> <p>A higher priority value increases the likelihood that the router becomes the designated router.</p> <p>A value of zero (0) indicates the router cannot become the designated router for the network.</p> <p>If more than one router uses the same priority value, the system uses the router ID to determine which router becomes the designated router.</p> <p>The default is 1.</p>
<i>retransmit-interval</i> <1-1800>	<p>Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface.</p> <p>The retransmit-interval value also applies to the retransmissions of database description and link-state request packets.</p> <p>The default is 5 seconds.</p>
<i>transit-delay</i> <1-1800>	<p>Specifies the estimated number of seconds required to transmit a link-state update packet over this interface.</p> <p>The default is 1 second.</p>

Configure OSPFv3 on the CLIP Interface for the Global Router

About This Task

Perform this task to configure an Open Shortest Path First Version 3 (OSPFv3) instance for a specific circuitless IP address on the Global Router (GRT).

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enter OSPF Configuration mode:
`router ospf`
3. Create the IPv6 OSPF area:
`ipv6 area {A.B.C.D}`
4. Return to Global Configuration mode:
`end`
5. Enable IPv6 OSPF router configuration:
`router ospf ipv6-enable`
6. Enter Loopback Interface Configuration mode:
`interface loopback <1-256>`
7. Create the IPv6 loopback interface address:
`ipv6 interface address WORD<0-255>`
8. Create the OSPFv3 instance:
`ipv6 ospf [<1-256>] area {A.B.C.D}`
9. Enable the OSPFv3 instance:
`ipv6 ospf [<1-256>] enable`

Example

Configuring the OSPFv3 instance on the loopback interface on the GRT:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#ipv6 area 192.2.0.1
Switch:1(config-ospf)#end
Switch:1(config)#router ospf ipv6-enable
Switch:1(config)#interface Loopback 2
Switch:1(config-if)#ipv6 interface address 2001::10:10:0:1/128
Switch:1(config-if)#ipv6 ospf 2 area 192.2.0.1
Switch:1(config-if)#ipv6 ospf 2 enable
```

Variable Definitions

The following table defines parameters for the **ipv6 area** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the area address.

The following table defines parameters for the **interface loopback** command.

Variable	Value
<i><1-256></i>	Specifies the interface ID.

The following table defines parameters for the **ipv6 interface address** command.

Variable	Value
<i>WORD<0-255></i>	Specifies the IPv6 address for the loopback interface.

The following table defines parameters for the **ipv6 ospf** command.

Variable	Value
<i><1-256></i>	Specifies the interface ID for this IPv6 ospf interface.
<i>area {A.B.C.D}</i>	Specifies the IPv6 OSPF area, typically formatted as an IP address. Enter the address in the <i>{A.B.C.D}</i> portion.

Configure OSPFv3 on the CLIP Interface for a VRF Instance

About This Task

Perform this task to configure an Open Shortest Path First Version 3 (OSPFv3) instance for a specific circuitless IP address on the Virtual Router Forwarding (VRF) instance.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create a VRF instance and specify a VRF name:
`ip vrf WORD<1-16>`
3. Enter VRF Router Configuration mode:
`router vrf WORD<1-16>`
4. Create an OSPFv3 instance on the specific VRF instance:
`ipv6 ospf`
5. Enable OSPFv3 on the specific VRF instance:
`ipv6 ospf enable`

6. Enter Loopback Interface Configuration mode:

```
interface Loopback <1-256>
```
7. Create the IPv6 loopback interface address for the specific VRF instance:

```
ipv6 interface address WORD<0-255> vrf WORD<1-16>
```
8. Create the OSPFv3 instance:

```
ipv6 ospf [<1-256>] area {A.B.C.D}
```
9. Enable the OSPFv3 instance:

```
ipv6 ospf [<1-256>] enable
```

Example

Configuring the OSPFv3 instance on the loopback interface on the VRF test:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip vrf test
Switch:1(config)#router vrf test
Switch:1(config-vrf)#ipv6 ospf
Switch:1(router-vrf)#ipv6 ospf enable
Switch:1(router-vrf)#interface Loopback 10
Switch:1(config-if)#ipv6 interface address 2001::10:10:0:1/128 vrf test
Switch:1(config-if)#ipv6 ospf 10 area 192.2.0.1
Switch:1(config-if)#ipv6 ospf 10 enable
```

Variable Definitions

The following table defines parameters for the **ip vrf** command and the **router vrf** command.

Variable	Value
<i>WORD<1-16></i>	Specifies the name of the VRF.

The following table defines parameters for the **interface loopback** command.

Variable	Value
<i><1-256></i>	Specifies the ID of the loopback interface.

The following table defines parameters for the **ipv6 interface** command.

Variable	Value
<i>address WORD<0-255></i>	Specifies the IPv6 loopback interface address. Enter the interface address in the <i>WORD<0-255></i> portion.
<i>vrf WORD<1-16></i>	Specifies the name of the VRF. Enter the VRF name in the <i>WORD<1-16></i> portion.

The following table defines parameters for the **ipv6 ospf** command.

Variable	Value
<1-256>	Specifies the interface ID value.
area {A.B.C.D}	Specifies the IPv6 OSPFv3 area address, typically formatted as an IP address. Enter the IP address in the {A.B.C.D} portion.

View OSPFv3 Information

View information about OSPF to view the current configuration.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View OSPF global information:
show ipv6 ospf [vrf WORD <1-16>] [vrfids WORD <0-512>]
3. View OSPF areas:
show ipv6 ospf area [vrf WORD <1-16>] [vrfids WORD <0-512>]
4. View OSPF interface information
show ipv6 ospf interface [gigabitEthernet {slot/port[sub-port]}|vlan <1-4059>] [vrf WORD <1-16>] [vrfids WORD <0-512>]
5. View OSPF interface timers:
show ipv6 ospf int-timers [vrf WORD <1-16>] [vrfids WORD <0-512>]
6. View the link-state database (LSDB) table:
show ipv6 ospf lsdb [adv-rtr <A.B.C.D>] [area <A.B.C.D>] [interface gigabitEthernet {slot/port[sub-port]}|vlan <1-4059>] [lsa-type <1-11>] [lsid <0-4294967295>] [scope <1-3>] [tunnel <1-2000>] [detail] [vrf WORD <1-16>] [vrfids WORD <0-512>]
7. View OSPF neighbors to see routers with interfaces to a common network, including neighbors on the virtual link to the OSPF backbone:
show ipv6 ospf neighbor [vrf WORD <1-16>] [vrfids WORD <0-512>]

Examples

```
Switch:1#show ipv6 ospf
```

```
=====
                        OSPFv3 Global Information - GlobalRouter
=====
```

```

router-id                : 170.76.84.0
admin-state              : DISABLE
version                  : 3
area-bdr-rtr-state       : FALSE
as-bdr-rtr-state         : FALSE
helper-mode              : ENABLED
as-scope-lsa-count       : 0
lsa-checksum             : 0
originate-new-lsas       : 0
```

```

rx-new-lsas                : 0
ext-lsa-count              : 0

    default-metric :
        ethernet - 100
        fast-ethernet - 10
        gig-ethernet - 1
        ten-gig-ethernet - 1
        forty-gig-ethernet - 1

        vlan - 10
Switch:1>show ipv6 ospf area

=====
                        OSPF Area - GlobalRouter
=====
AREA_ID          STUB_AREA  NSSA   IMPORT_SUM  TRANS_ROLE  TRANS_STATE
-----
0.0.0.0          false     false true         always     disabled
STUB_METRIC STUB_METRIC_TYPE SPF_RUNS  BDR_RTR_CNT ASBDR_RTR_CNT LSA_CNT LSACK_SUM
-----
10              ospfV3Metric    0         0           0           0         0

Switch:1#show ipv6 ospf interface
Total ospf areas: 1
Total ospf interfaces: 2
=====
                        OSPF Interface - GlobalRouter
=====
IFINDX(VID/BRT/CLIP) AREAID          ADM IFSTATE  METRIC PRI DR/BDR          IFTYPE
-----
192  (BRT-1/1      ) 0.0.0.0          dis DOWN    0       1  0.0.0.0          BROADCAST
                                0.0.0.0
1353 (CLIP-10      ) 0.0.0.0          ena LOOPBACK 0       1  0.0.0.0          PASSIVE
                                0.0.0.0
2148 (VID-100     ) 0.0.0.0          ena DOWN    10      1  0.0.0.0          BROADCAST

3 out of 3 Total Num of ospf interfaces displayed

Total ospf virtual interfaces: 0

=====
                        OSPF Virtual Interface - GlobalRouter
=====
AREAID          NBRIPADDR          STATE
-----

0 out of 0 Total Num of ospf virtual interfaces displayed

Switch:1#show ipv6 ospf int-timers

=====
                        OSPF Interface Timers - GlobalRouter
=====
IFINDX(VID/BRT/CLIP) AREAID          TRANSIT  RETRANS  HELLO  DEAD  POLL
                                DELAY    INTERVAL INTERVAL INTERVAL INTERVAL
-----
2059 (BRT-11      ) 0.0.0.0          1        5       10     40    120
2060 (CLIP-12     ) 0.0.0.0          1        5       10     40    120
=====

```

```

=====
                        OSPF Virtual Interface Timers - GlobalRouter
=====
AREAID                NBRIPADDR      TRANSIT  RETRANS  HELLO    DEAD
                        DELAY          INTERVAL INTERVAL INTERVAL
-----
Switch:1#show ipv6 ospf neighbor
=====
                        OSPF Neighbor - GlobalRouter
=====
IFINDX(VID/BRT/CLIP)  NBRROUTERID    NBRIPADDR                                STATE    TTL
-----
331      (1/19)        97.146.128.0    fe80:0:0:0:2ef4:c5ff:fe92:8a00    Restart  120

1 out of 1 Total Num of Neighbor Entries displayed.

=====
                        OSPF Virtual Neighbor - GlobalRouter
=====
NBRAREAID            NBRROUTERID    VIRTINTFID NBRIPV6ADDR                                STATE
-----
0 out of 0 Total Num of Virtual Neighbor Entries displayed.

=====
                        OSPF NBMA Neighbor - GlobalRouter
=====
INTERFACE NBRROUTERID    NBRIPADDR                                STATE
-----
0 out of 0 Total Num of NBMA Neighbor Entries displayed.
=====

```

Variable Definitions

Use the data in the following table to use the **show ipv6 ospf lsdb** commands.

Variable	Value
<i>adv-rtr</i> <A.B.C.D>	Shows information for the specified advertising router.
<i>area</i> <A.B.C.D>	Shows information for the specified area.
<i>detail</i>	Shows information beyond the basic information.
<i>interface</i> <i>gigabitethernet</i> { <i>slot/port</i> [/ <i>sub-port</i>]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>lsa-type</i> <1-11>	Shows information for the specified LSA type.
<i>lsid</i> <0-4294967295>	Shows information for the specified link-state ID.

Variable	Value
<i>scope</i> <1-3>	Shows information for the specified scope: <ol style="list-style-type: none"> 1. link-scope LSAs: View the link-scope LSDB to view the LSAs that are not flooded beyond the local link. 2. area-scope LSAs: View the area-scope LSDB to see the LSAs that are flooded in a single OSPF area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix-LSAs. 3. AS-scope LSAs: View the AS-scope LSDB to see the LSAs that are flooded through the routing domain. The AS scope is used for ASexternal- LSAs.
<i>tunnel</i> <1-2000>	Specifies the ID number of the tunnel.
<i>vrf</i> <i>WORD</i> <1-16>	Specifies the VRF name.
<i>vrfids</i> <i>WORD</i> <0-512>	Specifies VRF IDs.

View OSPFv3 Default Cost Information

About This Task

Use the following procedure to view the OSPF default cost information, to ensure accuracy.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the OSPF cost information:
show ipv6 ospf default-cost [*vrf* *WORD* <1-16>] [*vrfids* *WORD* <0-512>]

Example

```
Switch:1#show ipv6 ospf default-cost
=====
IPv6 OSPF Default Metric - GlobalRouter
=====
 10MbpsPortDefaultMetric: 100
 100MbpsPortDefaultMetric: 10
 1000MbpsPortDefaultMetric: 1
 10000MbpsPortDefaultMetric: 1
 25000MbpsPortDefaultMetric: 1
 40000MbpsPortDefaultMetric: 1
 100000MbpsPortDefaultMetric: 1
      VlanDefaultMetric: 10
```

Add an NBMA Neighbor

Add an NBMA neighbor for each interface that is eligible to become the DR.

An NBMA interface with a positive nonzero router priority is eligible to become the DR for the NBMA network and is configured with the identification of all attached routers, IPv6 addresses, and router priorities.

Before You Begin

- Identify the following information:
 - specific interfaces to include in the NBMA network
 - the IPv6 address for each interface
 - the router priority for each interface
 - the hello interval for the network
 - the router dead interval for the network
 - the poll interval for the network

About This Task

In contrast to a broadcast network where switches multicast (send to AllSPFRouters and AllDRouters) certain OSPF protocol packets, switches replicate and send NBMA packets to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination addresses AllSPFRouters and AllDRouters. Because the NBMA network does not broadcast, you must manually configure a list of neighbors and priorities for all routers in the network that can become the DR. Potential DRs use a positive nonzero router priority.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create a new NBMA neighbor:

```
ipv6 ospf nbma-nbr WORD<0-43> <0-255>
```

3. Change the priority of an existing NBMA neighbor:

```
ipv6 ospf nbma-nbr WORD<0-43> priority <0-255>
```

Example

Create an NBMA neighbor that will not become the DR:

```
Switch:1(config-if)#ipv6 ospf nbma-nbr fe80:0:0:0:8217:7dff:fe76:8a03 0
```


Variable definitions

Use the data in the following table to use the **ipv6 ospf nbma-nbr** command.

Variable	Value
<i>priority</i> <0-255>	Specifies the priority to use for this neighbor in the designated router election process. A value of 0 indicates the neighbor cannot become the designated router. The higher the priority value, the higher chance the switch will win the election process. The default is 1.
<i>WORD</i> <0-43>	Specifies the IPv6 address of the neighbor.

Configure Link LSA Suppression

About This Task

Use the following procedure to configure link LSA suppression on a port or a VLAN, to decrease unnecessary link LSA generation and flooding for non-broadcast and non-NBMA interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enter the following command:

```
ipv6 ospf area {A.B.C.D} network {p2p | p2mp} link-lsa-suppression
```

Example

Variable definitions

Following table describes the variables to the **ipv6 ospf area {A.B.C.D} network p2p link-lsa-suppression** command.

Variable	Description
<i>area</i> {A.B.C.D}	Create an IPv6 OSPF area.
<i>network</i>	Sets the type of interface.

Variable	Description
<i>[eth NBMA p2mp p2p passive]</i>	Specifies the type of interface.
<i>link-lsa-suppression</i>	Enables link LSA suppression.

Configure Route Redistribution to OSPFv3 in GRT Mode

Configure a redistribute entry to announce certain routes into the OSPFv3 domain, including static routes, direct routes, RIPng, OSPF routes, IS-IS routes, or Border Gateway Protocol (BGP) routes. Optionally, use a route policy to control the redistribution of routes.



Note

RIPng is not virtualized, therefore RIPng redistribution in OSPFv3 works only in GRT.

Before You Begin

- Enable OSPFv3 globally.
- Ensure that a route policy exists.
- Ensure that you set OSPFv3 as the boundary router.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For a redistribute policy (OSPFv3) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.
- You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace *ipv6* with *ipv6 ospf*.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create the redistribution instance:

```
ipv6 redistribute <bgp|direct|isis|rip|static>
```



Note

The switch loads the existing configuration when you upgrade to the current release. Once you have the configuration file saved using the current release, only the new configuration will be loaded.

3. Apply a route policy if required:

```
ipv6 redistribute <bgp|direct|isis|rip|static> route-map WORD<0-64>
```



Note

No inter-vrf Route Redistribution is supported for IPv6.

4. Configure other parameters, as required.

5. Enable the redistribution.

```
ipv6 redistribute <bgp|direct|isis|rip|static> enable
```

6. Ensure that the configuration is correct:

```
show ipv6 ospf redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution.

```
ipv6 ospf apply redistribute <bgp|direct|isis|rip|static> [vrf WORD<1-16>]
```

Changes do not take effect until you apply them.

9. View all routes that are redistributed into OSPFv3:

- a. View the routes that are redistributed from the GRT to OSPFv3:

```
show ipv6 ospf redistribute
```

- b. View the routes that are redistributed to OSPFv3 for a specific VRF instance:

```
show ipv6 ospf redistribute [vrf WORD<1-64>] [vrfids WORD<0-512>]
```

Example

Redistribute static routes from the GRT to OSPF.

Create the redistribution instance, apply a route policy, enable redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#ipv6 redistribute static

WARNING: Routes will not be injected until apply command is issued after enable command
Switch:1(config-ospf)#ipv6 redistribute static route-map policy1
Switch:1(config-ospf)#ipv6 redistribute static enable
Switch:1(config-ospf)#exit
Switch:1(config)#ipv6 ospf apply redistribute static
Switch:1(config)#show ipv6 ospf redistribute
```

OSPFv3 Redistribute List - GlobalRouter				
SRC	MET	MTYPE	ENABLE	RPOLICY
STAT	0	type2	TRUE	policy1

Variable Definitions

Use the data in the following table to use the **ipv6 redistribute** command.

Variable	Value
<i>enable</i>	Enables the OSPF route redistribution instance.
<i>metric <0-65535></i>	Configures the metric to apply to redistributed routes.
<i>metric-type <type1 type2></i>	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
<i>route-map WORD<0-64></i>	Configures the route policy to apply to redistributed routes.
<i><bgp direct isis rip static></i>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, or static.

Use the data in the following table to use the **ipv6 ospf apply redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF instance.
<i><bgp direct isis rip static></i>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, or static.

View the Status of OSPFv3 Redistribution

View the status of OSPFv3 route redistribution to verify the current configuration. You can redistribute directly connected routes and IPv6 static routes into OSPFv3.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the current configuration:
show ipv6 ospf redistribute [vrf WORD <1-16>] [vrfids WORD <0-512>]

Example

```
Switch:1#show ipv6 ospf redistribute
```

```
=====
OSPFv3 Redistribute List - GlobalRouter
=====
```

```

SOURCE  MET      MTYPE      ENABLE  RPOLICY
-----
Static  0          external  TRUE
-----

```

Variable Definitions

Use the data in the following table to use the **show ipv6 ospf redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF name.
<i>vrfids WORD<0-512></i>	Specifies VRF IDs.

View OSPFv3 Statistics

View OSPFv3 statistics to analyze trends.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View statistics:

```
show ipv6 ospf statistics [vrf WORD <1-16>] [vrfids WORD <0-512>]
```

Example

View IPv6 OSPF statistics:

```

Switch:1>enable
Switch:1#show ipv6 ospf statistics

=====
                        OSPFv3 Statistics - GlobalRouter
=====
      NumTxPkt: 9958
      NumRxPkt: 8982
    NumTxDropPkt: 33
    NumRxDropPkt: 0
    NumRxBadPkt: 0
      NumSpfRun: 42
    LastSpfRun: 0 day(s), 02:44:32
    LsdbTblSize: 45
    NumBadLsReq: 0
    NumSeqMismatch: 0
    NumOspfAdjacencies: 7

```

Variable Definitions

Use the data in the following table to use the **show ipv6 ospf stats** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF name.
<i>vrfids WORD<0-512></i>	Specifies VRF IDs.

Clear OSPFv3 Statistics

About This Task

Use the following procedure to clear all OSPFv3 statistics.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Clear the statistics:
`clear ipv6 ospf stats [vrf WORD<1-16>] [vrfids WORD<0-512>]`

Variable Definitions

Use the data in the following table to use the **clear ipv6 ospf stats** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies VRF IDs.

Disable Helper Mode for OSPFv3

Before You Begin

You configure OSPF on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router configuration mode and replace *ipv6* with *ipv6 ospf*.

About This Task

By default, OSPF Helper mode is enabled when OSPF is configured. You can disable helper mode by the following command and re-enable it again by using “no” or “default” commands.

Procedure

1. Enter OSPF Router Configuration mode:
`enable`

`configure terminal`

`router ospf`
2. Enter the following command to disable Helper mode:
`ipv6 helper-mode-disable`
3. Enter the following command to enable Helper mode:
`no ipv6 helper-mode-disable`

Or

`default ipv6 helper-mode disable`

Example

Disabling Helper mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#ipv6 helper-mode-disable
```

Enabling Helper mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
Switch:1(config-ospf)#no ipv6 helper-mode-disable
```

OSPF Configuration Using EDM

Configure Open Shortest Path First (OSPF) parameters so that the switch can participate in OSPF routing operations. The following section describes procedures that you use while you configure OSPF using Enterprise Device Manager (EDM).

Configure OSPF Globally

Configure OSPF parameters, such as automatic virtual links and OSPF metrics, so you can control OSPF behavior on the system.

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.
- Assign an IP address to the switch.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **General** tab.
4. Specify the OSPF router ID.
5. In AdminStart, select **enabled**.
6. (Optional) If required, configure the metrics that OSPF uses for different link speeds. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.
7. (Optional) To enable the switch to use OSPF SNMP traps, select the **TrapEnable** check box.
8. (Optional) To enable the automatic creation of virtual links, select the **AutoVirtLinkEnable** check box.
9. (Optional) Configure the OSPF holddown timer as required.
10. Select **Apply**.

General Field Descriptions

Use the data in the following table to use the **General** tab.

Name	Description
RouterId	Specifies the OSPF router ID. This variable has the same format as an IP address but distinguishes this router from other routers in the OSPF domain.
AdminStat	Shows the administrative status of OSPF for the router. Enabled denotes that the OSPF process is active on at least one interface; disabled disables it for all interfaces. The default is disabled.
VersionNumber	Specifies the OSPF version.
AreaBdrRtrStatus	Denotes if this router is an area border router (ABR). AreaBdrRtrStatus value must be true to create a virtual router interface.
ASBdrRtrStatus	Specifies ASBR status. If you select the ASBdrRtrStatus check box, the router is an autonomous system boundary router (ASBR).
ExternLsaCount	Shows the number of external (LS type 5) link-state advertisements in the link-state database.
ExternLsaCksumSum	Shows the 32-bit unsigned sum of the link-state checksums of the external link-state advertisements in the link-state database. This sum determines if a change occurred in a router link-state database and compares the link-state databases of two routers.
OriginateNewLsas	Shows the number of new link-state advertisements originated from this router. This number increments each time the router originates a new link-state advertisement (LSA).
RxNewLsas	Shows the number of received link-state advertisements that are new instances. This number does not include new instances of self-originated link-state advertisements.
Rfc1583Compatability	Controls the preference rules used when the router chooses among multiple autonomous system external (ASE) LSAs which advertise the same destination. If enabled, the preference rule is the same as that specified by RFC 1583. If disabled, the preference rule is as described in RFC 2328, which can prevent routing loops when ASE LSAs for the same destination originate from different areas. The default is disable.
OpaqueLsaSupport	Specifies support for Opaque LSA types.
10MbpsPortDefaultMetric	Indicates the default cost applied to 10 Mbps interfaces (ports). The default is 100.
100MbpsPortDefaultMetric	Indicates the default cost applied to 100 Mbps interfaces (ports). The default is 10.

Name	Description
1000MbpsPortDefaultMetric	Indicates the default cost applied to 1 Gbps interfaces (ports). The default is 1.
10000MbpsPortDefaultMetric	Indicates the default cost applied to 10 Gbps interfaces (ports). The default is 1.
25000MbpsPortDefaultMetric	Indicates the default cost applied to 25 Gbps interfaces (channelized 100 Gbps ports). The default is 1.
40000MbpsPortDefaultMetric	Indicates the default cost applied to 40 Gbps interfaces (ports). The default is 1.
100000MbpsPortDefaultMetric	Indicates the default cost applied to 100 Gbps interfaces (ports). The default is 1.
VlanDefaultMetric	Configures the VLAN interfaces default metric. The default is 10.
TrapEnable	Indicates whether to enable traps for OSPF. The default is false.
AutoVirtLinkEnable	Enables or disables the automatic creation of virtual links. The default is false.
SpfHoldDownTime	Specifies the OSPF holddown timer (3–60 seconds). The default is 10 seconds. The holddown timer delays a metric change due to a routing table update by x seconds. If you configure the timer to 0, OSPF accepts a new metric change immediately.
OspfAction	Initiates a new Shortest Path First (SPF) run to update the routing table. The default is none.
LastSpfRun	Indicates the time since the last SPF calculation made by OSPF.
HelperModeDisable	Disables the helper mode. It is enabled by default.

Enable OSPF Globally

Enable OSPF globally enabled to use the protocol on the router. If you disable OSPF globally, all OSPF actions cease.

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **General** tab.
4. For **AdminStat**, select the **enabled** or **disabled** option button, as required.

5. Select **Apply**.

Configure Global Default Metrics

Configure the metrics that OSPF uses for different link speeds. The lower the metric, the more likely that OSPF chooses the link to route an OSPF packet.

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **General** tab.
4. Change the metric for one or all of the following:
 - 10MbpsPortDefaultMetric
 - 100MbpsPortDefaultMetric
 - 1000MbpsPortDefaultMetric
 - 10000MbpsPortDefaultMetric
 - 25000MbpsPortDefaultMetric
 - 40000MbpsPortDefaultMetric
 - 100000MbpsPortDefaultMetric
5. Select **Apply**.

Configure an OSPF interface

Configure OSPF parameters, such as authentication and priority, so you can control OSPF interface behavior. You can specify the interface as passive, broadcast, Non-Broadcast Multiple Access (NBMA), or Point-to-Point (p2p).

Before You Begin

- Enable OSPF globally.
- Ensure that the interface exists (the port or VLAN has an IP address).
- You must know the network OSPF password to use password authentication.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Interfaces** tab.
4. Select **Insert**.

5. Select the IP address for the interface from the IP Address list.
6. To designate a router priority, in the **RtrPriority** box, type a new value.
7. In the **Type** area, select the type of OSPF interface you want to create.
8. Select the authentication type you want in the **AuthType** field.
9. If you chose **simplePassword**, in the **AuthKey** box, type a password of up to eight characters.
10. To change their values, select the current value in the **HelloInterval**, **RtrDeadInterval**, or **PollInterval** boxes, and then type new values.
11. Select **Insert**.
12. On the **Interfaces** tab, Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IP Address	Specifies the IP address of the current OSPF interface.
AddressLessIf	Designates whether an interface has an IP address: Interfaces with an IP address = 0 Interfaces without IP address = ifIndex
Areald	Specifies the OSPF area name in dotted-decimal format. For VLANs, keeping the default area setting on the interface causes link-state database (LSDB) inconsistencies. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdminStat	Specifies the current administrative status of the OSPF interface (enabled or disabled).
State	Specifies the current state of the OSPF interface. The value can be one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
RtrPriority	Specifies the OSPF priority to use during the election process for the designated router. The interface with the highest priority becomes the designated router. The interface with the second-highest priority becomes the backup designated router. If the priority is 0, the interface cannot become the designated router or the backup. The range is 0–255. The default is 1.
DesignatedRouter	Specifies the IP address of the designated router.

Name	Description
BackupDesignatedRouter	Specifies the IP address of the backup designated router.
Type	<p>Specifies the type of OSPF interface (broadcast, NBMA, and p2p).</p> <p>Note: To make it passive:</p> <ol style="list-style-type: none"> 1. Create the interface. 2. After interface creation, select VLAN > VLANs to select the VLAN that is created with the OSPF interface. 3. Select the IP tab and select the IP interface that is created with the OSPF interface. 4. Select the OSPF tab and select Passive for the IfType.
AuthType	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
AuthKey	Specifies the key (up to 8 characters) required when you specify simple password authentication in the AuthType parameter.
HelloInterval	<p>Specifies the length of time, in seconds, between hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds.</p> <p>After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.</p>

Name	Description
TransitDelay	Specifies the length of time, in seconds, required to transmit an LSA update packet over the interface. The default is 1.
RetransInterval	Specifies the length of time, in seconds, required between LSA retransmissions. The default is 5.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the Hello interval. To avoid interpretability issues, the RtrDeadInterval value for the OSPF interface must match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. The default is 120.
Events	Indicates the number of times this OSPF interface has changed state, or an error has occurred.
LsaCount	Specifies the total number of link state advertisements in this area LSDB, excluding AS-external LSAs.
LsaChecksumSum	Specifies the number of link-state advertisements. This sum excludes external (LS type 5) link-state advertisements. The sum determines if a change occurred in a router LSDB and compares the LSDB of two routers.
DesignatedRouterId	Specifies the router ID of the designated router.
BackupDesignatedRouterId	Specifies the router ID of the backup designated router.

Change an OSPF Non-Passive Interface Type

Change the interface type to designate the interface as either passive, NBMA, or broadcast.

Before You Begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- If the interface is currently an NBMA interface with manually configured neighbors, you must first delete all manually configured neighbors.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Interfaces** tab.

4. To disable the interface, double-click the **AdminStat** cell, and then select **disabled**.
5. Select **Apply**.
6. To change the interface type, double-click the **Type** cell, and then choose the new interface type.
7. Select **Apply**.
8. To enable the interface, double-click the **AdminStat** cell, and then select **enabled**.
9. Select **Apply**.

Change an OSPF Passive Interface Type

Change the interface type to designate the interface as either passive, NBMA, or broadcast.

Before You Begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- If the interface is currently an NBMA interface with manually configured neighbors, you must first delete all manually configured neighbors.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select on the VLAN where the OSPF interface is created.
4. Select **IP**.
5. Select the IP Address where the OSPF interface is created.
6. Select the **OSPF** tab.
7. Clear the **Enable** check box to disable the OSPF interface.
8. Select **Apply**.
9. Modify the interface type to passive.
10. Select the **Enable** check box.
11. Select **Apply**.

Configure NBMA Interface Neighbors

Configure NBMA neighbors so that the interface can participate in designated router election. All neighbors that you manually insert on the Neighbors tab are NBMA neighbors.

Before You Begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.

- Ensure that the interface type is NBMA.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Neighbors** tab.
4. Select **Insert**.
5. Enter the IP address and priority for the first neighbor.
6. Select **Insert**.
7. Add all required neighbors.
8. Select **Apply**.

Neighbors Field Descriptions

Use the data in the following table to use the **Neighbors** tab.

Name	Description
NbrIpAddr	Specifies the neighbor IP address.
AddressLessIndex	Indicates addressed and addressless interfaces. This value is 0 on an interface with an IP address. On addressless interfaces, the corresponding value of ifIndex in the Internet standard management information base (MIB).
NbrRtrId	Specifies the router ID of the neighboring router. The router ID has the same format as an IP address but identifies the router independent of its IP address.
Options	Specifies the bit mask that corresponds to the neighbor options parameter.
Priority	Specifies the priority.
State	Specifies the OSPF interface state.
Events	Specifies the number of state changes or error events that occur between the OSPF router and the neighbor router.
Retransmission Queue Length	Specifies the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
ospfNbmaNbrPermanence	Indicates whether the neighbor is a manually configured NBMA neighbor; permanent indicates it is an NBMA neighbor.
HelloSuppressed	Indicates whether hello packets to a neighbor are suppressed.
RestartHelperStatus	Specifies whether the router is acting as a graceful restart helper for the neighbor.

Name	Description
RestartHelperAge	Specifies the remaining time in the current OSPF graceful restart interval, if the router is acting as a restart helper for the neighbor.
RestartHelperExitReason	Specifies the outcome of the last attempt at acting as a graceful restart helper for the neighbor.

Configure OSPF Interface Metrics

Configure the metrics associated with the peer layer interface to control OSPF behavior. For finer control over port-specific metric speed, you can specify the metric speed when you configure OSPF on a port.

Before You Begin

- Enable OSPF globally.
- Ensure that the interface uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **If Metrics** tab.
4. Double-click the value cell, and type a new value.
5. Select **Apply**.

When you enable a port for OSPF routing, the default metric in the port tab is 0. A value of 0 means that the port uses the default metrics for port types that you specify on the OSPF General tab.

If Metrics Field Descriptions

Use the data in the following table to use the **If Metrics** tab.

Name	Description
IP Address	Specifies the IP address of the device used to represent a point of attachment in a TCP/IP internetwork.
AddressLessIf	Indicates addressed and addressless interfaces. This variable is 0 on interfaces with IP addresses and equals ifIndex for interfaces that have no IP address.
TOS	Specifies the type of service (TOS). The TOS is a mapping to the IP type of service flags as defined in the IP forwarding table management information base (MIB).

Name	Description
Value	Indicates the metric from the OSPF router to a network in the range.
Status	Specifies the status of the interface as active or not active. This variable is read-only.

View All OSPF Enabled Interfaces

View all OSPF enabled interfaces to determine which interfaces use OSPF routing.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Interfaces** tab.
4. To ensure the system displays the latest information, select **Refresh**.

View OSPF Statistics

View OSPF statistics.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Stats** tab.

Stats Field Descriptions

Use the data in the following table to use the OSPF **Stats** tab.

Name	Description
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.
LsdbTblSize	Displays the number of entries in the link state database table.
TxPackets	Displays the number of packets transmitted by OSPF.
RxPackets	Displays the number of packets received by OSPF.
TxDropPackets	Displays the number of packets dropped before transmitted by OSPF.
RxDropPackets	Displays the number of packets dropped before received by OSPF.

Name	Description
RxBadPackets	Displays the number of packets received by OSPF that are bad.
SpfRuns	Displays the total number of SPF calculations performed by OSPF, which includes the number of partial route table calculation for incremental updates.
BuffersAllocated	Displays the number of buffers allocated for OSPF.
BuffersFreed	Displays the number of buffers that are freed by the OSPF.
BufferAllocFailures	Displays the number of times that OSPF has failed to allocate buffers.
BufferFreeFailures	Displays the number of times that OSPF has failed to free buffers.
Routes	Displays the number of OSPF routes added to RTM.
Adjacencies	Displays how many adjacencies are learned through the interface.
Areas	Displays the number of areas configured.
Nbrs	Indicates the number of OSPF neighbors.
BadLsReqs	Indicates the number of bad link state requests.
SeqMismatches	Indicates the number of sequence mismatched packets.
NumAllocDDP	Indicates the number of database description (DD) packet buffers allocated for OSPF.
NumFreeDDP	Indicates the number of DD packet buffers that are freed by the OSPF.

Graph OSPF Statistics

Graph OSPF statistics. The graph formats available are: line chart, area chart, bar chart, and pie chart.

Procedure

1. In the navigation pane, expand **Configuration > IP**
2. Select **OSPF**.
3. Select the **Stats** tab.
4. (Optional) To refresh the values in the table, select **Clear Counters**.
5. To specify the polling interval, from the **Poll Interval** drop down menu, select a value. The options are:

Option	Description
5s	The polling interval is 5 seconds.
10s	The polling interval is 10 seconds.
30s	The polling interval is 30 seconds.
1m	The polling interval is 1 minute.

Option	Description
5m	The polling interval is 5 minutes.
30m	The polling interval is 30 minutes.
1h	The polling interval is 1 hour.

6. Select one value; for example, AbsoluteValue or Cumulative.

- Or, select two values; for example, AbsoluteValue and Cumulative.

To select a second value, press the **Ctrl** key, then select the second value. You cannot select more than two values.

7. From the toolbar, select a chart icon. The options are:

Option	Description
Line Chart	Displays a line chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Area Chart	Displays an area chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Bar Chart	Displays a bar chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Pie Chart	Displays a pie chart for the values you selected against the polling interval.

The Chart Legend uses different colors to identify the values you selected that are plotted on the graph.

8. To switch the horizontal and vertical axes values, on the chart toolbar, select **Horizontal**.
9. To switch views of the log scale from high to low values, or low to high values, on the chart toolbar, select **Log Scale**.
10. To switch to another chart using the same values, on the chart toolbar, select a chart icon.

Configure OSPF on a Port

Configure OSPF parameters on a port so you can control OSPF behavior on the port.



Important

When you enable OSPF on a port, the switch automatically creates an area 0.0.0.0, and advertises it on the specific port, by default. To avoid this behavior, you must manually configure the port into a properly configured area on the switch.

Before You Begin

- Enable OSPF globally .
- Ensure that the port uses an IP address.
- Ensure that the ospf_md5key.txt file is on the switch to use MD5 authentication.
- You must know the network OSPF password to use password authentication.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **OSPF** tab.
5. Select **Enable**.
6. Specify the hello interval.
7. Specify the router dead interval.
8. Designate a router priority.
9. Configure a metric.
10. If you want, select an authentication type.
11. If you select **simplePassword** authentication, type a password in the **AuthKey** box.
12. Configure the area ID.
13. If desired, select the **AdvertiseWhenDown** check box.
14. Select an interface type.
15. Type a value in the **PollInterval** box.
16. For IfMtuIgnore, select either **enable** or **disable**.
17. For **BfdEnable**, select **enable**.
18. Select **Apply**.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
Enable	Enables or disables OSPF routing on the specified port. The default is false.
HelloInterval	<p>Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds.</p> <p>After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.</p>
RtrDeadInterval	<p>Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet, and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.</p>
DesigRtrPriority	<p>Specifies the priority of this port in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.</p>
Metric	<p>Specifies the metric for the type of service (TOS) on this port. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1.</p> <ul style="list-style-type: none"> FFFF—No route exists for this TOS. IPCP links—Defaults to 0. 0—Use the interface speed as the metric value when the state of the interface is up.

Name	Description
AuthType	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha1—Specifies secure hash algorithm (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. You can only access and enable the SHA-1 authentication type after you enable enhanced secure mode. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
AuthKey	Specifies the key (up to 8 characters) when you specify simple password authentication in the port AuthType variable.
AreaId	<p>Specifies the OSPF area name in dotted-decimal format.</p> <p>The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).</p>
AdvertiseWhenDown	<p>Advertises the network on this port as up, even if the port is down. The default is false.</p> <p>After you configure a port with no link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even if the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.</p>
IfType	<p>Specifies the type of OSPF interface (broadcast, NBMA, passive, or p2p).</p> <p>Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.</p>
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must have the same poll interval.

Name	Description
IfMtuIgnore	Specifies whether the interface ignores the global maximum transmission unit (MTU) configuration. To allow the switch to accept OSPF database description (DD) packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
BfdEnable	Enable Bidirectional Forwarding Detection (BFD) for OSPF.

Enable OSPF on a CLIP Interface

Before You Begin

- You must globally enable OSPF.
- The OSPF area must already exist.

About This Task

Enable Open Shortest Path First (OSPF) on a CLIP interface so that it can participate in OSPF routing.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select the required CLIP interface.
5. Select **OSPF**.
6. Select the **Enable** check box.
You must enable OSPF on the CLIP interface for CLIP to function.
7. In the current **AreaId** field, enter the IP address of the OSPF backbone area.
8. Select **Apply**.

Circuitless OSPF Field Descriptions

Use the data in the following table to use the **Circuitless OSPF** tab.

Name	Description
Enable	Enables OSPF on the CLIP interface.
AreaId	Specifies the OSPF area ID.

View Port OSPF Statistics

View port OSPF statistics to manage network performance.

Procedure

1. On the Device Physical View, select a port.

2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **OSPF** tab.

OSPF Field Descriptions

The following table describes parameters on the **OSPF** tab.

Name	Description
VersionMismatches	Specifies the number of version mismatches received by this interface.
AreaMismatches	Specifies the number of area mismatches received by this interface.
AuthTypeMismatches	Specifies the number of authentication type mismatches received by this interface.
AuthFailures	Specifies the number of authentication failures.
NetmaskMismatches	Specifies the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Specifies the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Specifies the number of dead interval mismatches received by this interface.
OptionMismatches	Specifies the number of option mismatches in the hello interval or dead interval fields received by this interface.
RxHellos	Specifies the number of hello packets received by this interface.
RxDBDescrs	Specifies the number of database descriptor packets received by this interface.
RxLSUpdates	Specifies the number of link state update packets received by this interface.
RxLSReqs	Specifies the number of link state request packets received by this interface.
RxLSAcks	Specifies the number of link state acknowledge packets received by this interface.
TxHellos	Specifies the number of hello packets transmitted by this interface.
TxDBDescrs	Specifies the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Specifies the number of link state update packets transmitted by this interface.
TxLSReqs	Specifies the number of link state request packets transmitted by this interface.
TxLSAcks	Specifies the number of link state acknowledge packets transmitted by this interface.

Graph OSPF Statistics for a Port

Use statistics to help you monitor OSPF performance on a port. You can also graph statistics for all OSPF packets.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **OSPF** tab.
5. Select **Graph**.
6. Select one or more values.
7. Select the type of graph.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
VersionMismatches	Indicates the number of version mismatches received by this interface.
AreaMismatches	Indicates the number of area mismatches received by this interface.
AuthTypeMismatches	Indicates the number of AuthType mismatches received by this interface.
AuthFailures	Indicates the number of authentication failures.
NetMaskMismatches	Indicates the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Indicates the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Indicates the number of dead interval mismatches received by this interface.
OptionMismatches	Indicates the number of options mismatches received by this interface.
RxHellos	Indicates the number of hello packets received by this interface.
RxDBDescrs	Indicates the number of database descriptor packets received by this interface.
RxLSUpdates	Indicate the number of Link state update packets received by this interface.
RxLsReqs	Indicates the number of Link state request packets received by this interface.
RxLSAcks	Indicates the number of Link state acknowledge packets received by this interface.

Name	Description
TxHellos	Indicates the number of hello packets transmitted by this interface.
TxDBDescrs	Indicates the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Indicate the number of Link state update packets transmitted by this interface.
TxLSReqs	Indicates the number of Link state request packets transmitted by this interface.
TxLSAcks	Indicates the number of Link state acknowledge packets transmitted by this interface.

Configure OSPF on a VLAN

Configure OSPF parameters on a VLAN to control OSPF behavior on the VLAN.



Important

When you enable OSPF on a VLAN, the switch automatically creates an area 0.0.0.0, and advertises it on the specific VLAN, by default. To avoid this behavior, you must manually configure the VLAN into a properly configured area on the switch.

Before You Begin

- Enable OSPF globally.
- Ensure that the VLAN uses an IP address.
- Ensure that the ospf_md5key.txt file is on the switch to use MD5 authentication.
- Ensure that you know the network OSPF to use password authentication.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **OSPF** tab.
The information on the OSPF tab applies only to a routed port or VLAN, which means the VLAN uses an IP address.
7. To enable OSPF on the VLAN interface, select the **Enable** check box.
8. To change their values, select the current value in the **HelloInterval**, **RtrDeadInterval**, or **PollInterval** boxes, and then type new values.

9. To designate a router priority, in the **DesigRtrPriority** box, type the new value.
10. Select the authentication type in the **AuthType** field.
11. If you chose **simplePassword**, in the **AuthKey** box, type a password of up to eight characters.
12. Select the interface type you want to create.
13. Select **Apply**.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
Enable	Enables or disables OSPF routing on the specified VLAN. The default is false.
HelloInterval	Specifies the length of time, in seconds, between the transmission of hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. After you change the hello interval values, you must save the configuration file, and then restart the switch. After the switch restarts, it restores the values and checks for consistency.
RtrDeadInterval	Specifies the interval used by adjacent routers to determine if the router was removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the hello interval. To avoid interoperability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
DesigRtrPriority	Specifies the priority of this VLAN in multiaccess networks to use in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. If a tie occurs, routers use their router ID as a tie breaker. The default is 1.
Metric	Specifies the metric for this TOS on this VLAN. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. <ul style="list-style-type: none"> FFFF—No route exists for this TOS. IPCP links—Defaults to 0. 0—Use the interface speed as the metric value when the state of the interface is up.

Name	Description
AuthType	<p>Specifies the type of authentication required for the interface.</p> <ul style="list-style-type: none"> • none—Specifies that no authentication required. • simple password—Specifies that all OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey parameter. • MD5 authentication—Specifies that all OSPF updates received by the interface must contain the MD5 key. • sha1—Specifies secure hash algorithm 1 (SHA-1), which is a cryptographic hash function that produces a 160-bit hash value, usually given in a hexadecimal number, 40 digits long. You can only access and enable the SHA-1 authentication type after you enable enhanced secure mode. • sha-2—Specifies SHA-2, which offers the hash function SHA-256. <p>Note: sha-2, an update of SHA-1, can offer six hash functions that include SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA 512/256, with hash values that are 224, 256, 384, or 512 bits. However, the current release supports only SHA-256.</p>
AuthKey	Specifies the key (up to eight characters) when you specify simple password authentication in the VLAN AuthType variable.
AreaId	Specifies the OSPF area name in dotted-decimal format. The area name is not related to an IP address. You can use a suitable value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdvertiseWhenDown	<p>Advertises the network even if the port is down. If true, OSPF advertises the network on this VLAN as up, even if the port is down. The default is false.</p> <p>After you configure a port without a link and enable AdvertiseWhenDown, it does not advertise the route until the port is active. Then, OSPF advertises the route even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown.</p>
IfType	<p>Specifies the type of OSPF interface (broadcast, NBMA, passive, or p2p).</p> <p>Before you change an OSPF interface type, you must first disable the interface. If the interface is an NBMA interface, you must also delete all configured neighbors.</p>
PollInterval	Specifies the length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must use the same poll interval.

Name	Description
IfMtuIgnore	Specifies whether the VLAN ignores the MTU configuration. To allow the switch to accept OSPF DD packets with a different MTU size, enable MtuIgnore. The interface drops incoming OSPF DD packets if their MTU is greater than 1500 bytes.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for OSPF.

View Graphs for OSPF on a VLAN

Perform the following procedure to view graphs for OSPF on a VLAN. You can use graph statistics to monitor OSPF performance on a VLAN. The graph formats available are: line chart, area chart, bar chart, and pie chart.

Before You Begin

- OSPF must be enabled.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN, and then select **IP**.
5. Select the **OSPF** tab.
6. Select **Graph**.
7. (Optional) To refresh the values in the table, select **Clear Counters**.
8. To specify the polling interval, from the **Poll Interval** drop down menu, select a value. The options are:

Choice Option	Choice Description
5s	The polling interval is 5 seconds.
10s	The polling interval is 10 seconds.
30s	The polling interval is 30 seconds.
1m	The polling interval is 1 minute.
5m	The polling interval is 5 minutes.
30m	The polling interval is 30 minutes.
1h	The polling interval is 1 hour.

9. Select one or two values.
To select two values, for example, AbsoluteValue and Cumulative. Select the first value, and then press the **Control** key to select the second value. You cannot select more than two values.

10. From the toolbar, select a chart icon. The options are:

Choice Option Choice Description

Line Chart	Displays a line chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Area Chart	Displays an area chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Bar Chart	Displays a bar chart for the values you selected against the polling interval. The X axis represents time. The vertical axis represents the logging scale.
Pie Chart	Displays a pie chart for the values you selected against the polling interval.

The Chart Legend uses different colors to identify the values you selected that are plotted on the graph.

11. To switch the horizontal and vertical axes values, on the chart toolbar, select **Horizontal**.
12. To switch views of the log scale from high to low values, or low to high values, on the chart toolbar, select **Log Scale**.
13. To switch to another chart using the same values, on the chart toolbar, select a chart icon.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the OSPF-Graph tab.
Average/sec	Displays the average value for each second.
Minimum/sec	Displays the minimum value for each second.
Maximum/sec	Displays the maximum value for each second.
LastVal/sec	Displays the last value for each second.
VersionMismatches	Displays the number version mismatches received by this interface.
AreaMismatches	Displays the number area mismatches received by this interface.
AuthTypeMismatches	Displays the number AuthType mismatches received by this interface.
AuthFailures	Displays the number Authentication failures.

Name	Description
NetMaskMismatches	Displays the number net mask mismatches received by this interface.
HelloIntervalMismatches	Displays the number hello interval mismatches received by this interface.
DeadIntervalMismatches	Displays the number dead interval mismatches received by this interface.
OptionMismatches	Displays the number options mismatches received by this interface.
RxHellos	Displays the number hello packets received by this interface.
RxDBDescrs	Displays the number database descriptor packets received by this interface.
RxLSUpdates	Displays the number Link state update packets received by this interface.
RxLSReqs	Displays the number Link state request packets received by this interface.
RxLSAcks	Displays the number Link state acknowledge packets received by this interface.
TxHellos	Displays the number hello packets transmitted by this interface.
TxDBDescrs	Displays the number database descriptor packets transmitted by this interface.
TxLSUpdates	Displays the number Link state update packets transmitted by this interface.
TxLSReqs	Displays the number Link state request packets transmitted by this interface.
TxLSAcks	Displays the number Link state acknowledge packets transmitted by this interface.

Create Stubby or Not-so-Stubby OSPF Areas

Import information from other areas to learn their OSPF relationships. Perform this procedure to create normal, stubby, or not-so-stubby areas (NSSA).

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About This Task

Place stubby areas or NSSAs at the edge of an OSPF routing domain. Ensure that you configure all routers in the stubby or NSSA as stubby or NSSA, respectively.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Areas** tab.
The backbone ID has an area ID of 0.0.0.0.
4. Select **Insert**.
5. Configure the area ID.
6. Select an option in the ImportAsExtern area.
To add a not-so-stubby (NSSA) area, select **importNssa**. To import external LSAs (create a normal OSPF area), select **importExternal**. To not import external LSAs (create a stubby area), select **importNoExternal**.
7. Select **Apply**.

Areas Field Descriptions

Use the data in the following table to use the **Areas** tab.

Name	Description
AreaId	Specifies a 32-bit integer that uniquely identifies an area. Area ID 0.0.0.0 is the OSPF backbone. For VLANs, using the default area on the interface causes LSDB inconsistencies.
ImportAsExtern	Specifies the method to import ASE link-state advertisements. The value can be importExternal (default), importNoExternal, or importNssa.
SpfRuns	Specifies the number of SPF calculations performed by OSPF.
AreaBdrRtrCount	Specifies the number of area border routers reachable within this area. Each SPF pass calculates this value, initially zero.
AsBdrRtrCount	Specifies the number of autonomous system border routers reachable within this area. Each SPF pass calculates this value, initially zero.
AreaLsaCount	Specifies the total number of link state advertisements in this area LSDB, excluding AS-external LSAs.
AreaLsaCksumSum	Specifies the number of link-state advertisements. This sum excludes external (LS type 5) link-state advertisements. The sum determines if a change occurred in a router LSDB and compares the LSDB of two routers.
AreaSummary	Specifies whether to send summary advertisements in a stub area.
ActiveIfCount	Specifies the number of active interfaces in this area.
AreaNssaTranslatorRole	Indicates if the NSSA border router can perform NSSA translation of Type 7 LSAs to Type 5 LSAs. The possible values are always or candidate. The default is candidate.

Name	Description
AreaNssaTranslatorState	Indicates if and how an NSSA border router translates Type 7 LSAs to Type 5 LSAs. The possible values are: <ul style="list-style-type: none">• enabled — the border router always translates the LSAs• elected — a candidate border router translates the LSAs• disabled — a candidate border router does not translate the LSAs
AreaNssaTranslatorStabilityInterval	Specifies number of seconds after an elected translator determines its services are no longer required.
AreaNssaTranslatorEvents	Specifies the number of translator state changes that have occurred since the last boot up.

Configure Stub Area Metrics Advertised by an ABR

Configure metrics to control the use of routes in a routing domain.

Before You Begin

- Enable OSPF globally.
- Ensure that the port uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Stub Area Metrics** tab.
4. Double-click the metric value to specify a new metric speed for the required stub areas.
5. Select **Apply**.

Stub Area Metrics *Field Descriptions*

Use the data in the following table to use the **Stub Area Metrics** tab.

Name	Description
AreaId	Specifies the 32-bit identifier for the stub area.
TOS	Specifies the type of service associated with the metric.
Metric	Specifies the metric value applied at the indicated type of service. By default, the value equals the lowest metric value at the type of service among the interfaces to other areas.
Status	Specifies the status of the stub area. This variable is read-only.

Insert OSPF Area Aggregate Ranges

Use aggregate area ranges to reduce the number of link-state advertisements required within the area. You can also control advertisements.

Before You Begin

- Enable OSPF globally.
- Ensure that the port uses an IP address.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Area Aggregate** tab.
4. Select **Insert**.
5. Type the area ID.
6. Select the type of link-state database.
7. Type the IP address of the network.
8. Type the subnet mask.
9. Select the effect.
10. In **AdvertiseMetric**, type a cost to advertise for the OSPF area range.
11. Select **Insert**.

Area Aggregate Field Descriptions

Use the data in the following table to use the **Area Aggregate** tab.

Name	Description
AreaID	Specifies the area in which the address exists.
LsdbType	Specifies the LSDB type: <ul style="list-style-type: none">• summaryLink—aggregated summary link• nssaExternalLink—not so stubby area link
IP Address	Specifies the IP address of the network or subnetwork indicated by the range.
Mask	Specifies the network mask for the area range.
Effect	Specifies advertisement methods: <ul style="list-style-type: none">• advertiseMatching means advertise the aggregate summary LSA with the same LSID.• doNotAdvertiseMatching means suppress all networks that fall within the entire range.• advertiseDoNotAggregate means advertise individual networks.

Name	Description
AdvertiseMetric	Changes the advertised metric cost for the OSPF area range.
ExtRouteTag	Specifies the external route tag to be included in NSSA (type-7) LSAs.

Enable Automatic Virtual Links

Use automatic virtual links to provide an automatic, dynamic backup link for vital OSPF traffic.

Before You Begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **General** tab.
4. Select **AutoVirtLinkEnable**.
5. Select **Apply**.

Configure a Manual Virtual Interface

Use manual virtual links (interfaces) to provide a backup link for vital OSPF traffic with a minimum of resource use.

Before You Begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Virtual If** tab.
4. Select **Insert**.
5. Specify the area ID of the transit area.
The transit area is the common area between two ABRs.
6. Specify the neighbor ID.
The neighbor ID is the IP router ID of the ABR that the other ABR needs to reach the backbone.

7. Select **Insert**.
8. To verify that the virtual link is active, select **Refresh** and check the **State** column.
If the state is point-to-point, the virtual link is active. If the state is down, the virtual link configuration is incorrect.

Virtual If Field Descriptions

Use the data in the following table to use the **Virtual If** tab.

Name	Description
Areald	Specifies the transit area ID that the virtual link traverses.
Neighbor	Specifies the router ID of the virtual neighbor.
TransitDelay	Specifies the estimated number of seconds required to transmit a link-state update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between link-state advertisement, and retransmissions for adjacencies that belong to this interface. This variable also applies to DD and link-state request packets. This value must exceed the expected round-trip time. The default is 5.
HelloInterval	Specifies the length of time, in seconds, between the hello packets that the router sends on the interface. This value must be the same for the virtual neighbor. The default is 10.
RtrDeadInterval	Specifies the number of seconds that expires before neighbors declare the router down. This value must be a multiple of the hello interval. This value must be the same for the virtual neighbor. The default is 60.
State	Specifies the OSPF virtual interface state.
Events	Specifies the number of state changes or error events on this virtual Link.
AuthType	Specifies the authentication type specified for a virtual interface. You can locally assign additional authentication types. The default is none.
AuthKey	Specifies the authentication password. If AuthType is a simple password, the device adjusts and zeros fill the eight octets. Unauthenticated interfaces need no authentication key, and simple password authentication cannot use a key with more than eight octets.
LsaCount	Specifies the total number of link state advertisements in this virtual interface LSDB.
LsaCksumSum	Specifies the number of link-state advertisements. The sum determines if a change occurred in a virtual interface LSDB and compares the virtual interfaceLSDB of the virtual neighbors.

View Virtual Neighbors

View virtual neighbors to view the area and virtual link configuration for the neighboring device.

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Virtual Neighbors** tab.

Virtual Neighbors Field Descriptions

Use the data in the following table to use the **Virtual Neighbors** tab.

Name	Description
Area	Specifies the sub-network in which the virtual neighbor resides.
RtrId	Specifies the 32-bit integer (represented as an IP address) that uniquely identifies the neighbor router in the autonomous system.
IP Address	Specifies the IP address of the virtual neighboring router.
Options	Specifies the bit mask that corresponds to the neighbor options parameter.
State	Specifies the OSPF interface state.
Events	Specifies the number of state changes or error events that occurred between the OSPF router and the neighbor router.
LsRetransQLen	Specifies the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
HelloSuppressed	Specifies whether hello packets from the neighbor are suppressed.
RestartHelperStatus	Specifies whether the router is acting as a graceful restart helper for the neighbor.
RestartHelperAge	Specifies the remaining time in the current OSPF graceful restart interval, if the router is acting as a restart helper for the neighbor.
RestartHelperExitReason	Specifies the outcome of the last attempt at acting as a graceful restart helper for the neighbor.

Configure Host Routes

Configure host routes when the switch resides in a network that uses routing protocols other than OSPF. A host route is a more-specific route and is used even if it is higher cost than a network route.

Before You Begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About This Task

You can specify which hosts directly connect to the router and the metrics and types of service to advertise for the hosts.

Use a host route to create a custom route to a specific host to control network traffic.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Hosts** tab.
4. To insert a new host, select **Insert**.
5. In **IP Address**, type the area IP address of the new host.
6. In **Metric**, type the metric to advertise.
7. Select **Insert**.
8. Select **Apply**.

Hosts Field Descriptions

Use the data in the following table to use the **Hosts** tab.

Name	Description
IpAddress	Specifies the IP address of the host that represents a point of attachment in a TCP/IP internetwork.
TOS	Specifies the type of service of the route.
Metric	Specifies the metric advertised to other areas. The value indicates the distance from the OSPF router to a network in the range.
AreaID	Specifies the area where the host is found. By default, the area that submits the OSPF interface is in 0.0.0.0.

Enable ASBR Status

Enable the ASBR status to make the switch an autonomous system boundary router (ASBR). Use ASBRs to advertise nonOSPF routes into OSPF domains so that the routes pass through the domain. A router can function as an ASBR if one or more of its

interfaces connects to a non-OSPF network, for example, Routing Information Protocol (RIP), BGP, or Exterior Gateway Protocol (EGP).

Before You Begin

- Enable OSPF globally.
- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About This Task

To conserve resources, you can limit the number of ASBRs on your network or specifically control which routers perform as ASBRs to control traffic flow.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **General** tab.
4. Select **ASBdrRtrStatus**.
5. Select **Apply**.

Manage OSPF Neighbors

View or delete OSPF neighbors to control OSPF operations.

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About This Task

The OSPF Hello protocol initiates and maintains neighbor relationships. The exception is that, in an NBMA network, you must manually configure permanent neighbors on each router eligible to become the DR. You can add neighbors for NBMA interfaces, but all other neighbors are dynamically learned.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Neighbors** tab.
4. To delete a manually configured neighbor, select the neighbors with a value of **permanent** in the **ospfNbmaNbrPermanence** column.
5. Select **Delete**.
6. Select **Apply**.

View the Link-State Database

View the area advertisements and other information in the LSDB to ensure correct OSPF operations.

Before You Begin

- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.
- The VRF must have an RP trigger of OSPF.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Link State Database** tab.

Link State Database Field Descriptions

Use the data in the following table to use the **Link State Database** tab.

Name	Description
AreaId	Identifies the area. The OSPF backbone uses the area ID 0.0.0.0.
Type	Specifies the OSPF interface type. Broadcast LANs, such as Ethernet and IEEE 802.5, use broadcast; X.25 and similar technologies use NBMA; and links that are point-to-point use pointToPoint.
Lsid	Identifies the piece of the routing domain that the advertisement describes.
RouterId	Identifies the router in the autonomous system.
Sequence	Identifies old and duplicate link-state advertisements.
Age	Specifies the age, in seconds, of the link-state advertisement.
Checksum	Contains the checksum of the complete contents of the advertisement, except for the age parameter. The checksum does not include the age parameter so that advertisement age increments without updating the checksum.

View the Local Link-State Database

About This Task

View the link-local advertisements for each virtual interface in the link-state database.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Local Lsdb** tab.

Local Lsdb Field Descriptions

Use the data in the following table to use the **Local Lsdb** tab.

Name	Description
IpAddress	Specifies the IP address of the interface from which the link-state advertisement was received if the interface is numbered.
AddressLessIf	Specifies the interface index of the interface from which the LSA was received if the interface is unnumbered.
Type	Specifies the type of the link-state advertisement.
Lsid	Specifies the piece of the routing domain that the advertisement describes.
RouterId	Specifies the originating router in the Autonomous System.
Sequence	Identifies old and duplicate link-state advertisements.
Age	Specifies the age, in seconds, of the link-state advertisement.
Checksum	Contains the checksum of the complete contents of the advertisement, except for the age parameter. The checksum does not include the age parameter so that advertisement age increments without updating the checksum.
Advertisement	Specifies the size of the entire link-state advertisement.

View AS-Scope Link-State Advertisements

About This Task

View the AS-scope link-state advertisements.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **As Lsdb** tab.

As Lsdb Field Descriptions

Use the data in the following table to use the **As Lsdb** tab.

Name	Description
Type	Specifies the type of the link-state advertisement.
Lsid	Specifies the piece of the routing domain that the advertisement describes.
RouterId	Specifies the originating router in the Autonomous System.
Sequence	Specifies old and duplicate link-state advertisements.
Age	Specifies the age, in seconds, of the link-state advertisement.
Checksum	Contains the checksum of the complete contents of the advertisement, except for the age parameter. The checksum does not include the age parameter so that advertisement age increments without updating the checksum.
Advertisement	Specifies the size of the entire link-state advertisement.

View Virtual Links in the Link-State Database

About This Task

View virtual links in the link-state database.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Virt Local Lsdb** tab.

Virt Local Lsdb Field Descriptions

Use the data in the following table to use the **Virt Local Lsdb** tab.

Name	Description
TransitArea	Specifies the ID for the transit area that the virtual link traverses.
Neighbor	Specifies the router ID of the of the virtual neighbor.
Type	Specifies the type of the link-state advertisement.
Lsid	Specifies the piece of the routing domain that the advertisement describes.

Name	Description
RouterId	Specifies the originating router in the Autonomous System.
Sequence	Specifies old and duplicate link-state advertisements.
Age	Specifies the age, in seconds, of the link-state advertisement.
Checksum	Contains the checksum of the complete contents of the advertisement, except for the age parameter. The checksum does not include the age parameter so that advertisement age increments without updating the checksum.
Advertisement	Specifies the size of the entire link-state advertisement.

View the Area LSA Count

About This Task

View the link-state advertisements for a given area.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Area Lsa Count** tab.

Area Lsa Count Field Descriptions

Use the data in the following table to use the **Area Lsa Count** tab.

Name	Description
AreaId	Specifies the OSPF area.
LsaType	Specifies the link-state advertisement type.
Number	Specifies the number link-state advertisements of a given type for a given area.

Configure Route Redistribution to OSPF

Before You Begin

- Enable OSPF globally.
- Ensure that a route policy exists.

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About This Task

Configure a redistribute entry to announce routes of a certain source protocol type into the OSPF domain, for example, static, RIP, or direct. Optionally, use a route policy to control the redistribution of routes.



Important

Changing the OSPF redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform this procedure. As a best practice, if you want to change default preferences for an OSPF redistribute context, do so before you enable the protocols.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. Select an option for the route source.
6. Select **enable**.
7. Select a route policy.
8. Configure the metric type.
9. Configure the subnet.
10. (Optional) Select **AutoFilterTag** to enable route filtering using tags.
11. Select **Insert**.

Redistribute Field Descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrflid	Specifies the destination virtual router forwarding instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrflid	Specifies the source VRF instance.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) an OSPF redistribute entry for a specified source type.

Name	Description
RoutePolicy	Configures the route policy (by name) to use for detailed redistribution of external routes from a specified source into an OSPF domain. Select the ellipsis (...) and choose from the list in the dialog box.
Metric	Configures the OSPF route redistribution metric for basic redistribution. A value of 0 indicates to use the original cost of the route.
MetricType	Configures the OSPF route redistribution metric type. The default is type 2. The cost of a type 2 route is the external cost, regardless of the interior cost. A type 1 cost is the sum of both the internal and external costs.
Subnets	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.
AutoFilterTag	Enables route filtering to ensure tagged routes are not redistributed into BGP. The default is disabled.

Force Shortest-Path Calculation Updates

Manually initiate an SPF run, or calculation, to immediately update the OSPF LSDB. This configuration is useful if

- you need to immediately restore a deleted OSPF-learned route
- the routing table entries and the LSDBs do not synchronize

Before You Begin

- Change the VRF instance as required to configure OSPF on a specific VRF instance. The VRF must have an RP trigger of OSPF. Not all parameters are configurable on non-default VRFs.

About This Task



Important

This process is computationally intensive. Use this command only if required.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **OSPF**.
3. Select the **General** tab.
4. In **OspfAction**, select **runSpf**.
5. Select **Apply**.
6. Select **Yes** to force an SPF run.

After you initiate an SPF run, wait at least 10 seconds before you initiate another SPF run.

OSPFv3 Configuration using EDM

Use the procedures in this section to configure OSPFv3 using EDM.

Configure OSPFv3 Globally

Configure OSPFv3 globally to enable it on the system and to configure the router ID.

Before You Begin

- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.



Note

Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPsec on OSPFv3 virtual link interfaces

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Globals** tab.
4. Type the router ID, in the format of an IPv4 address.
5. Select **enabled**.
6. (Optional) Select **ASBdrRtrStatus** to make the router an AS boundary router.
Enable the ASBR if the router attaches at the edge of the OSPF network, and has one or more interfaces that run an interdomain routing protocol. The default is disabled.
7. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.



Note

Different hardware platforms support different port speeds. For more information, see your hardware documentation.

Name	Description
RouterId	Specifies a 32-bit integer that identifies the router in the autonomous system. This value must be unique. The default value will be one of the IPv4 interface addresses.
AdminStat	Enables or disables OSPFv3 on the router. If you disable OSPFv3 globally, you disable it on all interfaces. The default is disabled.

Name	Description
VersionNumber	Shows the OSPF version number, which for IPv6 is version 3.
AreaBdrRtrStatus	Shows if the router is an area border router.
ASBdrRtrStatus	Configures the router as an autonomous system boundary router. The default is disabled (clear).
HelperModeDisable	Disables Graceful Restart Helper Mode feature.
AsScopeLsaCount	Shows the number of AS-external link-state advertisements in the LSDB.
AsScopeLsaCksumSum	Shows the sum of the checksums for the link-scope LSAs in the LSDB. Use the sum to determine if a change in the LSDB occurs, and to compare the LSDBs of the two routers.
OriginateNewLsas	Shows the number of new link-state advertisements. The number increases each time the router originates a new LSA.
RxNewLsas	Shows the number of new link-state advertisements received. This number does not include new instances of self-originated link-state advertisements.
ExtLsaCount	Shows the number of external (LS type 0x4005) LSAs in the LSDB.
10MbpsPortDefaultMetric	Indicates the default cost applied to 10 Mbps interfaces (ports). The default is 100.
100MbpsPortDefaultMetric	Indicates the default cost applied to 100 Mbps interfaces (ports). The default is 10.
1000MbpsPortDefaultMetric	Indicates the default cost applied to 1 Gbps interfaces (ports). The default is 1.
10000MbpsPortDefaultMetric	Indicates the default cost applied to 10 Gbps interfaces (ports). The default is 1.
25000MbpsPortDefaultMetric	Indicates the default cost applied to 25 Gbps interfaces (channelized 100 Gbps ports). The default is 1.
40000MbpsPortDefaultMetric	Indicates the default cost applied to 40 Gbps interfaces (ports). The default is 1.
100000MbpsPortDefaultMetric	Indicates the default cost applied to 100 Gbps interfaces (ports). The default is 1.
vlanDefaultMetric	Indicates the default cost applied to VLAN interfaces. The default is 10.

Create an OSPFv3 Area

Create an area to subdivide the autonomous system (AS) into areas that group contiguous networks, routers that connect to these networks, and attached hosts. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic compared to treating the entire AS like a single link-state domain.

Before You Begin

- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.



Note

Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPsec on OSPFv3 virtual link interfaces

About This Task

A stub area does not receive advertisements for external routes, which reduces the size of the link-state database (LSDB). A stub area uses only one area border router (ABR). Any packets destined for outside the area are routed to the area border exit point, examined by the ABR, and forwarded to a destination.

A not so stubby area (NSSA) prevents the flooding of AS-External link-state advertisements into the area by replacing them with a default route. NSSAs also import small stub (non- OSPF) routing domains into OSPF.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Areas** tab.
4. Select **Insert**.
5. Type the area ID.
6. Select **Insert**.

Areas Field Descriptions

Use the data in the following table to use the **Areas** tab.

Name	Description
Id	Specifies a 32-bit integer to uniquely identify an area. Use 0.0.0.0 for the OSPFv3 backbone.
ImportasExtern	<p>Indicates the support for importing AS-external LSAs:</p> <ul style="list-style-type: none"> • importExternal—normal area • importNoExternal—stub area • importNssa—not-so-stubby-area <p>AS-scope LSAs are not imported into stub areas or NSSAs. NSSAs import AS-External data at Type 7 LSAs, which use area scope. importExternal is the default.</p>
SpfRuns	Shows the number of times the intra-area route table was calculated using the LSDB of this area.
BdrRtrCount	Shows the number of reachable ABRs in this area. The value starts at zero (0). The system calculates this value in each SPF run.
AsBdrRtrCount	Shows the number of reachable ASBRs in this area. The value starts at zero (0). The system calculates this value in each SPF run.
ScopeLsaCount	Shows the number of area-scope LSAs in the LSDB for this area.
ScopeLsaCksumSum	Shows the sum of the checksums for the area-scope LSAs in the LSDB. Use the sum to determine if a change in the LSDB occurs, and to compare the LSDBs of the two routers.
Summary	Controls the import of inter-area LSAs into a stub area. If the value is noAreaSummary , the router does not originate nor propagate inter-area LSAs into the stub area. If the value is sendAreaSummary (the default), the router both summarizes and propagates inter-area LSAs.
StubMetric	Configures the metric value advertised for the default route to stub and NSSA areas.
NssaTranslatorRole	Indicates if the NSSA border router can perform NSSA translation of Type 7 LSAs to Type 6 LSAs. The possible values are always or candidate. The default is candidate.

Name	Description
NssaTranslatorState	Indicates if and how an NSSA border router translates Type 7 LSAs to Type 5 LSAs. The possible values are <ul style="list-style-type: none"> • enabled—The border router always translates the LSAs. • elected—A candidate border router translates the LSAs. • disabled—A candidate border router does not translate the LSAs.
StubMetricType	Specifies the type of metric advertised as a default route. The possible values are: <ul style="list-style-type: none"> • ospfv3Metric—OSPF metric • comparableCost—external Type 1 • nonComparable—external Type 2 The default is ospfv3Metric.

Create OSPFv3 Area Ranges

Create an area address range on the OSPFv3 router to reduce the number of area border router (ABR) advertisements into other OSPF areas. An area address range is an implied contiguous range of area network addresses for which the ABR advertises a single summary route.

Before You Begin

- You must create the OSPF area.
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.



Note

Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPSec on OSPFv3 virtual link interfaces

About This Task

If you create two ranges, and one range is a subset of the other, the router uses the most specific match.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Area Aggregate** tab.
4. Select **Insert**.
5. Select the area ID.

6. Select the type of area.

interAreaPrefixLsa generates an aggregated summary.

nssaExternalLsa generates an NSSA link summary.

7. Type the prefix for the IPv6 area address.

8. Type the number of bits from the IPv6 address that you want to advertise.

9. Select **Insert**.

Area Aggregate *Field Descriptions*

Use the data in the following table to use the **Area Aggregate** tab.

Name	Description
AreaID	Specifies the area in which the address aggregate exists. Use dotted decimal notation to specify the area name.
AreaLsdbType	Specifies the area LSDB type to which the address aggregate applies. interAreaPrefixLsa generates an aggregated summary. nssaExternalLsa generates an NSSA link summary.
Prefix	Specifies the IPv6 prefix. The prefix and prefix length define the range.
PrefixLength	Specifies the length of the prefix, in bits. The prefix cannot be shorter than 3 bits. The prefix and prefix length define the range.
Effect	Specifies the advertisement mode for prefixes in the range. advertiseMatching advertises the aggregate summary LSA with the same link-state ID. doNotAdvertiseMatching does not advertise networks that fall within the range.
AdvertiseMetric	Specifies a cost value to advertise for the OSPF area range. This value applies to summary LSAs (Type 3). If the value is 0, OSPF uses the cost to the farthest point in the network that is summarized.

Create an OSPFv3 Virtual Link

Create a virtual link if the switch does not connect directly to the backbone. The switch can create automatic virtual links or you can perform this procedure to create virtual links manually. Manual virtual links conserve resources and provide specific control over virtual link placement in your OSPFv3 configuration.

Before You Begin

- The router must be an ABR to create a virtual router interface.

- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.

**Note**

Non-default VRFs do not support the configuration of the following parameters:

- OSPFv3 interfaces over IPv6 tunnels
- IPsec on OSPFv3 virtual link interfaces

About This Task

Virtual linking is similar to backup redundancy. The switch creates a virtual link for vital traffic paths in your OSPFv3 configuration if traffic is interrupted, such as when an interface cable that provides a connection to the backbone (either directly or indirectly) is disconnected from the switch. Automatic virtual linking ensures that a link is created by using another switch.

OSPF routes cannot be learned through an ABR unless it connects to the backbone directly or through a virtual link.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Virtual If** tab.
4. Select **Insert**.
5. Specify the ID for the transit area.

The transit area is the common area between two ABRs.

6. Specify the router ID for the virtual neighbor.

The neighbor ID is the IP router ID of the ABR through which the other ABR must route traffic destined for the backbone.

7. Select **Insert**.
8. Select **Refresh** to verify that the virtual link is active.

If the state is point-to-point, the virtual link is active. If the state is down, the virtual link is configured incorrectly.

Virtual If Field Descriptions

Use the data in the following table to use the **Virtual If** tab.

Name	Description
Areaid	Specifies the ID for the transit area that the virtual link traverses. Do not use 0.0.0.0.
Neighbor	Specifies the router ID of the virtual neighbor.

Name	Description
TransitDelay	Specifies the estimated number of seconds to transmit a link-state update packet over this interface. The default is 1 second.
RetransInterval	Specifies the number of seconds between link-state advertisement retransmissions for adjacencies that belong to this interface. This value also applies to the retransmissions of database description and link-state request packets. The default is 5 seconds.
HelloInterval	Specifies the number of seconds between hello packets that the router sends on this interface. Configure the same value on the virtual neighbor. The default is 10 seconds.
RtrDeadInterval	Specifies the number of seconds after which the neighbor declares the router down if it does not receive hello packets. Configure this value as a multiple of the hello interval. You must configure the same value on the virtual neighbor. The default is 60 seconds.
State	Shows the state of the virtual interface: either down or pointToPoint.
Events	Shows the number of state changes or error events on the virtual link.
LinkScopeLsaCount	Shows the number of link-scope LSAs in the LSDB for the virtual link.
LinkLsaCksumSum	Shows the sum of the checksums for the link-scope LSAs in the LSDB. Use the sum to determine if a change in the LSDB occurs, and to compare the LSDBs of the two routers.

Create an OSPF Interface on a Brouter Port

Configure the OSPF protocol on an IPv6 interface to support dynamic routing on the interface. Perform this procedure to create an OSPF interface on a brouter port.

If you want to modify existing OSPFv3 interfaces, see [Modify an OSPFv3 Interface](#) on page 2920. To configure OSPFv3 on an IPv6 VLAN, see [Create an OSPF VLAN Interface](#) on page 2909.

Before You Begin

- The IPv6 interface must exist.

Procedure

1. In the Device Physical view, select a port.

2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select the **IPv6 OSPF Interface** tab.
5. Select **Insert**.
6. Select the area ID.
7. Select **enabled**.
8. Select **Insert**.

IPv6 OSPFv3 Interface Field Descriptions

Use the data in the following table to use the **IPv6 OSPFv3 Interface** tab.

Name	Description
Index	Specifies the interface index for the IPv6 interface on which OSPFv3 is configured.
Areaid	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Type Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none"> • broadcast • NBMA • point-to-point • point-to-multipoint • passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.

Name	Description
RetransInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this parameter to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this parameter to the same value for all routers attached to a common network. The default is 40.
PollInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.

Name	Description
MetricValue Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies whether Link LSA suppression is enabled.

Configure OSPFv3 on a CLIP Interface

Configure the Open Shortest Path First Version 3 (OSPFv3) instance on circuitless IP (CLIP) interfaces for the Global Router or a specific Virtual Router Forwarding (VRF) instance.

Before You Begin

- Ensure the CLIP interface exists.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

You can also perform this configuration using [Create an OSPF Interface on a CLIP](#) on page 2907.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**
2. Select **IPv6**.
3. Select the **Circuitless IPv6** tab.
4. Select a CLIP interface.
5. Select **OSPF**.
6. Select **enabled**.
7. Select the ellipsis (...) to choose an area ID, and then select **Ok**.
8. Select **Apply**.

OSPF Field Descriptions

Use the data in the following table to use the **OSPF** tab.

Name	Description
AdminStat	Enables the OSPFv3 instance on the CLIP interface.
Areald	Assigns an OSPF area to the OSPFv3 instance.

Create an OSPF Interface on a CLIP

Before You Begin

- The CLIP interface must exist.

About This Task

Configure the OSPF protocol on an IPv6 interface to support dynamic routing on the interface. Perform this procedure to create an OSPFv3 interface on a circuitless IP (CLIP).

If you want to modify existing OSPFv3 interfaces, see [Modify an OSPFv3 Interface](#) on page 2920.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Interfaces** tab.
4. Select **Insert**.
5. Select **IPv6 Interfaces**.
6. Select the index associated to a CLIP interface, and then select **Ok**.
7. Select **enabled** from **AdminStat**.
8. Select **Insert**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Index	Specifies the interface index for the IPv6 interface on which OSPFv3 is configured.
Areald	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.

Name	Description
Type Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none"> • broadcast • NBMA • point-to-point • point-to-multipoint • passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this parameter to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this parameter to the same value for all routers attached to a common network. The default is 40.

Name	Description
PollInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies whether Link LSA suppression is enabled.

Create an OSPF VLAN Interface

Configure the OSPF protocol on an IPv6 VLAN to support dynamic routing on the interface.

If you want to modify existing OSPFv3 interfaces, see [Modify an OSPFv3 Interface](#) on page 2920.

Before You Begin

- The IPv6 interface must exist.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.

3. Select the **Basic** tab.
4. Select a **VLAN**.
5. Select **IPv6**.
6. Select the **IPv6 OSPF Interface** tab.
7. Select **Insert**.
8. Select the area ID.
9. Select **enabled**.
10. Select **Insert**.

IPv6 OSPF Interface *Field Descriptions*

Use the data in the following table to use the **IPv6 OSPF Interface** tab.

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.
AreaId	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Type	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none"> • broadcast • NBMA • point-to-point • point-to-multipoint • passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.

Name	Description
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this parameter to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this parameter to the same value for all routers attached to a common network. The default is 40.
PollInterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.

Name	Description
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.
BfdEnable	Enables Bidirectional Forwarding Detection (BFD) for OSPF.

Create an OSPF Interface on a Tunnel

Configure the OSPF protocol on an IPv6 interface to support dynamic routing on the interface. Perform this procedure to create an OSPF interface on a tunnel.

If you want to modify existing OSPFv3 interfaces, see [Modify an OSPFv3 Interface](#) on page 2920. To configure OSPFv3 on an IPv6 VLAN, see [Create an OSPF VLAN Interface](#) on page 2909.

Before You Begin

- The IPv6 interface must exist.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **Tunnel**.
3. Select the **Tunnel Config** tab.
4. Select a configured tunnel.
5. Select **IPv6 OSPF**.
6. Select **Insert**.
7. Select the area ID.
8. Select **enabled**.
9. Select **Insert**.

OSPF Interface Field Descriptions

Use the data in the following table to use the **OSPF Interface** tab.

Name	Description
Index	Shows the interface index for the IPv6 interface on which OSPFv3 is configured.
Areaid	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Type	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none">• broadcast• NBMA• point-to-point• point-to-multipoint

Name	Description
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this field to the same value for all routers attached to a common network. The default is 10.
RtrDeadInterval	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this field to the same value for all routers attached to a common network. The default is 40.
PollInterval	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.

Name	Description
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression	Specifies whether Link LSA suppression is enabled.

View the AS-Scope Link-State Database

Before You Begin

- Change the VRF instance as required to view OSPFv3 on a specific VRF instance.

About This Task

View the AS-scope link-state database (LSDB) to see the LSAs that are flooded through the routing domain. The AS scope is used for AS external-LSAs.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **AS-scope LSDB** tab.

AS-scope LSDB *Field Descriptions*

Use the data in the following table to use the **AS-scope LSDB** tab.

Name	Description
Type	Shows the type of the link-state advertisement. Each link state type has a separate advertisement format. AS-scope LSAs not recognized by the router can be stored in the database.
RouterId	Shows the 32 bit number that uniquely identifies the originating router in the autonomous system.
Lsid	Identifies the piece of the routing domain that is being described by the advertisement.
Sequence	Shows a signed 32-bit integer that detects old and duplicate link-state advertisements. The larger the sequence number, the more recent the advertisement.
Age	Shows the age of the link-state advertisement in seconds.
Checksum	Indicates the checksum of the complete contents of the advertisement, except the age field. The age field is not affected so that the advertisement age value increments without updating the checksum. The checksum used is the same for ISO connectionless datagrams, the Fletcher checksum.

View the Area-Scope LSDB

Before You Begin

- Change the VRF instance as required to view OSPFv3 on a specific VRF instance.

About This Task

View the area-scope LSDB to see the LSAs that are flooded in a single OSPFv3 area. Area scope is used in router LSAs, network LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix-LSAs.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Area-scope LSDB** tab.

Area-scope LSDB Field Descriptions

Use the data in the following table to use the **Area-scope LSDB** tab.

Name	Description
Areaid	Specifies the area ID from which the LSA is received. Area ID 0.0.0.0 is the OSPF backbone.
Type	Specifies the type of the link-state advertisement. Each link-state type has a separate advertisement format. Area-scope LSAs unrecognized by the router are also stored in this database.
RouterId	Specifies the originating router in the autonomous system.
Lsid	Specifies the piece of the routing domain that is described by the advertisement.
Sequence	Shows a signed 32-bit integer that detects old and duplicate link-state advertisements. The larger the sequence number, the more recent the advertisement.
Age	Shows the age of the link-state advertisement in seconds.
Checksum	Specifies the checksum of the complete contents of the advertisement, except the age field. The age field is not affected so that the advertisement age value increments without updating the checksum. The checksum used is the same for ISO connectionless datagrams, the Fletcher checksum.

View the Link-Scope LSDB

Before You Begin

- Change the VRF instance as required to view OSPFv3 on a specific VRF instance.

About This Task

View the link-scope LSDB to view the LSAs that are not flooded beyond the local link.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Link-scope LSDB** tab.

Link-scope LSDB Field Descriptions

Use the data in the following table to use the **Link-scope LSDB** tab.

Name	Description
IfIndex	Shows the identifier of the link from which the LSA was received.
Type	Shows the type of the link-state advertisement. Each link state type has a separate advertisement format. Link-scope LSAs not recognized by the router can be stored in the database.
RouterId	Shows the 32 bit number that uniquely identifies the originating router in the autonomous system.
Lsid	Specifies the piece of the routing domain that is being described by the advertisement.
Sequence	Shows a signed 32-bit integer that detects old and duplicate link-state advertisements. The larger the sequence number, the more recent the advertisement.
Age	Shows the age of the link-state advertisement in seconds.
Checksum	Specifies the checksum of the complete contents of the advertisement, except the age field. The age field is not affected so that the advertisement age value increments without updating the checksum. The checksum used is the same for ISO connectionless datagrams, the Fletcher checksum.

Add an NBMA Neighbor

Add an NBMA neighbor for each interface that is eligible to become the DR.

An NBMA interface with a positive nonzero router priority is eligible to become the DR for the NBMA network and is configured with the identification of all attached routers, IPv6 addresses, and router priorities.

Before You Begin

- Identify the following information:
 - specific interfaces to include in the NBMA network
 - the IPv6 address for each interface
 - the router priority for each interface
 - the hello interval for the network

- the router dead interval for the network
- the poll interval for the network
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance.

About This Task

In contrast to a broadcast network where switches multicast (send to AllSPFRouters and AllDRouters) certain OSPF protocol packets, switches replicate and send NBMA packets to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination addresses AllSPFRouters and AllDRouters. Because the NBMA network does not broadcast, you must manually configure a list of neighbors and priorities for all routers in the network that can become the DR. Potential DRs use a positive nonzero router priority.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **NBMA Neighbors** tab.
4. Select **Insert**.
5. Select the IPv6 port or VLAN interface.
6. Specify the IPv6 address for the neighbor.
7. Specify the priority for the neighbor.
8. Select **Insert**.

NBMA Neighbors Field Descriptions

Use the data in the following table to use the **NBMA Neighbors** tab.

Name	Description
IfIndex	Specifies the link ID for the link over which the switch reaches the neighbor.
Address	Specifies the IPv6 address of the neighbor.
Priority	Specifies the priority to use for this neighbor in the designated router election process. A value of 0 indicates the neighbor cannot become the designated router. The higher the priority value, the higher chance the switch will win the election process. The default is 1.

Name	Description
RtrId	Identifies the neighboring router in the autonomous system. The value is 0.0.0.0 until the switch receives a hello message from the neighbor.
State	Identifies the state of the relationship with the neighbor. The state can be one of the following: <ul style="list-style-type: none">• down• attempt• init• twoWay• exchangeStart• exchange• loading• full

Configure Route Redistribution to OSPFv3

Configure a redistribute entry to announce routes of a certain source protocol into the OSPFv3 domain. Optionally, use a route policy to control the redistribution of routes.

Before You Begin

- Enable OSPFv3 globally.
- Ensure that a route policy exists if you intend to use a route policy.
- Change the VRF instance as required to configure OSPFv3 on a specific VRF instance. The VRF must have an RP trigger of OSPFv3. Not all parameters are configurable on non-default VRFs.

About This Task



Important

Changing the OSPFv3 redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform this procedure. As a best practice, change the default preferences for an OSPFv3 redistribute context, you must do so before you enable the protocols.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. Select an option for the route source.
6. Select **enable**.
7. Select a route policy.
8. Configure the metric type.

9. Select **Insert**.

Redistribute *Field Descriptions*

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrflid	Specifies the destination virtual router forwarding instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrflid	Specifies the source VRF instance. You cannot configure this variable.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) an OSPF redistribute entry for a specified source type.
RoutePolicy	Configures the route policy (by name) to use for detailed redistribution of external routes from a specified source into an OSPF domain. Select the ellipsis (...) and choose from the list in the dialog box.
Metric	Configures the OSPF route redistribution metric for basic redistribution. A value of 0 indicates to use the original cost of the route.
MetricType	Configures the OSPF route redistribution metric type. The default is type 2. The cost of a type 2 route is the external cost, regardless of the interior cost. A type 1 cost is the sum of both the internal and external costs.

Modify an OSPFv3 Interface

Configure the OSPF protocol on IPv6 interface to support dynamic routing on the interface. The IPv6 interface can be a tunnel, port, or VLAN.

Before You Begin

- The OSPFv3 interface must exist.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Interfaces** tab.

4. Double-click a cell to edit the value.
5. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
Index	Specifies the interface index for the IPv6 interface on which OSPFv3 is configured.
AreaId	Specifies the area ID to which the IPv6 interface connects. Use 0.0.0.0 for the OSPFv3 backbone.
Type Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the OSPFv3 interface type as one of the following: <ul style="list-style-type: none"> • broadcast • NBMA • point-to-point • point-to-multipoint • passive
AdminStat	Specifies the administrative status for the OSPFv3 interface. If you enable the status, it is advertised as an internal route to some areas. If you disable the status, the interface is external to OSPFv3. The default is enabled.
RtrPriority Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the priority of this interface. Multiaccess networks use the priority in the designated router election. A higher priority value increases the chance the router becomes the designated router. A value of zero (0) indicates the router cannot become the designated router for the network. If more than one router uses the same priority value, the router ID determines the designated router. The default is 1.
TransitDelay Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the estimated number of seconds to transmit a link-state-update packet over this interface. The default is 1.
RetransInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between retransmission of link-state advertisements for the adjacencies that belong to this interface, and for database description and link-state request packets. The default is 5.
HelloInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between the hello packets that the router sends on this interface. You must configure this parameter to the same value for all routers attached to a common network. The default is 10.

Name	Description
RtrDeadInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds after which to declare a router down if no hello packets are received. You must configure this parameter to the same value for all routers attached to a common network. The default is 40.
PollInterval Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the number of seconds between hello packets sent to an inactive NBMA neighbor. The default is 120.
State	Shows the state of the OSPFv3 interface as one of the following: <ul style="list-style-type: none"> • down • loopback • waiting • pointToPoint • designatedRouter • backupDesignatedRouter • otherDesignatedRouter
DesignatedRouter	Shows the router ID for the designated router.
BackupDesignatedRouter	Shows the router ID for the backup designated router.
MetricValue Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies the cost for the interface. The default value for a brouter port or VLAN is 1. The default value for a tunnel is 100. Note: If you do not specify a cost for the interface, the switch dynamically updates the interface cost with the configured global OSPF default cost. The global OSPF default cost depends on the speed of the interface.
LinkLsaSuppression Note: Value is not configurable for OSPFv3 CLIP interfaces.	Specifies whether Link LSA suppression is enabled.

View OSPFv3 Neighbors

About This Task

View OSPFv3 neighbors to see routers with interfaces to a common network.

The OSPFv3 hello protocol maintains and dynamically discovers neighbor relationships.

The exception is an NBMA network; you manually configure permanent neighbors on each router eligible to become the designated router (DR).

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Neighbors** tab.

Neighbors Field Descriptions

Use the data in the following table to use the **Neighbors** tab.

Name	Description
IfIndex	Displays the local-link ID of the link over which the neighbor can be reached.
RtrId	Identifies the neighboring router in the Autonomous System. The value is the router ID of the neighboring router, which in OSPF uses the same format as an IPv6 address but identifies the router independent of IPv6 address.
Address	Displays the IPv6 address for the neighbor associated with the local link.
Options	Displays the bit mask that corresponds to the options field on the neighbor.
State	Displays the state of the relationship with the neighbor. The value can be one of the following: <ul style="list-style-type: none">• down• attempt• init• twoWay• exchangeStart• exchange• loading• full

Name	Description
NbrIfId	Displays the interface ID that the neighbor advertises in its hello packets on this link.
DeadIntCnt	<p>Displays the Dead interval Count or TTL (time to live) field that indicates how many seconds remain before the system declares the Neighbor down.</p> <p>The starting value is the Router Dead Interval value and it decrements to 0 if no Hello is received for that neighbor within the interval. If no Hello is received within the interval, then the system declares the neighbor down.</p> <p>When a hello is received for the neighbor, the system resets the value to the Router Dead Interval value.</p>

View Virtual Neighbors

About This Task

View information about the neighbors on the virtual link to the OSPFv3 backbone.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select the **Virtual Neighbors** tab.

Virtual Neighbors Field Descriptions

Use the data in the following table to use the **Virtual Neighbors** tab.

Name	Description
Area	Shows the ID for the transit area.
RtrId	Shows the ID for the neighboring router in the autonomous system.
LocalIfIndex	Shows the local interface ID for the virtual link over which the switch can reach the neighbor.
AddressType	<p>Shows the type of address as one of the following:</p> <ul style="list-style-type: none"> • ipv4 • ipv6 • ipv4z • ipv6z • dns <p>ipv4z and ipv6z indicate a scope zone.</p>

Name	Description
Address	Shows the IPv6 address that this virtual neighbor advertises. This value must be a global scope address.
Options	Shows a bit mask that corresponds to the OSPF options field of the neighbor.
State	Shows the state of the virtual neighbor relationship. The value can be one of the following: <ul style="list-style-type: none">• down• attempt• init• twoWay• exchangeStart• exchange• loading• full

View OSPFv3 Statistics

View OSPF statistics to analyze trends. You can also graph statistics for all OSPF packets transmitted by the switch.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **OSPFv3**.
3. Select **Stats**.

Stats Field Descriptions

Use the data in the following table to use the **Stats** tab.

Name	Description
TxPackets	Shows the count of sent packets.
RxPackets	Shows the count of received packets.
TxDropPackets	Shows the count of sent, dropped packets.
RxDropPackets	Shows the count of received, dropped packets.
RxBadPackets	Shows the count of received, bad packets.
SpfRuns	Shows the count of intra-area route table updates with calculations using this area link-state database.
LastSpfRun	Shows the count of the most recent SPF run.

Name	Description
LsdbTblSize	Shows the size of the link state database table.
BadLsReqs	Shows the count of bad link requests.
SeqMismatches	Shows the count of sequence mismatched packets.
Routes	Shows the number of OSPF routes added to the routing table.
Adjacencies	Shows the number of existing adjacencies.
Areas	Shows the number of configured areas.
Nbrs	Shows the number of OSPF neighbors.



Port Performance Management

[Digital Diagnostic Monitoring](#) on page 2927

[Port Performance Management Using CLI](#) on page 2927

[Port Performance Management using EDM](#) on page 2933

The topics in this section provide information to configure and view Digital Diagnostic Monitoring (DDM).

Digital Diagnostic Monitoring

Use Digital Diagnostic Monitoring (DDM) to monitor laser operating characteristics such as temperature, voltage, current, and power. This feature works at any time during active laser operation without affecting data traffic.

The following optical transceivers support DDM:

- 1 Gbps Small Form Factor Pluggable (SFP)
- 10 Gbps Small Form Factor Pluggable plus (SFP+)
- 25 Gbps Small Form Factor Pluggable 28 (SFP28)
- 40 Gbps Quad Small Form Factor Pluggable plus (QSFP+)
- 100 Gbps Quad Small Form Factor Pluggable 28 (QSFP28)
- 400 Gbps Quad Small Form-factor Pluggable Double Density (QSFP56-DD)

A Digital Diagnostic Interface (DDI) is an interface that supports DDM. These devices provide real-time monitoring of individual DDI transceivers. The DDM software provides warnings or alarms after the temperature, voltage, laser bias current, transmitter power or receiver power fall outside of vendor-specified thresholds during initialization.

For information about optical transceivers, see [Extreme Optics](#) website.

Port Performance Management Using CLI

This section contains procedures to monitor individual DDI transceivers using the CLI.

Configure DDM

Configure Digital Diagnostic Monitoring to get information concerning the status of the transmitted and received signals to allow better fault isolation and error detection.

About This Task

When you enable DDM, you see the internal performance condition (temperature, voltage, bias, Tx power and Rx power) of the pluggable transceiver. The default is disabled.

For information about how to reset a transceiver for troubleshooting purposes, see [Reset a Transceiver](#) on page 3938.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. (Optional) Enable the port to go down when an alarm occurs:
`pluggable-optical-module ddm-alarm-portdown`
3. (Optional) Configure the DDM interval:
`pluggable-optical-module ddm-monitor-interval <5-60>`
4. (Optional) Enable the sending of trap messages when an alarm occurs:
`pluggable-optical-module ddm-traps-send`
5. Enable DDM:
`pluggable-optical-module ddm-monitor`

Example

Configure the interval to 10 seconds and enable DDM.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#pluggable-optical-module ddm-monitor-interval 10
Switch:1(config)#pluggable-optical-module ddm-monitor
```

Variable Definitions

Use the data in the following table to use the **pluggable-optical-module** command

Variable	Value
<i>ddm-alarm-portdown</i>	When enabled, the port goes down when any alarm occurs. The default is disabled.
<i>ddm-monitor</i>	Enables DDM. When enabled, you see the internal performance condition (temperature, voltage, bias, Tx power and Rx power) of the pluggable transceiver. The default is disabled.

Variable	Value
<i>dmd-monitor-interval</i> <5-60>	Configures the DDM monitor interval. If an alarm occurs, the log message is received within the specific interval. The default value is 5 seconds.
<i>dmd-traps-send</i>	Enables or disables the sending of trap messages. When enabled, the trap message is sent to the device manager any time the alarm occurs. The default is enabled.

View DDI Port Information

Perform this procedure to view basic manufacturing information and characteristics, and the current configuration.

About This Task

This command displays information for DDI transceivers.



Note

Transceiver support differs across hardware platforms. For information about supported parts, see [Extreme Optics](#) website.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View basic manufacturing information and characteristics:
show pluggable-optical-modules basic [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
3. View configuration information:
show pluggable-optical-modules config
4. View detailed manufacturing information and characteristics:
show pluggable-optical-modules detail [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]

Examples

```
Switch:1#show pluggable-optical-modules config

=====
Pluggable Optical Module Global Configuration
=====

      ddm-monitor : disabled
ddm-monitor-interval : 5
      ddm-traps-send : enabled
ddm-alarm-portdown : disabled
```

The following example includes Extreme Networks extended diagnostic information. Not all parts support this information. A value of 254 indicates the part has not been initialized.

```
Switch:1#show pluggable-optical-modules detail 1/29
=====
                        Pluggable Optical Module Info 1/29 Detail
=====
Port: 1/29
Type: 40GbSR4
DDM Supported : TRUE
Vendor Name   : EXTREME NETWORKS      Partnumber : EQPT404SR4VCM100
Vendor REV    : A                    Vendor SN   : 18260006CNFN01
Vendor Date   : 06/26/18
Wavelength    : 850.00 nm

Digital Diagnostic Interface Supported

Optics Status      : Ok
Calibration        : Internal
RX Power Measurement : Average
Auxiliary 1 Monitoring : Not Implemented
Auxiliary 2 Monitoring : Not Implemented

-----
              LOW_ALARM LOW_WARN  ACTUAL  HIGH_WARN HIGH_ALARM THRESHOLD
              THRESHOLD THRESHOLD  VALUE   THRESHOLD THRESHOLD  STATUS
-----
Temp (C)      -5.0      0.0      31.3710   70.0      75.0      Normal
Voltage (V)    2.9700    3.1000    3.2996    3.4650    3.6300    Normal
Tx1Bias (mA)   2.0      3.0      7.6800    13.0      14.0      Normal
Tx2Bias (mA)   2.0      3.0      7.5520    13.0      14.0      Normal
Tx3Bias (mA)   2.0      3.0      7.6800    13.0      14.0      Normal
Tx4Bias (mA)   2.0      3.0      7.4240    13.0      14.0      Normal
Tx1Power (dBm) -11.3000   -7.3000   -2.6000    0.0       3.0      Normal
Tx2Power (dBm) -11.3000   -7.3000   -2.5000    0.0       3.0      Normal
Tx3Power (dBm) -11.3000   -7.3000   -2.9000    0.0       3.0      Normal
Tx4Power (dBm) -11.3000   -7.3000   -2.7000    0.0       3.0      Normal
Rx1Power (dBm) -13.9000   -9.9000   -4.1000    0.0       3.0      Normal
Rx2Power (dBm) -13.9000   -9.9000   -3.4000    0.0       3.0      Normal
Rx3Power (dBm) -13.9000   -9.9000   -2.9000    0.0       3.0      Normal
Rx4Power (dBm) -13.9000   -9.9000   -3.2000    0.0       3.0      Normal

Extreme Extended Diagnostic Information

-----
Power On Counter (48 hours) : 5
Tx1 DDM Initial (dBm)       : 7
Tx1 DDM Last Gasp (dBm)     : 7
Tx2 DDM Initial (dBm)       : 7
Tx2 DDM Last Gasp (dBm)     : 7
Tx3 DDM Initial (dBm)       : 7
Tx3 DDM Last Gasp (dBm)     : 7
Tx4 DDM Initial (dBm)       : 6
Tx4 DDM Last Gasp (dBm)     : 6
Rx1 DDM Initial (dBm)       : 254
Rx1 DDM Last Gasp (dBm)     : 254
Rx2 DDM Initial (dBm)       : 254
Rx2 DDM Last Gasp (dBm)     : 254
```


Rx3 DDM Initial (dBm)	: 254
Rx3 DDM Last Gasp (dBm)	: 254
Rx4 DDM Initial (dBm)	: 254
Rx4 DDM Last Gasp (dBm)	: 254

Variable Definitions

Use the data in the following table to use the **show pluggable-optical-modules basic** and **show pluggable-optical-modules detail** commands.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

View DDI Temperature Information

About This Task

This command displays information for DDI transceivers.



Note Transceiver support differs across hardware platforms. For information about supported parts, see [Extreme Optics](#) website.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View temperatures:
`show pluggable-optical-modules temperature [{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}`

Example

Switch:1#show pluggable-optical-modules temperature							
=====							
Pluggable Optical Module Temperature (C)							
=====							
PORT	LOW_ALARM	LOW_WARN	ACTUAL	HIGH_WARN	HIGH_ALARM	THRESHOLD	
NUM	THRESHOLD	THRESHOLD	VALUE	THRESHOLD	THRESHOLD	STATUS	

1/2	7.0	1.1250	65.2539	0.0	3.0156	Low Alarm
1/3	7.0	1.1250	65.2539	0.0	3.0156	Low Alarm
1/9	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
1/15	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
2/1	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
2/2	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm
2/3	7.0625	0.0	65.2539	0.0	3.0156	Low Alarm

Variable Definitions

Use the data in the following table to use the **show pluggable-optical-modules temperature** command.

Variable	Value
<code>{slot/port[/sub-port]}</code> <code>[-slot/port[/sub-port]]</code> <code>[,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View DDI Voltage Information

About This Task

This command displays information for DDI transceivers.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View voltages:

```
show pluggable-optical-modules voltage [{slot/port[/sub-port]}[-slot/  
port[/sub-port]][,...][slot/all][all]]
```

Example

```
Switch:1#show pluggable-optical-modules voltage
```

```
=====
                        Pluggable Optical Module Voltage (V)
=====
```

PORT NUM	LOW_ALARM THRESHOLD	LOW_WARN THRESHOLD	ACTUAL VALUE	HIGH_WARN THRESHOLD	HIGH_ALARM THRESHOLD	THRESHOLD STATUS
1/2	0.1281	0.0	1.2596	0.5376	1.6396	Normal
1/3	0.0001	0.0	1.2596	0.3072	1.6396	Normal
1/9	0.0006	0.0	1.2596	2.6368	0.0	Normal
1/15	0.0006	0.0	1.2596	2.6368	0.0	Normal
2/1	0.0006	0.0	1.2596	2.6368	0.0	Normal

2/2	0.0006	0.0	1.2596	2.6368	0.0	Normal
2/3	0.0006	0.0	1.2596	2.6368	0.0	Normal

Variable Definitions

Use the data in the following table to use the **show pluggable-optical-modules voltage** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Port Performance Management using EDM

This section contains procedures to monitor individual DDI transceivers using EDM.

View DDI Information

About This Task

You can view Digital Diagnostic Interface (DDI) information, for example, port information, temperature, and voltages for DDI transceivers.



Note

Transceiver support differs across hardware platforms. For information about supported parts, see the [Extreme Optics](#) website.

Procedure

1. On the Physical Device view, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **DDI/SFP** tab.

DDI/SFP Field Descriptions

Use the data in the following table to use the **DDI/SFP** tab.

Name	Description
ConnectorType	Specifies the type of connector.
SupportsDDM	Specifies if the transceiver supports DDM.
DdmStatusMask	Specifies the DDM status. A value other than ddm-ok represents a specific error.
CLEI	Specifies the Telcordia register assignment CLEI code.
VendorName	Specifies the name of the manufacturer.
VendorPartNumber	Specifies the part number for the transceiver. Select the part number to access the technical specifications for the part.
VendorRevNumber	Specifies the manufacturer revision level for the transceiver.
VendorSN	Specifies the manufacturer serial number for the transceiver.
VendorDateCode	Specifies the manufacturer date code for the transceiver.
Wavelength	Specifies the wavelength in nm. This field is valid for optical transceivers only.
Calibration	Specifies if the calibration is internal or external.
PowerMeasure	Specifies Rx power measurement as average or OMA.
Aux1Monitoring	Specifies if auxiliary monitoring is implemented for the transceiver.
Aux2Monitoring	Specifies if auxiliary monitoring is implemented for the transceiver.
TemperatureLowAlarmThreshold	Specifies the low alarm threshold in degrees Celsius.
TemperatureLowWarningThreshold	Specifies the low warning threshold in degrees Celsius.
Temperature	Specifies the current temperature in degrees Celsius of the transceiver.
TemperatureHighWarningThreshold	Specifies the high warning threshold in degrees Celsius.
TemperatureHighAlarmThreshold	Specifies the high alarm threshold in degrees Celsius.
TemperatureStatus	Specifies if any temperature thresholds were exceeded.
VoltageLowAlarmThreshold	Specifies the low alarm threshold in volts.
VoltageLowWarningThreshold	Specifies the low warning threshold in volts.

Name	Description
Voltage	Specifies the current voltage in volts.
VoltageHighWarningThreshold	Specifies the high warning threshold in volts.
VoltageHighAlarmThreshold	Specifies the high alarm threshold in volts.
VoltageStatus	Specifies if any voltage thresholds were exceeded.
BiasLowAlarmThreshold	Specifies the bias current low alarm threshold in mA.
BiasLowWarningThreshold	Specifies the bias current low warning threshold in mA.
Bias	Specifies the laser bias current in mA.
BiasHighWarningThreshold	Specifies the bias current high warning threshold in mA.
BiasHighAlarmThreshold	Specifies the bias current high alarm threshold in mA.
BiasStatus	Specifies if any bias thresholds were exceeded.
TxPowerLowAlarmThreshold	Specifies the low alarm threshold in dBm for the Tx power.
TxPowerLowWarningThreshold	Specifies the low warning threshold in dBm for the Tx power.
TxPower	Specifies the current Tx power in dBm.
TxPowerHighWarningThreshold	Specifies the high warning threshold in dBm for the Tx power.
TxPowerHighAlarmThreshold	Specifies the high alarm threshold in dBm for the Tx power.
TxPowerStatus	Specifies if any Tx power thresholds were exceeded.
RxPowerLowAlarmThreshold	Specifies the low alarm threshold in dBm for the Rx power.
RxPowerLowWarningThreshold	Specifies the low warning threshold in dBm for the Rx power.
RxPower	Specifies the current Rx power in dBm.
RxPowerHighWarningThreshold	Specifies the high warning threshold in dBm for the Rx power.
RxPowerHighAlarmThreshold	Specifies the high alarm threshold in dBm for the Rx power.
RxPowerStatus	Specifies if any Rx power thresholds were exceeded.
Aux1LowAlarmThreshold	Specifies the low alarm threshold auxiliary 1 reading.
Aux1LowWarningThreshold	Specifies the low warning threshold auxiliary 1 reading.

Name	Description
Aux1	Specifies the current auxiliary 1 reading.
Aux1HighWarningThreshold	Specifies the high warning threshold auxiliary 1 reading.
Aux1HighAlarmThreshold	Specifies the high alarm threshold auxiliary 1 reading.
Aux1Status	Specifies if any auxiliary 1 thresholds were exceeded.
Aux2LowAlarmThreshold	Specifies the low alarm threshold auxiliary 2 reading.
Aux2LowWarningThreshold	Specifies the low warning threshold auxiliary 2 reading.
Aux2	Specifies the current auxiliary 2 reading.
Aux2HighWarningThreshold	Specifies the high warning threshold auxiliary 2 reading.
Aux2HighAlarmThreshold	Specifies the high alarm threshold auxiliary 2 reading.
Aux2Status	Specifies if any auxiliary 2 thresholds were exceeded.
ExtremeExtraFeatures	Specifies if the device supports the extra Extreme Networks features. If the value is true, you can select Extended DDM to view more information.



Note

1. Threshold and actual values for TxBias, TxPower, and RxPower are provided for all four channels in SFP56-DD, QSFP+, and QSFP28 optical transceivers.
2. Auxiliary monitoring does not apply to SFP56-DD, QSFP+ or QSFP28.

View Extended DDI Information

About This Task

You can view extended information for DDI transceivers that support extra Extreme Networks features.



Note

Transceiver support differs across hardware platforms. For information about supported parts, see the [Extreme Optics](#) website.

Procedure

1. On the Physical Device view, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.

4. Select the **DDI/SFP** tab.
5. Select **Extended DDM**.
6. Select the **Extended Diagnostic** tab.

Extended Diagnostic *Field Descriptions*

Use the data in the following table to use the **Extended Diagnostic** tab.

Name	Description
PowerOnCounter	Tracks the power-on-life of the part. This value increments by 1 every 48 hours of consecutive power. This value is never decremented or cleared. If you remove or reinsert the part, reset or reboot the chassis, reset the slot, or channelize or dechannelize the port, then the 48-hour time period restarts.
TxDdmInitial	Specifies the Tx dB from low alarm, which is the difference between the TxPower dBm and the TxPower low alarm. The host provides the value after the first 48 consecutive hours of operation.
TxDdmLastGasp	Specifies the last gasp of Tx dB from low alarm. The host updates this value every 48 hours.
RxDdmInitial	Specifies the Rx dB from low alarm, which is the difference between the RxPower dBm and the RxPower low alarm. The host provides the value after the first 48 consecutive hours of operation after a power cycle.
RxDdmLastGasp	Specifies the last gasp of Rx dB from low alarm. The host updates this value every 48 hours.



Power over Ethernet

[PoE Overview](#) on page 2944

[PoE Detection Types](#) on page 2945

[Power Usage Threshold](#) on page 2946

[Port Power Limit](#) on page 2947

[Port Power Priority](#) on page 2947

[PoE/PoE+ Allocation Using LLDP](#) on page 2949

[Fast PoE and Perpetual PoE](#) on page 2950

[PoE Ports in EDM](#) on page 2950

[Power over Ethernet Configuration using CLI](#) on page 2951

[Power over Ethernet configuration using EDM](#) on page 2961

Table 220: Power over Ethernet product support

Feature	Product	Release introduced
Power over Ethernet (PoE)	4220 Series	Fabric Engine 9.2 Not supported on 4220-12T-4X, 4220-24T-4X, or 4220-8X
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-8XE, 5320-48T-8XE, 5320-24T-4X-XT, or 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4 Not supported on: <ul style="list-style-type: none"> • 5420F-24T-4XE • 5420F-48T-4XE • 5420F-24S-4XE • 5420M-24T-4YE • 5420M-48T-4YE
	5520 Series	VOSS 8.2.5 Supported on: <ul style="list-style-type: none"> • 5520-12MW-36W • 5520-24W • 5520-48W
	5720 Series	Fabric Engine 8.7 Not supported on 5720-VIM-2CE or 5720-VIM-6YE
	7520 Series	Not Applicable
	7720 Series	Not Applicable
	7830 Series	Not Applicable
	VSP 4900 Series	VOSS 8.1 Supported on all fixed ports on VSP4900-48P and ports 1/1-1/12 on VSP4900-12MXU-12XE
	VSP 7400 Series	Not Applicable

Table 220: Power over Ethernet product support (continued)

Feature	Product	Release introduced
PoE/PoE+ allocation using LLDP	4220 Series	Not Applicable
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-8XE or 5320-48T-8XE
	5420 Series	VOSS 8.4 Not supported on: <ul style="list-style-type: none"> • 5420F-24T-4XE • 5420F-48T-4XE • 5420F-24S-4XE • 5420M-24T-4YE • 5420M-48T-4YE
	5520 Series	VOSS 8.2.5 Supported on: <ul style="list-style-type: none"> • 5520-12MW-36W • 5520-24W • 5520-48W
	5720 Series	Fabric Engine 8.7 Not supported on 5720-VIM-2CE or 5720-VIM-6YE
	7520 Series	Not Applicable
	7720 Series	Not Applicable
	7830 Series	Not Applicable
	VSP 4900 Series	VOSS 8.1 Supported on all fixed ports on VSP4900-48P and ports 1/1-1/12 on VSP4900-12MXU-12XE
	VSP 7400 Series	Not Applicable

Table 220: Power over Ethernet product support (continued)

Feature	Product	Release introduced
Fast PoE	4220 Series	Fabric Engine 9.2 Not supported on 4220-12T-4X, 4220-24T-4X, or 4220-8X
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-8XE or 5320-48T-8XE
	5420 Series	VOSS 8.4 Not supported on: <ul style="list-style-type: none"> • 5420F-24T-4XE • 5420F-48T-4XE • 5420F-24S-4XE • 5420M-24T-4YE • 5420M-48T-4YE
	5520 Series	VOSS 8.2.5 Supported on: <ul style="list-style-type: none"> • 5520-12MW-36W • 5520-24W • 5520-48W
	5720 Series	Fabric Engine 8.7 Not supported on 5720- VIM-2CE or 5720-VIM-6YE
	7520 Series	Not Applicable
	7720 Series	Not Applicable
	7830 Series	Not Applicable
	VSP 4900 Series	VOSS 8.1 Supported on all fixed ports on VSP4900-48P and ports 1/1-1/12 on VSP4900-12MXU-12XE
	VSP 7400 Series	Not Applicable

Table 220: Power over Ethernet product support (continued)

Feature	Product	Release introduced
Perpetual PoE	4220 Series	Fabric Engine 9.2 Not supported on 4220-12T-4X, 4220-24T-4X, or 4220-8X
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-8XE or 5320-48T-8XE
	5420 Series	VOSS 8.4 Not supported on: <ul style="list-style-type: none"> • 5420F-24T-4XE • 5420F-48T-4XE • 5420F-24S-4XE • 5420M-24T-4YE • 5420M-48T-4YE
	5520 Series	VOSS 8.2.5 Supported on: <ul style="list-style-type: none"> • 5520-12MW-36W • 5520-24W • 5520-48W
	5720 Series	Fabric Engine 8.7 Not supported on 5720-VIM-2CE or 5720-VIM-6YE
	7520 Series	Not Applicable
	7720 Series	Not Applicable
	7830 Series	Not Applicable
	VSP 4900 Series	VOSS 8.1 Supported on all fixed ports on VSP4900-48P and ports 1/1-1/12 on VSP4900-12MXU-12XE
	VSP 7400 Series	Not Applicable

Table 220: Power over Ethernet product support (continued)

Feature	Product	Release introduced
Ability to configure the Detection Type per port	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Fabric Engine 9.2 Not supported on: <ul style="list-style-type: none"> • 5420F-24T-4XE • 5420F-48T-4XE • 5420F-24S-4XE • 5420M-24T-4YE • 5420M-48T-4YE
	5520 Series	Fabric Engine 9.2 Supported on: <ul style="list-style-type: none"> • 5520-12MW-36W • 5520-24W • 5520-48W
	5720 Series	Not Supported
	7520 Series	Not Applicable
	7720 Series	Not Applicable
	7830 Series	Not Applicable
	VSP 4900 Series	Not Supported
	VSP 7400 Series	Not Supported

Power over Ethernet (PoE) is the implementation of IEEE 802.3af, IEEE 802.3at, and IEEE 802.3bt (Type 3 and Type 4), which allows for both data and power to pass over a copper Ethernet LAN cable. Typical power devices include wireless Access Points and VoIP telephones.

Depending on the technology and application requirements, PoE is classified into classes. Depending on the power requirements, the PoE devices are categorized by type. Classes range from Class 0 to 8 whereas types range from Type 1 to 4. Each type associates with an IEEE 802.3 PoE standard. These standards provide signaling between the power sourcing equipment (PSE) and the powered device (PD). PSE devices, such as switches, provide power on the network cable. The devices that PSE provides power to are called PDs, such as VoIP phones, wireless access points, and IP surveillance cameras.

To know which ports support PoE, see the hardware documentation.

The switch uses the Dynamic Power Allocation scheme when supplying power to devices. Only the power being consumed by the device is allocated, improving efficiency and enabling support for more number of devices.

You can configure PoE from CLI and Enterprise Device Manager (EDM).

PoE Overview

You can plug any IEEE 802.3af-compliant, 802.3at-compliant, or 802.3bt-compliant (Type 3 and Type 4) for PWR+ powered device into a front-panel port and receive power in that port. Data also can pass simultaneously on that port. This capability is called PoE.

The second power supply unit (PSU) in the switch adds budget to the total POE budget. See the product data sheet for information about the POE budget based on PSU selection.

The IEEE 802.3af draft standard regulates a maximum of 15.4 W of power for each port; that is, a power device cannot request more than 15.4 W of power. As different network devices require different levels of power, the overall available power budget of the switch depends on your power configuration and the particular connected network devices. If you connect an IP device that requires more than 16 W of power, you see an error on that port notifying you of an overload.

The switch automatically detects each IEEE 802.3af-compliant powered device attached to each front-panel port and immediately sends power to that appliance. The switch also automatically detects how much power each device requires and supply the required DC voltage at a set current based on the load conditions and current availability. The switch supports both PoE and standard LAN devices.

Similarly, the switch automatically detects any IEEE 802.3at-compliant (maximum 25.5 W) or IEEE 802.3bt-compliant (maximum 60W for Type 3, maximum 90W for Type 4) powered device attached to any PoE front panel port and immediately sends power to that appliance.

The power detection function of the switch operates independently of the data link status. A device that is already operating the link for data or a device that is not yet operational can request power. That is, the switch provides power to a requesting device even if the data link for that port is disabled. The switch monitors the connection and automatically disconnects power from a port when you remove or change the device, as well as when a short occurs.

The switch automatically detects devices that require no power connections from them, such as laptop computers or other switching devices, and sends no power to those devices. You control the supply of power to specific ports by setting the maximum allowed power to each port in 1 W increments, from 3 W to 98W.

**Important**

Wait for 30 seconds between unplugging and replugging an IP device to the switch to enable the IP device to discharge. If you attempt to connect earlier, the switch cannot detect the IP device.

The switch provides the capability to configure a PoE power threshold, which lets you specify a percentage of the total PoE power usage at which the switch sends a

warning message. If the power consumption is below the threshold, the switch logs an information message.

**Important**

- After you enable Fast PoE or Perpetual PoE or both, you must save the running configuration file.
- If Fast PoE or Perpetual PoE are enabled and you change any other global or port-specific PoE settings (for example, power limits, or port power priority), you must resave the running configuration file.

PoE Detection Types

The configured detection type specifies the following versions of the IEEE standard to support. Standards are arranged from oldest (top) to newest (bottom).

Detection Type	Power Mode
802.3af	Normal
802.3af and legacy	Normal
802.3at	High
802.3at and legacy	High
802.3bt type 3	Normal
802.3bt type 4	Normal
802.3bt and legacy	Normal

The default detection type is automatically configured for the highest PoE standard supported on the switch.

**Note**

802.3at is backwards compatible with 802.3af. Therefore, both normal power and high power devices are supported in this mode.

Changing from a newer IEEE standard to an older IEEE standard is permitted, however power delivery is interrupted during this operation, and all PoE devices are reset. There is no service interruption when changing from an older IEEE standard to a newer IEEE standard.

The following table identifies product support for the different PoE types.

Table 221: PoE type support

Product	Highest PoE Standard
4220-12P-4X 4220-24P-4X	802.3at Type 2 PoE (30W) on BASE-T ports
4220-4MW-8P-4X 4220-4MW-20P-4X	802.3at Type 2 PoE (30W) on 10/100/1000 BASE-T ports 802.3bt Type 4 PoE (90W) on 4 1G/2.5G/5G ports
5320-16P-2MXT-2X 5320-16P-4XE 5320-16P-4XE-DC 5320-24P-8XE 5320-48P-8XE	802.3at Type 2 PoE (30W) on BASE-T ports
5420F-48P-4XE 5420F-48P-4XL 5420F-24P-4XE	802.3at Type 2 PoE (30W)
5420F-8W-16P-4XE	802.3bt Type 4 PoE (90W) on 8 ports 802.3at Type 4 PoE (30W) on 16 ports
5420F-16W-32P-4XE 5420F-16MW-32P-4XE	802.3bt Type 4 PoE (90W) on 16 ports 802.3at Type 4 PoE (30W) on 32 ports
5420M-24W-4YE 5420M-48W-4YE	802.3bt Type 4 PoE (90W)
5420M-16MW-32P-4YE	802.3bt Type 4 PoE (90W) on 16 ports 802.3at Type 4 PoE (30W) on 32 ports
5520-12MW-36W	802.3bt Type 4 PoE (90W)
5520-24W	802.3bt Type 4 PoE (90W)
5520-48W	802.3bt Type 4 PoE (90W)
5720-24MW 5720-48MW 5720-24MXW 5720-48MXW	802.3bt Type 4 PoE (90W)

Power Usage Threshold

The power usage threshold is a chassis configurable percent of the total power available on the switch. When the POE power consumption exceeds this threshold, a log message is generated to warn such an event. When power consumption transitions below the threshold, an informational log message is logged. The default threshold is 80%.

Port Power Limit

Each PoE port has a configurable power limit. This configuration attribute limits the amount of power supplied on a particular port and varies across different hardware platforms. If a PD requires more than the configured limit, the device will not connect properly or is forced to run at a lower limit.

The following table lists the power limit for different hardware platforms:

Table 222: Power Limits

Platform	Power Limit
4220 Series	32 watts: <ul style="list-style-type: none"> • 4220-12P-4X • 4220-24P-4X 98 watts: <ul style="list-style-type: none"> • 4220-4MW-8P-4X: 4MW ports only; 8P ports support 32 watts • 4220-4MW-20P-4X: 4MW ports only; 20P ports support 32 watts
5320 Series	32 watts
5420 Series	32 watts: <ul style="list-style-type: none"> • 5420F-24P-4XE • 5420F-48P-4XE • 5420F-48P-4XL 98 watts: <ul style="list-style-type: none"> • 5420F-8W-16P-4XE: 8W ports only; 16P ports support 32 watts • 5420F-16MW-32P-4XE: 16MW ports only; 32P ports support 32 watts • 5420F-16W-32P-4XE: 16W ports only; 32P ports support 32 watts • 5420M-16MW-32P-4YE: 16MW ports only; 32P ports support 32 watts • 5420M-24W-4YE • 5420M-48W-4YE
5520 Series	98 watts
5720 Series	98 watts

Port Power Priority

You can configure the power priority of each port by choosing low, high, or critical power priority settings.

The switch automatically drops low-priority ports when the power requirements exceed the available power budget. When the power requirements become lower than the switch power budget, the power returns to the dropped port. When several ports

have the same priority and the power budget is exceeded, the ports with the highest interface number are dropped until the consumption is within the power budget.

The priority methods are:

1. Port configured PoE priority

- Low: (default) standard priority for standard devices
- High: higher priority than low for important devices
- Critical: highest priority for critical devices like wireless APs

2. Port number priority where the lower port numbers have a higher priority.

PD Classification

The PDs are classified during initial connection establishment as defined in IEEE 802.3 standards. The classification defines the amount of power a device is expected to consume.

Table 223: Classification Chart for Various Standards

Name	Standard	Class	Type	PSE Min Output Power (in W)	PD Min Input Power (in W)	Example of Supported PDs
PoE	IEEE 802.3af	0	1	15.4	12.95	Static surveillance camera, VoIP phones, wireless access points
		1	1	4	3.84	IP phones
		2	1	7	6.49	IP camera
		3	1	15.4	12.95	Wireless access points
PoE+	IEEE 802.3at	4	2	30	25.5	High power PD
PoE++	IEEE 802.3bt	5	3	45	40	Video conferencing equipment, multi-radio wireless access points
		6	3	60	51	

Table 223: Classification Chart for Various Standards (continued)

Name	Standard	Class	Type	PSE Min Output Power (in W)	PD Min Input Power (in W)	Example of Supported PDs
PoE++	IEEE 802.3bt	7	4	75	62	Laptops, flat screens
		8	4	90	73	

PoE/PoE+ Allocation Using LLDP

Power over Ethernet/Power over Ethernet Plus allocation using Link Layer Discovery Protocol (LLDP) supports Ethernet switches, which do not support hardware-level power negotiation. With this feature, these switches support IEEE-based PoE and play the role of power sourcing equipment (PSE).

The devices that are powered using PoE, such as IP Phone and Video Surveillance Cameras, are classified as Powered Devices (PD). The maximum allowed continuous output power per cable for each standard is:

- IEEE 802.3af: 15.4 W
- IEEE 802.3at: 25.5 W
- IEEE 802.3bt Type 3: 60 W
- IEEE 802.3bt Type 4: 90 W

The negotiation of actual power supply and demand between a PSE and a PD can be executed at either the physical layer or at the data link layer. After the link is established at the physical layer, the PSE can use the IEEE 802.1AB LLDP protocol to repeatedly query the PD to discover its power needs. Communication using LLDP allows for a finer control of power allocation, making it possible for the PSE to dynamically supply the exact power levels needed by individual PDs, and globally for all PDs that are attached. Using LLDP is optional for the PSE, however, it is mandatory for a Type 2 PD that requires more than 12.95 watts of power.



Important

LLDP supports PoE discovery and power allocation because some switches do not support hardware-level power negotiation. This allows Type 2 PDs such as PTZ (pan-tilt-zoom) Video Surveillance Cameras to be fully functional when connected to one of these switches. This functionality is enabled by default and is not configurable.



Note

Some switches feature a hardware design that supports hardware-level detection. Therefore, they do not require LLDP.

Fast PoE and Perpetual PoE

Fast PoE minimizes the PoE controller recovery time in case of a power failure. With Fast PoE, the PoE controller initializes the moment the switch powers on, which results in a faster recovery period.

Perpetual PoE provides uninterrupted power to all connected devices during a switch reboot.



Important

- After you enable Fast PoE or Perpetual PoE or both, you must save the running configuration file.
- If Fast PoE or Perpetual PoE are enabled and you change global or port-specific PoE configuration, you must save the running configuration file.

PoE Ports in EDM

The front panel view of Enterprise Device Manager (EDM) provides additional information for Power over Ethernet (PoE) ports for switches that support PoE.

This additional information is in the form of a colored **P** that displays inside the graphical representation of the port. This colored P represents the current power aspect of the PoE port.

Table 224: Power Aspect color codes

Color	Description
Green	The port is currently delivering power.
Red	The power and detection mechanism for the port is disabled.
Orange	The power and detection mechanism for the port is enabled. The port is not currently delivering power.
White/Gray	The power and detection mechanism for the port is unknown.



Important

The data and power aspect coloring schemes are independent of each other. You can view the initial status for both data and power aspect for the port. To refresh the power status, right-click the unit, and then select **Refresh Port Tooltips** from the shortcut menu.

Power over Ethernet Configuration using CLI

This section provides details to configure PoE settings using CLI.



Important

- After you enable Fast PoE or Perpetual PoE or both, you must save the running configuration file.
- If Fast PoE or Perpetual PoE are enabled and you change any other global or port-specific PoE settings, you must resave the running configuration file.

Disable PoE on a Port

About This Task

Perform the following procedure to disable PoE on a port. The Ethernet connected device does not receive any power over Ethernet if you shutdown PoE on the port.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Disable PoE on the port:

```
poe poe-shutdown [port <portlist>]
```

<portlist> is the port on which you want to disable PoE. The default is `enable`.


What to Do Next

To return power to the port, enter `no poe-shutdown [port <portlist>]`.

Configure the PoE Detection Type Globally

Perform the following procedure to configure the PoE powered device (PD) detection type globally on the switch. The default is automatically configured for the highest PoE

standard supported on the switch. For more information, see [PoE Detection Types](#) on page 2945.



Important

Changing from a newer IEEE standard to an older IEEE standard is permitted, however power delivery is interrupted during this operation, and all PoE devices are reset. There is no service interruption when changing from an older IEEE standard to a newer IEEE standard.

Procedure

1. Enter Global Configuration mode:
- ```
enable

configure terminal
```
2. Configure PoE detection type:
- ```
poe poe-pd-detect-type {802dot3af | 802dot3af_and_legacy | 802dot3at |
802dot3at_and_legacy | 802dot3bt_type3 | 802dot3bt_type4}
```

Variable Definitions

The following table defines parameters for the **poe-pd-detect-type** command.

Variable	Value
<code>{802dot3af 802dot3af_and_legacy 802dot3at 802dot3at_and_legacy 802dot3bt_type3 802dot3bt_type4}</code>	<div>Configures the detection type to one of the following values:</div> <ul style="list-style-type: none">802dot3af: Set PD detection mode in 802.3af802dot3af_and_legacy: Set PD detection mode in 802.3af and legacy802dot3at: Set PD detection mode in 802.3at802dot3at_and_legacy: Set PD detection mode in 802.3at and legacy802dot3bt_type3: Set PD detection mode in 802dot3bt Type3802dot3bt_type4: Set PD detection mode in 802dot3bt Type4 <div>Note: The default detection type is automatically configured for the highest PoE standard supported on the switch.</div>

Configure the PoE Detection Type on a Port

Perform the following procedure to configure the PoE powered device (PD) detection type on a specific port or ports.

About This Task

**Note**

This procedure applies to 5420 Series and 5520 Series.

Ensure the configuration aligns with hardware capabilities. The default is automatically configured for the highest PoE standard supported on the switch. For more information, see [PoE Detection Types](#) on page 2945.

Port-level configuration takes precedence over a global PoE PD detection type configuration.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the PD detection type:

```
poe poe-pd-detect-type [{slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]] {802dot3af | 802dot3af_and_legacy | 802dot3at  
| 802dot3at_and_legacy | 802dot3bt_type3 | 802dot3bt_type4}
```

Variable Definitions

The following table defines parameters for the **poe poe-pd-detect-type** command.

Variable	Value
{802dot3af 802dot3af_and_legacy 802dot3at 802dot3at_and_legacy 802dot3bt_type3 802dot3bt_type4}	<p>Configures the detection type to one of the following values:</p> <ul style="list-style-type: none"> 802dot3af: Specifies the PD detection mode as 802.3af 802dot3af_and_legacy: Specifies the PD detection mode as 802.3af and legacy 802dot3at: Specifies the PD detection mode as 802.3at 802dot3at_and_legacy: Specifies the PD detection mode as 802.3at and legacy 802dot3bt_type3: Specifies the PD detection mode as 802dot3bt Type3 802dot3bt_type4: Specifies the PD detection mode as 802dot3bt Type4 <p>Note: The default detection type is automatically configured for the highest PoE standard supported on the switch.</p>
port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure PoE Power Usage Threshold

About This Task

Perform the following procedure to configure the PoE power usage threshold limit globally as a percentage on the switch. The switch logs a warning message when a PoE PD power usage exceeds the configured threshold. The switch logs an informational message when a PoE PD power usage is below the configured threshold.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the power usage threshold:
`poe poe-power-usage-threshold <1-99>.`

Variable Definitions

The following table defines parameters for the **poe-power-usage-threshold** command.

Variable	Value
<1-99>	Specifies the PoE usage threshold in the range of 1—99 percent.

Configure Power Limits for Channels

About This Task

Perform the following procedure to configure the PoE power limit for specific ports or channels. You can limit the PoE wattage available from an individual port or list of ports.

Procedure

1. Enter Interface Configuration mode:
`enable`
`configure terminal`
`interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]`
`[,...][slot/all][all]} or interface vlan <1-4059>`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure PoE channel limits:
`poe poe-limit [port <portlist>] <power_limit>`

Variable definitions

The following table defines parameters for the **poe-limit** command.

Variable	Value
<code><portlist></code>	Identifies the ports on which the limit is set.
<code><power_limit></code>	Specifies the configurable power limit, in watts on a particular port. To see the available range for the switch, use the CLI Help.

Configure Port Power Priority

About This Task

Perform the following procedure to configure the PoE power priority for a port or list of ports. You can configure the PoE power priority of ports to manage availability of the connected PDs. If the switch needs to shut down PDs because PoE exceeds the power limit threshold, low priority devices are shut down before high priority, and high priority are shut down before critical.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure port power priority:

```
poe poe-priority [port <portlist>] {critical| high| low}
```

Variable Definitions

The following table defines parameters for the **poe-priority** command.

Variable	Value
<code><portlist></code>	Identifies the ports to set priority for.
<code>{low high critical}</code>	Identifies the PoE priority.

Enable Fast PoE Globally

About This Task

Perform this procedure to enable Fast PoE on the switch. After you enable Fast PoE, you must save the running configuration file.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable Fast PoE:

```
poe fast-poe-enable
```

3. Save the configuration file:

```
save config
```

Enable Perpetual PoE Globally

About This Task

Perform this procedure to enable Perpetual PoE on the switch. After you enable Perpetual PoE, you must save the running configuration file.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable Perpetual PoE:

```
poe perpetual-poe-enable
```

3. Save the configuration file:

```
save config
```

Enable Fast PoE on a Port

About This Task

Perform this procedure to enable Fast PoE on a specific copper port of the switch. After you enable Fast PoE, you must save the running configuration file.

Procedure

- 1. Enter Interface Configuration mode:
`enable`

`configure terminal`

`interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>`



Note If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- 2. Enable Fast PoE on the copper port:
`poe fast-poe-enable [port {slot/port[-slot/port] [,...][slot/all]
[all]]}`
- 3. Save the configuration file:
`save config`

Variable Definitions

The following table defines parameters for the **fast-poe-enable** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/ all][all]}</code>	Specifies the port or ports to be configured.

Enable Perpetual PoE on a Port

About This Task

Perform this procedure to enable Perpetual PoE on a specific copper port of the switch. After you enable Perpetual PoE, you must save the running configuration file.

Procedure

- 1. Enter Interface Configuration mode:
`enable`

`configure terminal`

`interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>`



Note If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable Fast PoE on the copper port:

```
poe perpetual-poe-enable [port {slot/port[-slot/port]][,...][slot/all]
[all]]
```

3. Save the configuration file:

```
save config
```

Variable Definitions

The following table defines parameters for the **fast-poe-enable** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/ all][all]}</code>	Specifies the port or ports to be configured.

Display Global PoE Configuration

About This Task

Perform the following procedure to display the global PoE configuration. You can view the global PoE status, power consumption, power limit threshold, and more.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the global configuration:
show poe-main-status

Example

```
Switch:1#show poe-main-status
=====
PoE Main Status - Stand-alone
=====
Available DTE Power      : 1855 Watts
DTE Power Status         : NORMAL
DTE Power Consumption    : 92 Watts
DTE Power Usage Threshold : 80
PD Detect Type           : 802.3at and Legacy
Power Source Present     : AC Only
Primary Power Status     : Present and operational
Redundant Power Status   : Present and Operational
Fast POE Status          : Enabled
Perpetual POE Status     : Enabled
POE Firmware Version:    : 3.0.0.6
```

Display PoE Port Status

About This Task

Perform the following procedure to display the PoE status for each port. You can use this information to view the status, classification, watts, and priority for each PoE port.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the port status:
show poe-port-status

Example

```
Switch:1#show poe-port-status
```

=====									
POE Port Status									
=====									
PORT	ADMIN STATUS	CURRENT STATUS	CLASSIFICATION	LIMIT (Watts)	PRIORITY	FAST POE	PERPETUAL POE	PD DETECTION TYPE	

1/1	Enable	DeliveringPower	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/2	Enable	DeliveringPower	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/3	Enable	DeliveringPower	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/4	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/5	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/6	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/7	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/8	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/9	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	
1/10	Enable	Searching	Class0	32	Low	Disable	Disable	802.3at and Legacy	



Note

The PoE status of all ports is displayed. The preceding output is a sample of the full output.

Display Port Power Measurement

About This Task

Perform the following procedure to display the PoE power measurement. You can view the voltage, amperage, and wattage for every PoE port. PoE ports without a PD in use are measured as zeros.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View measurement information:
show poe-power-measurement

Example

```
Switch:1#show poe-power-measurement
```

=====			
POE Port Measurement			
=====			
PORT	Volt (V)	CURRENT (mA)	POWER (Watt)

1/1	34.0	117	6.200
1/2	34.0	94	5.000

1/3	34.0	535	28.500
1/4	0.0	0	0.000
1/5	0.0	0	0.000
1/6	34.0	525	27.900
1/7	34.0	152	8.100
1/8	34.0	49	2.600

**Note**

The PoE port measurement for all ports is displayed. The preceding output is a sample of the full output.

Power over Ethernet configuration using EDM

This section provides details to configure PoE settings using EDM.

Configure PoE Globally

About This Task

Configure PoE usage threshold and device type settings, and enable Fast PoE and Perpetual PoE globally on a switch.

**Important**

- After you enable Fast PoE or Perpetual PoE or both, you must save the running configuration file.
- If Fast PoE or Perpetual PoE are enabled and you change any other global or port-specific PoE settings, you must resave the running configuration file.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **PoE** tab.
4. For **Usage Threshold%**, type the value in percentage for comparing the measured power.
5. For **DetectType**, select the PoE detection type depending on the capabilities of the switch.
6. Select **FastPoeEnable**.
7. Select **PerpetualPoeEnable**.
8. Select **Apply**.

PoE Field Descriptions

Use the data in the following table to use the **PoE** tab.

Name	Description
Power(watts)	Specifies the nominal power of the Power Sourcing Entity expressed in Watts.
OperStatus	Specifies the operational status of the main Power Sourcing Entity.
ConsumptionPower(watts)	Specifies the measured usage power expressed in Watts.
UsageThreshold%	Configures the usage threshold in percent for comparing the measured power and initiating an alarm if the threshold is exceeded.
DetectType	<p>Configures the mechanism used to detect powered ethernet devices attached to a powered ethernet port. The options are:</p> <ul style="list-style-type: none"> • 802.3af • 802.3af and legacy support • 802.3at • 802.3at and legacy support • 802.3bt Type 3 • 802.3bt Type 4 <p>Note: The default detection type is automatically configured for the highest PoE standard supported on the switch.</p>
PowerPresent	<p>Specifies the current power source present on the switch. Available power sources are AC and DC.</p> <p>A value of acOnly indicates that the only power supply is AC.</p> <p>A value of dcOnly indicates that the only power supply is DC.</p> <p>A value of acDc indicates that the two power supplies, AC and DC are supplying power.</p>
FastPoeEnable	Enables Fast PoE on the switch. The default is disabled.
PerpetualPoeEnable	Enables Perpetual PoE on the switch. The default is disabled.

Configure PoE on Ports

About This Task

Enable or disable PoE on a port, and configure PoE priority and power limit settings.



Important

If Fast PoE or Perpetual PoE are enabled and you change any other global or port-specific PoE settings, you must resave the running configuration file.

For information about which ports support PoE, see the hardware documentation.

Procedure

1. In the Device Physical View, select one or more ports that support PoE.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **PoE** tab.
5. Configure the fields as required.
6. Select **Apply**.

PoE Field Descriptions

Use the data in the following table to use the **PoE** tab.

Name	Description
AdminEnable	Enabled or disables PoE on this port.
FastPoeEnable	Enables or disables Fast PoE on this port. The default is disabled.
PerpetualPoeEnable	Enables or disables Perpetual PoE on this port. The default is disabled.
DetectionStatus	Specifies the operational status of the power device detecting mode on this port: <ul style="list-style-type: none">• Disabled—detecting function disabled• Searching—detecting function is enabled and the system is searching for a valid powered device on this port• DeliveringPower—detection found a valid powered device and the port is delivering power• Fault (OtherFault)—a power-specific fault has been detected on the port• Test—detecting device is in test mode
PowerClassifications	Specifies the power classification of the device connected to this port. Power classification tags different terminals on the Power over LAN network according to their power consumption. Devices such as IP telephones, WLAN access points, and others can be classified according to their power requirements.

Name	Description
PowerPriority	Configures the power priority for this port: <ul style="list-style-type: none">• critical• high• low
PowerLimit(Watts)	Configures the maximum power that the switch can supply to a port.
DetectType Note: Exception: only supported on 5420 Series and 5520 Series.	Configures the PoE powered device (PD) detection type on a specific port or ports. The default detection type is automatically configured for the highest PoE standard supported on the switch.

View Power Consumption for a PoE Port

About This Task

Perform the following procedure to display the PoE power measurement. You can view the voltage, amperage, and wattage for every PoE port. PoE ports without a powered device (PD) in use are measured as zeros.

For information about which ports support PoE, see the hardware documentation.

Procedure

1. In the Device Physical View, select one or more ports that support PoE.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select **PoE Power Measurement**.

PoE Power Measurement *Field Descriptions*

Use the data in the following table to use the **PoE Power Measurement** tab.

Name	Description
Voltage(volts)	Displays the power measured in volts.
Current(amps)	Displays the power measured in amps.
Power(watts)	Displays the power measured in watts.



Power Savings

[Energy Saver](#) on page 2966

[Energy Efficient Ethernet](#) on page 2968

[Power Savings Configuration Using CLI](#) on page 2969

[Power Savings Configuration Using EDM](#) on page 2978

Power savings allow you to reduce network infrastructure power consumption during periods of low data activity without impacting network connectivity.

Depending on the power saving option you choose, you can implement power savings on a switch-wide, or on a per-port basis. The following power saving options are supported:

- Energy Saver (switch-wide or per-port)
- Energy Efficient Ethernet (EEE) (per-port only)

You must choose either Energy Saver, or Energy Efficient Ethernet (EEE)—you cannot use both options together.

The following sections describe the Energy Saver and Energy Efficient Ethernet (EEE) features, and how to configure them.

Energy Saver

Table 225: Energy Saver product support

Feature	Product	Release introduced
Energy Saver	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4 Not supported on 5420F-24S-4XE
	5520 Series	VOSS 8.2.5 Not supported on 5520-24X or 5520-48SE Supported on fixed ports only.
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10 Supported on 7520-48XT-6C
	7720 Series	Not Supported
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1 Supported on VSP4900-48P and ports 1/1 to 1/12 on VSP4900-12MXU-12XE
	VSP 7400 Series	Not Supported

To reduce direct power consumption by up to 40%, Energy Saver uses intelligent-switching capacity reduction in off-peak mode by controlling port link speeds and optionally powering off low priority Power over Ethernet (PoE) devices during off-peak periods. You can schedule Energy Saver to activate during multiple specific time periods. These time periods can be as short as one minute, or can last a week, a weekend, or individual days.



Note

- Energy Saver is supported only on copper ports that have auto-negotiation enabled on them.
- Energy Saver does not operate on the 5720 Series when you configure local port Auto-Negotiation advertisements to *2500-full* or *5000-full*.
- If auto-negotiation is disabled on a port and a custom port speed is configured, Energy Saver will not change the speed of that port.

**Important**

- Configuring the port link speed to a low value impacts the overall network performance. The best practice is to use the Energy Saver feature during the hours when the network is not overburdened.
- If a switch is reset while Energy Saver is activated, the PoE power-saving calculation might not accurately reflect the power saving, and in some cases might display zero savings. This problem occurs because the switch did not have sufficient time to record PoE usage between the reset of the switch and the reactivation of Energy Saver. When Energy Saver is next activated, the PoE power saving calculation is correctly updated.
- When Energy Saver is active and you replace a unit, that unit will not be in Energy Saver mode. You must configure Energy Saver directly after replacing a unit.

Interaction with PoE

Energy Saver can use Power over Ethernet (PoE) port-power priority levels to shut down low-priority PoE ports and provide power savings. The power consumption savings of each switch is determined by the number of ports with Energy Saver enabled, and by the power consumption of PoE ports that are powered off. If Energy Saver is disabled on a port, the port is not powered off, irrespective of the PoE configuration. Energy Saver turns off the power to a port only when PoE is enabled globally, the port Energy Saver is enabled, and the PoE priority for the port is configured to Low.

Configuration Fundamentals

To fully configure and use Energy Saver, you must first enable Energy Saver on ports, create a schedule, and then enable Energy Saver globally.

Alternatively, you can configure Energy Saver using the Efficiency Mode quick configuration method, which enables Energy Saver on all ports, creates a default schedule, and enables Energy Saver globally.

You can manually deactivate and reactivate Energy Saver at any time, without affecting the port configurations.

**Note**

- Energy Saver is supported only on copper ports that have auto-negotiation enabled.
- Network Time Protocol (NTP) must be enabled and configured to use Energy Saver.

Energy Efficient Ethernet

Table 226: Energy Efficient Ethernet product support

Feature	Product	Release introduced
Energy Efficient Ethernet (EEE)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4 Supported on fixed copper ports only Not supported on 5420F-24S-4XE
	5520 Series	VOSS 8.2.5 Supported on fixed copper ports only Not supported on 5520-24X or 5520-48SE
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10 Supported on 7520-48XT-6C
	7720 Series	Not Supported
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1 Supported on all fixed copper ports on VSP4900-48P and ports 1/1-1/12 on VSP4900-12MXU-12XE
	VSP 7400 Series	Not Supported

Energy Efficient Ethernet (EEE) supports the IEEE 802.3az standard for power savings in Ethernet networks for a select group of physical layer devices. A physical device that can support low power idle (LPI) mode is considered EEE-capable. Legacy devices that do not support EEE can be made EEE-compliant with an EEE-compliant PHY and an SDK version that allows the MAC device to interact with the PHY EEE functionality.

In a typical configuration, the EEE protocol communicates with the switch and the physical device to determine when to enter LPI mode during a period of inactivity, and to exit LPI mode when data transmission resumes.



Note

EEE is supported only on copper ports that have auto-negotiation enabled on them.

Power Savings Configuration Using CLI

Configure Energy Saver or Energy Efficient Ethernet using the command line interface (CLI).

Enable Energy Saver on Ports

Before You Begin

- If you have previously enabled Energy Saver globally, you must disable it globally before enabling Energy Saver on individual ports.
- If you have previously enabled Efficiency Mode, you must disable it before enabling Energy Saver on individual ports.

About This Task

Perform this procedure to enable Energy Saver on a specific port or range of ports.



Note

Energy Saver does not operate on the 5720 Series when you configure local port Auto-Negotiation advertisements to *2500-full* or *5000-full*.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format *slot/port/sub-port*.

2. Enable Energy Saver on the specified port:

Identifies the slot and port in one of the following formats:

- a single slot and port (slot/port)
- a range of slots and ports (slot/port-slot/port)
- a series of slots and ports (slot/port,slot/port,slot/port)
- all ports on the same slot (slot/all)
- all ports on the switch (all)

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format *slot/port/sub-port*.

```
energy-saver port {slot/port[/sub-port] [-slot/port[/sub-port]][,...]  
[slot/all][all]} enable
```

Example

Enable energy savings on slot 1 port 2:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 1/2
Switch:1(config-if)#energy-saver port 1/2 enable
```

What to Do Next

Configure an Energy Saver schedule.

Variable Definitions

The following table defines parameters for the **energy-saver** command.

Variable	Value
<i>enable</i>	Enables energy savings on ports. The default is disabled.
<i>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Create an Energy Saver Schedule

Before You Begin

- If you have previously enabled Energy Saver globally, you must disable it globally before creating a schedule.
- If you have previously enabled Efficiency Mode, you must disable it before creating a schedule. You cannot change the default Efficiency Mode schedule entries when Efficiency Mode is enabled.
- You must enable Energy Saver on every port affected by the schedule.

About This Task

Perform this procedure to configure scheduled time intervals during which the switch will operate in low power state. This time interval can be for a week, weekend, or individual days.



Note

You can configure a maximum of 84 entries in the Energy Saver schedule.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the Energy Saver schedule:

```
energy-saver schedule {friday | monday | saturday | sunday | thursday  
| tuesday | wednesday | weekday | weekend} <hhmm> {activate |  
deactivate}
```

Example

Configure an Energy Saver schedule:

```
Switch:1>enable  
Switch:1# configure terminal  
Switch:1(config)#energy-saver schedule weekend 0735 activate  
Switch:1(config)#energy-saver schedule monday 0600 deactivate
```

What to Do Next

Enable Energy Saver globally.

Variable Definitions

The following table defines parameters for the **energy-saver schedule** command.

Variable	Value
<i>{activate deactivate}</i>	Activates or deactivates the scheduled event.
<i><hhmm></i>	Specifies the hour and minutes to enable Energy Saver feature on the switch.
<i>{friday monday saturday sunday thursday tuesday wednesday weekday weekend}</i>	Specifies the day(s) to enable Energy Saver feature on the switch.

Enable Energy Saver Globally

Before You Begin

- You must enable Energy Saver on individual ports.
- You must create an Energy Saver schedule.

About This Task

Perform this procedure to enable Energy Saver globally on the switch. You can optionally configure PoE power savings, to power off low priority PoE devices during off-peak times.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. (Optional) Configure PoE power savings:

```
energy-saver poe-power-saving
```



Note

You must configure PoE power savings before you enable Energy Saver globally.

3. Enable Energy Saver:

```
energy-saver enable
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#energy-saver poe-power-saving
Switch:1(config)#energy-saver enable
```

Variable Definitions

The following table defines parameters for the **energy-saver** command.

Variable	Value
<i>enable</i>	Enables Energy Saver feature. The default is disabled.
<i>poe-power-saving</i>	Enables PoE power saving. The default is disabled.

Enable and Configure Energy Saver using Quick Configuration

About This Task

Perform this procedure to enable and configure Energy Saver globally using the quick configuration Efficiency Mode. Efficiency Mode automatically configures the following:

- enables Energy Saver on all ports.
- creates a default schedule with a weekday schedule of Energy Saver activated from 6:00 p.m. to 7:30 a.m., and during weekends. You cannot change this default schedule while Efficiency Mode is enabled.
- enables Energy Saver globally.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable efficiency mode:

```
energy-saver efficiency-mode
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#energy-saver efficiency-mode
```

Variable Definitions

The following table defines parameters for the **energy-saver** command.

Variable	Value
<i>efficiency-mode</i>	Enables efficiency mode. The default is disabled.

Activate or Deactivate Energy Saver Manually

Before You Begin

Before you can change any previously saved Energy Saver settings, you must disable Energy Saver globally.

About This Task

Perform this procedure to activate or deactivate Energy Saver on the switch at any time. Energy Saver is deactivated by default.

Activating Energy Saver reduces the port speed to the minimum value supported by the switch and enables PoE power saving, even if PoE is globally disabled. Deactivating Energy Saver restores the previous configuration.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Activate or deactivate Energy Saver:

```
energy-saver {activate | deactivate}
```

Example

```
Switch:1> enable
```

Activate Energy Saver:

```
Switch:1# energy-saver activate
```

Deactivate Energy Saver:

```
Switch:1# energy-saver deactivate
```

Variable Definitions

The following table defines parameters for the **energy-saver** command.

Variable	Value
<i>activate</i>	Activates Energy Saver manually.
<i>deactivate</i>	Deactivates Energy Saver manually.

Energy Saver Show Commands

Use the procedures in this section to display specific information about Energy Saver configuration on the switch.

Display Energy Saver Global Information

About This Task

Perform this procedure to display information about Energy Saver global configuration.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display global configuration:
`show energy-saver global`

Example

```
Switch:1#show energy-saver global
Energy Saver: Disabled
Energy Saver PoE Power Saving Mode: Disabled
Energy Saver Efficiency-Mode Mode: Disabled
Day/Time: Wednesday 02:31:12
Current Energy Saver state: Energy Saver is Inactive
```

Display Energy Saver Interface Information

About This Task

Perform this procedure to display information about Energy Saver configuration on the ports.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display information about all ports or specify a particular port:
`show energy-saver interface [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]`

Example

```
Switch:1>enable
Switch:1#show energy-saver interface
Port ES State PoE Savings PoE Priority
```

```

-----
1/1      Enabled  Enabled  Low
1/2      Enabled  Enabled  Low
1/3      Enabled  Enabled  Low
1/4      Enabled  Enabled  Low
1/5      Enabled  Enabled  Low
1/6      Enabled  Enabled  Low
1/7      Enabled  Enabled  Low
1/8      Enabled  Enabled  Low
1/9      Enabled  Enabled  Low
1/10     Enabled  Enabled  Low
1/11     Enabled  Enabled  Low
1/12     Enabled  Enabled  Low
1/13     Disabled N/A      N/A
1/14     Disabled N/A      N/A
1/15     Disabled N/A      N/A

```

Display Energy Saver Power Savings Information

About This Task

Perform this procedure to display information about Energy Saver power savings on the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about energy savings:
show energy-saver savings

Example

```

Switch:1#show energy-saver savings
-----
Unit Model          Switch Capacity      Saving PoE Saving
-----
8404C                0.0 watts            N/A
=====

```

Display Energy Saver Schedule Information

About This Task

Perform this procedure to display information about Energy Saver schedules configured on the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about Energy Saver schedules:
show energy-saver schedule

Example

```

Switch:1#show energy-saver schedule
-----
Day      Time  Action
-----

```

Monday	18:00	Activate
Monday	07:00	Deactivate

Enable Energy Efficient Ethernet (EEE)

About This Task

Perform this procedure to enable Energy Efficient Ethernet (EEE) on a port. The default is disabled.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:


```
enable
```



```
configure terminal
```



```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
```

```
[,...][slot/all][all]}
```
-  **Note**
If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
2. Enable Energy Efficient Ethernet:

```
energy-saver eee enable
```

Display Energy Efficient Ethernet (EEE) Statistics

About This Task

Perform this procedure to display information about Energy Efficient Ethernet (EEE) statistics for all ports on a switch, or for a specific port.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about all ports or specify a particular port:

```
show energy-saver eee statistics [{slot/port[/sub-port]}[-slot/port[/sub-port]]
```

```
[,...][slot/all][all]}
```

Example

Switch:1>show energy-saver eee statistics					
=====					
EEE Port Status					
=====					
PortId	EEE Status	Tx LPI Events	Tx Idle Duration (micro seconds)	Rx LPI Events	Rx Idle Duration (micro seconds)

1/1	enabled	847	963657920	115	965100020

Variable Definitions

Use the data in following table to use the `show energy-saver eee statistics` command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Clear Energy Efficient Ethernet (EEE) Statistics

About This Task

Perform this procedure to clear information about Energy Efficient Ethernet (EEE) statistics for all ports on a switch, or for a specific port.

Procedure

1. Enter Privileged EXEC mode:

enable
2. Clear Energy Efficient Ethernet statistics on all ports or specify a particular port:

clear energy-saver eee stats [port {slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}

Variable Definitions

Use the data in following table to use the `clear energy-saver eee stats` command.

Variable	Value
<code>port</code>	Specifies one or more ports.
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Power Savings Configuration Using EDM

Use the following procedures to configure either Energy Saver or Energy Efficient Ethernet EEE using Enterprise Device Manager (EDM).

Enable Energy Saver Globally

About This Task

Perform this procedure to enable Energy Saver globally.

Procedure

1. In the navigation pane, expand **Configuration > Power Management**.
2. Select **Energy Saver**.
3. Select the **Energy Saver Globals** tab.
4. Configure the fields as required.
5. Select **Apply**.

Energy Saver Globals Field Descriptions

Use the data in the following table to use the **Energy Saver Globals** tab.

Name	Description
EnergySaverEnabled	Enables Energy Saver globally on the switch. The default is disabled.
PoePowerSavingEnabled	Enables Energy Saver PoE power saving. The default is disabled.

Name	Description
EfficiencyModeEnabled	Enables Energy Saver efficiency mode. The default is disabled. Efficiency mode enables Energy Saver globally and on all ports, it also enables PoE power saving. It also creates a weekday schedule that starts at 6:00 p.m. and ends at 7:30 a.m., and during the weekend Energy Saver is always activated.
EnergySaverActive	Activates Energy Saver on the switch. Energy Saver is deactivated by default.

Configure Energy Saver Schedule

Before You Begin

- You must disable Energy Saver globally.
- You must enable Energy Saver on every port affected by the schedule.
- You must deactivate Energy Saver efficiency-mode.

About This Task

Perform this procedure to configure a scheduled time interval during which the switch will operate in low power state. This time interval can be for a week, weekend, or individual days.



Note

- You can configure maximum 84 entries in the Energy Saver schedule.
- If efficiency mode is enabled, you cannot configure any other entries in the Energy Saver schedule.

Procedure

1. In the navigation pane, expand **Configuration > Power Management**.
2. Select **Energy Saver**.
3. Select the **Energy Saver Schedules** tab.
4. Select **Insert**.
5. Configure the fields as required.
6. Select **Insert**.
7. Select **Apply**.

Energy Saver Schedules *Field Descriptions*

Use the data in the following table to use the **Energy Saver Schedules** tab.

Name	Description
ScheduleDay	Specifies the day on which Energy Saver is activated or deactivated. The options are: <ul style="list-style-type: none"> • monday • tuesday • wednesday • thursday • friday • saturday • sunday • weekdays • weekend
ScheduleHour	Specifies the hour at which Energy Saver is activated or deactivated. The range is 0 to 23. Note: 0 is equivalent to 12 a.m., and 12 is equivalent to 12 p.m.
ScheduleMinute	Specifies the minute at which Energy Saver is activated or deactivated. The range is 0 to 59.
ScheduleAction	Specifies if Energy Saver is activated or deactivated. The options are: <ul style="list-style-type: none"> • activate • deactivate

Enable Energy Saver or EEE on Ports

You can enable Energy Saver or EEE using the Energy Saver tab accessed by the Edit navigation path, or the Power Management navigation path. Use one of the following procedures to enable either Energy Saver or EEE on ports.

*Enable Energy Saver or EEE on Ports***About This Task**

Perform this procedure to enable Energy Saver or EEE on one or more ports.

Procedure

1. On the Device Physical View tab, select one or more ports.
2. In the navigation pane, expand **Edit > Port**.
3. Select **General**.
4. Select the **Energy Saver** tab.

5. Enable either Energy Saver or EEE:
 - To enable Energy Saver, select **EnergySaverEnabled**.
 - To enable EEE, select **EnergySaverEEEEEnable**.
6. Select **Apply**.

Energy Saver Field Descriptions

Use the data in the following table to use the **Energy Saver** tab.

Name	Description
Port	Specifies the port number.
EnergySaverEnabled	Configures whether Energy Saver is enabled on the specific port.
EnergySavedPoeStatus	Specifies the Energy Saver PoE status for the specific port.
EnergySaverEEEEEnable	Configures whether EEE is enabled on the specific port.

Enable Energy Saver or EEE on Ports

About This Task

Perform this procedure to enable Energy Saver or EEE on one or more ports.

Procedure

1. In the navigation pane, expand **Configuration > Power Management**.
2. Select **Energy Saver**.
3. Select the **Ports** tab.
4. Enable either Energy Saver or EEE:
 - To enable Energy Saver, in the **EnergySaverEnabled** column, double-click the field associated with the specific ports, and then select **true**.
 - To enable EEE, in the **EnergySaverEEEEEnable** column, double-click the field associated with the specific ports, and then select **true**.
5. Select **Apply**.

Ports Field Descriptions

Use the data in the following table to use the **Ports** tab.

Name	Description
Port	Specifies the port number.
EnergySaverEnabled	Configures whether Energy Saver is enabled on the port.
EnergySaverPoEStatus	Displays Energy Saver PoE status for the port.
EnergySaverEEEEEnable	Configures whether EEE is enabled on the port.

View Energy Savings

About This Task

Perform this procedure to view the amount of switch capacity and PoE power being saved on the units.

Procedure

1. In the navigation pane, expand **Configuration > Power Management**.
2. Select **Energy Saver**.
3. Select the **Energy Savings** tab.

Energy Savings Field Descriptions

Use the data in the following table to use the **Energy Savings** tab.

Name	Description
UnitIndex	Specifies the unit number.
UnitSavings(1/10 watts)	Specifies the amount of switch capacity power being saved on the specific unit.
PoeSavings(1/10 watts)	Specifies the amount of PoE power being saved on the specific unit.

Display Energy Efficient Ethernet Statistics

Perform this procedure to display information about Energy Efficient Ethernet (EEE) statistics for all ports on a switch, or for a specific port.

Procedure

1. In the navigation pane, expand **Configuration > Power Management**.
2. Select **Energy Saver**.
3. Select the **EEE Statistics** tab.
4. (Optional) Select **Clear Stats** to clear information about EEE statistics.

EEE Statistics Field Descriptions

Use the data in the following table to use the **EEE Statistics** tab.

Name	Description
Port	Shows the port number.
State	Shows the state of Energy Efficient Ethernet (EEE) on the port.
TxEvents	Shows the EEE Tx event count for the port.
TxDuration	Shows the EEE Tx durations for the port.
RxEvents	Shows the EEE Rx event count for the port.
RxDuration	Shows the EEE Rx durations for the port.



Quality of Service

[QoS Fundamentals](#) on page 2983

[Basic DiffServ configuration using CLI](#) on page 3012

[Basic DiffServ configuration using EDM](#) on page 3017

[QoS Configuration Using CLI](#) on page 3019

[QoS Configuration Using EDM](#) on page 3034

This section provides conceptual and procedural information about Quality of Service (QoS).

QoS Fundamentals

Use the information in this section to help you understand Quality of Service (QoS).

This section describes a range of features that you can use on the switch to manage traffic flowing through your network. You can configure your network to prioritize specific types of traffic to ensure that the traffic receives the appropriate QoS level. For those cases where traffic levels are so high that congestion occurs despite management, the switch provides additional congestions handling features that are described in this section.

QoS refers to the ability to control network flows either by prioritizing traffic or by guaranteeing performance levels. QoS does not refer to a specifically achieved service quality. To provide QoS, you can use some combination of the switch's traffic management tools to help deliver provisioned network QoS. It is up to the network administrator to accurately analyze a given situation and select the proper tool(s) for the task.

Introduction to QoS

The switch comes with a set of traffic management tools that you can use to provide QoS for Layer 2 (bridged) or Layer 3 (routed) traffic flows. Many of these flows are multiplexed across a set of network switches and compete for network resources at convergence points. Without traffic management, the congested data flows compete for resources and the result is unpredictable. The resulting QoS can only be described as best-effort. The opposite is also true, without congestion there are sufficient network resources for all traffic to pass without competition. Without congestion, traffic management is not required. In this sense, you cannot separate discussions about QoS and traffic management from those on congestion. The switch provides a set of tools

that you can use to provide network services that are far superior to best-effort thereby enabling the delivery of provisioned QoS.

To deliver QoS, the switch uses two types of traffic management tools:

- Congestion management
- Congestion handling

Congestion management acts to prevent congestion by prioritizing traffic flows through priority queuing and priority-aware servicing methods. Other functions, such as policing, are also considered congestion management. Policing indirectly prioritizes some traffic by limiting the rates of other traffic.

Congestion handling alleviates existing congestion by dropping lower priority traffic before higher priority traffic. The switch handles the congestion by queue-specific tail dropping. The basic QoS architecture of the switch identifies three primary functional areas:

- Ingress QoS identification and classification
- the switch's internals and queuing architecture
- QoS marking/remarking for downstream use



Important

Remarking packets with an ACL filter does not change the internal QoS level of the packets. You must add the **permit internal-qos [value]** statement to the ACL filter.

The QoS architecture is coherent end-to-end across a network. The QoS at any particular network element can be marked in the relevant Layer 2 or Layer 3 protocol fields and provided to the next hop. The receiving next hop can then use this information to classify its own ingress traffic, apply its specific internal traffic management features, and remark the results for subsequent hops.

Depending on the product, the QoS implementation on the switch can support the following options:

1. Ingress priority mappings including: DSCP to internal QoS, 802.1p-bits to internal QoS, and port-level QoS configuration.
2. Egress priority mappings including: internal QoS to DSCP and internal QoS to 802.1p-bits.
3. Port-based rate limiting or ingress policers



Note

7520 Series and 7720 Series do not support port-based rate limiting or ingress policers.

4. Port-based broadcast and multicast rate limiting
5. Port-based egress shaping
6. Egress queue rate limiting

For information about feature support, see *Fabric Engine and VOSS Feature Matrix*.

Traffic Management

Prioritized traffic handling requires QoS classification first. The switch typically classifies traffic by using the endpoint switch configuration in conjunction with the protocol elements of the incoming frame such as *priority*. You can add additional classification by using access control lists (ACLs) and other filtering functionality.

The switch's internal traffic management functions use the results of classification to determine the prioritization of traffic. Examples of internal functions that prioritize traffic would be both strict and weighted round robin (WRR) queue scheduling. These mechanisms prioritize data by favorable scheduling or weighting.

The disposition of a particular data frame is not necessarily fully determined as a result of classification. You can apply additional traffic management functions such as Ingress Port Rate Limiting or Policing depending on your switch features. Ingress Port Rate Limiting is a congestion management mechanism to limit the traffic rate accepted by the specified ingress port. Policing is a congestion management method that limits incoming traffic at the ingress that could cause congestion at the egress. While queuing strategies affect prioritization by favoring some traffic; policing is essentially the opposite. Policing works not by favoring some traffic, but by penalizing the bursty traffic. Queue scheduling and policing are only two of the available tools for congestion management. Additional details are presented in subsequent sections of this document.

In addition to the traffic management tools which aid in the prevention of congestion, tools are also provided to handle congestion as it occurs. Congestion handling tools monitor congestion levels at convergence points in the switch and selectively discard frames if congestion begins to increase. Per queue tail dropping is the primary congestion handling function of the switch. You can also use ACLs and filtering as congestion handling tools.

Differentiated Services (DiffServ)

Table 227: Differentiated Services product support

Feature	Product	Release introduced
Differentiated Services (DiffServ) including Per-Hop Behavior	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Differentiated Services (DiffServ) is a traffic management tool that classifies network traffic into eight traffic classes, and then gives each class differentiated treatment. DiffServ networks map the traffic's class into a set of packet forwarding behavior, referred to as a *Per-Hop Behavior (PHB)*. A PHB could specify which egress queue to use. For example, a switch may classify a packet by determining its protocol to be IPv4, subsequently extract the DSCP value, and apply a PHB by directing the packet to a specific queue. DiffServ does not prescribe a set of traffic classes and does not predetermine which types of traffic should be handled by a given class. DiffServ simply provides a generic means of classifying packets so they may be treated differently.

DiffServ applies to IP packets only.

DiffServ Access and DiffServ Core

A fundamental characteristic of DiffServ networks is the distinction made between switches at the network edges and those residing in the network interior. The switch refers to this distinction as DiffServ Access (edge) and DiffServ Core (interior), respectively.

It is important to note that the switch operates simultaneously as both a DiffServ Access switch and a DiffServ Core switch. The architectural premise is that the edge or access nodes perform the bulk of the work (classification, policing, etc.) and mark the packet for downstream processing. In theory this would permit the interior or core switches to bypass much of the edge processing as they would “trust” the classification and marking performed by the access switch. The notion of trust is key to the access/core switch distinction.

- If you configure a port as an access port, the system does not trust packet markings.
- If you configure a port as a core port, the system trusts packet markings.

On the access side, malicious users can send packets into a network with intent to cause serious harm (e.g., denial of service attacks). However, on core switches, the only traffic sources are one's own upstream switches. As such, a core switch has the opportunity to trust the classification, markings (and implied PHB) determined by the previous hop.

The following figure shows DiffServ network operations. The devices are on the network edge where they perform classification, marking, policing, and shaping functions.




Figure 200: DiffServ network core and edge devices

Use a DiffServ Access port at the edge of a DiffServ network. The access port classifies traffic according to port QoS. Outgoing packet DSCP and 802.1p values are derived

from port QoS and QoS maps. The system strips Dot1Q headers at ingress, and adds them back at egress only if you configure the egress port as a tagged or trunk port.

A DiffServ Core port does not change packet classification or markings; the port trusts the incoming traffic markings. A core port preserves the DSCP marking of all incoming packets, and uses these markings to assign the packet to an internal QoS level. For tagged packets, the port honors the 802.1p bits within a Dot1Q header, and uses these bits to classify ingress traffic. You can control the honoring (or not) of 802.1p bits by configuring the 802.1p override.

Per-Hop Behavior (PHB)

Traffic entering the DiffServ network enters a queue according to their marking, which determines the PHB of the packets. For example, if the system marks a video stream to receive the highest priority, it enters a high-priority queue. As these packets traverse the DiffServ network, the system forwards the video stream before other packets.

As a standard, DiffServ is described in the context of Layer 3. Classification is accomplished by mapping a packet priority field from the packet and then applying a per-hop behavior. DiffServ standards define the IPv4 header's Differentiated Services Code Point (DSCP) field to determine classification and subsequent per-hop behavior.

The RFC2598 standard provides only the following four fundamental per-hop behaviors:

- Default (DF) — This PHB provides best-effort forwarding behavior.
- Expedited Forwarding (EF) — This PHB provides performance-critical forwarding.
- Assured Forwarding (AF) — This PHB classifies traffic based upon **class** (priority) only.



Important

The switch never classifies nor takes action based upon drop precedence. In response to congestion, the only drops available for a given traffic class are tail drops.

- Class Selector (CS) — This PHB provides a simple mapping of a DSCP to one of eight traffic classes. While the switch provides all four PHBs, the CS PHB is most analogous to the switch's internal processing – classification occurs to derive priority, which subsequently determine the per-hop behavior (e.g., queuing).

DiffServ and Filters

QoS (DiffServ) and filters operate independently; you do not have to use filters to provide QoS. However, filters can override QoS operations. For more information, see [Traffic Filtering Fundamentals](#) on page 3681.

DiffServ and VXLAN Gateway

DSCP bits in the outer IP header of VXLAN-encapsulated packets are always derived from the internal QoS, irrespective of the ingress port DiffServ configuration. Customer packet IP DSCP bits are not modified as part of VXLAN encapsulation. For more information about VXLAN, see [VXLAN Gateway](#) on page 4170.

Traffic Traversing the Switch

The switch's traffic management capabilities are best understood by examining the functionality that is invoked as packets flow from ingress ports, through the switch, to egress ports. The following list includes the set of features and processing that the switch performs as flows traverse the switch:

- Classification and ingress mapping
- Filtering
- Rate Limiting or Policing
- Queueing
- Remarking
- Shaping

The switch classifies packets to determine their priority. While DiffServ is traditionally defined as Layer 3 functionality, the switch extends the logical concept to Layer 2. The switch can, based upon user configuration, determine a packet's priority from either its Layer 2 (p-bits) or Layer 3 (DSCP) information.

- If the packet arrives on an untrusted port (DiffServ Access), then the packet's priority comes from user-configured parameters such as port priority.
- If the packet arrives on a trusted port, priority comes from information contained in the packet's header (p-bits or DSCP).

After the switch determines the packet's configured or marked priority, it maps that value to be used internally. The QoS level used by the switch is referred to as the **Internal QoS Level (IQL)**. The IQL is the internal numerical value that the switch uses to determine the packets per-hop behaviors such as queue selection and bandwidth guarantee.

The following list identifies the order of DiffServ operations for a packet:

- Packet classification: IEEE 802.1p and DSCP markings classify (map) the packet to its appropriate PHB and QoS level.
- Remarking: The switch can remark packets according to QoS actions you configure on the switch (internal QoS mappings).
- Shaping: The switch provides port-based shaping. Port-based shaping shapes all outgoing traffic to a specific rate.

Classification and Mapping

Traffic classification includes functions that examine a packet to determine further actions according to defined rules. Classification involves identifying flows so that the router can modify the packet contents or Per-Hop Behavior (PHB), apply conditioning treatments to the packet, and determine how to forward the packet to the egress interface. Packet classification depends on the service type of the packet and the point in the traffic management process where the classification occurs.

The device classifies traffic as it enters the DiffServ network, and assigns appropriate PHB based on the classification. To differentiate between classes of service, the device marks the DiffServ (DS) parameter in the IP packet header, as defined in RFC2474

and RFC2475. The DSCP marking defines the forwarding treatment of the packet at each network hop. This marking (or classification) occurs at the edge of the DiffServ domain, and is based on the policy (or filter) associated with the particular microflow or aggregate flow.

You can configure the mapping of DSCP-to-forwarding behaviors and DSCP re-markings. Re-marking the DSCP resets the treatment of packets based on new network specifications or desired levels of service.

Layer 3 marking uses the DSCP parameter. Layer 2 (Ethernet) marking involves the 802.1p-bits parameter.

For Layer 2 packets, priority bits (or 802.1p bits) define the traffic priority of the Ethernet packet. You can configure an interface to map DSCP or 802.1p bits to internal QoS levels on ingress. You can configure an interface to map internal QoS levels to DSCP or 802.1p bits at egress. 802.1p bit mapping provides the Ethernet VLAN QoS requirements.

Within the network, a packet PHB associated with the DSCP determines how a device forwards the packet to the next hop—if at all. Consequently, nodes can allocate buffer and bandwidth resources to each competing traffic stream. The initial DSCP value is based on network policies for the type of service required. The objective of DSCP-to-Service Class mapping is to translate the QoS characteristics defined by the packet DSCP marker to a Service Class. The DSCP-to-Service Class mapping occurs at ingress. For each received packet, the mapping function assigns a Service Class.

The switch maintains four mapping tables. These tables translate the ingress 802.1p-bits or DSCP markings to an internal QoS level, and then retranslate the internal QoS level to an egress DSCP or 802.1p-bits marking as follows:

- ingress 802.1p-bits to QoS level
- ingress DSCP to QoS level
- QoS level to egress 802.1p-bits
- QoS level to egress DSCP

Service Classes

Service classes define a standard architecture to provide end-to-end QoS on a broad range of Ethernet switching and voice products. They function as default QoS policies built into the product. They incorporate the various QoS technologies to provide a complete end-to-end QoS behavioral treatment. The switch includes a built-in QoS implementation for service classes.

The switch includes eight preconfigured queues (corresponding to the eight service classes) on each port of an interface module.

A service class domain classifies traffic as one of the following:

- Network control traffic (Critical/Network)
- Subscriber traffic (Premium, Metal, or Standard)

Queue 7 — Critical/Network Service Class (PHB of CS6/CS7)

The switch uses the Critical/Network Service Class for traffic within a single administrative network domain. If such traffic does not get through, the network cannot function.

Queue 6 — Premium Service Class (PHB of CS5/EF)

The switch uses the Premium Service Class for IP telephony services, and provides the low latency and low jitter required to support such services. IP telephony services include Voice over IP (VoIP), voice signaling, Fax over IP (FoIP), and voice-band data services over IP (for example, analog modem). The switch can also use the Premium Service Class for Circuit Emulation Services over IP (CESoIP).

Metal Service Classes

The Platinum, Gold, Silver, and Bronze Service Classes are collectively referred to as the metal classes. The metal Service Classes provide a minimum bandwidth guarantee and are for variable bit rate or bursty types of traffic. Applications that use the metal Service Class support mechanisms that dynamically adjust their transmit rate and burst size based on congestion (packet loss) detected in the network. The following list describes the individual metal classes:

- Queue 5 — Platinum Service Class (PHB of CS4/AF41)

The switch uses the Platinum Service Class for applications that require low latency, for example, real-time services such as video conferencing and interactive gaming. Platinum Service Class traffic provides the low latency required for interhuman (interactive) communications. The Platinum Service Class provides a minimum bandwidth assurance for Assured Forwarding (AF) 41 and Class Selector (CS) 4-marked flows.

- Queue 4 — Gold Service Class (PHB of CS3/AF31)

The switch uses the Gold Service Class for applications that require near-real-time service and are not as delay-sensitive as applications that use the Platinum service. Such applications include streaming audio and video, video on demand, and surveillance video.

The Gold Service Class assumes that traffic buffers at the source and destination and, therefore, the traffic is less sensitive to delay and jitter. By default, the Gold Service Class provides a minimum bandwidth assurance for AF31, AF32, AF33 and CS3-marked flows.

- Queue 3— Silver Service Class (PHB of CS2/AF21)

The switch uses the Silver Service Class for responsive (typically client- and server-based) applications. Such applications include Systems Network Architecture (SNA) terminals (for example, a PC or Automatic Teller Machine) to mainframe (host) transactions that use Data Link Switching (SNA over IP), Telnet sessions, web-based ordering and credit card processing, financial wire transfers, and Enterprise Resource Planning applications.

Silver Service Class applications require a fast response and have asymmetrical bandwidth needs. The client sends a short message to the server and the server

responds with a much larger data flow back to the client. For example, after a user selects a hyperlink (that sends a few dozen bytes) on a webpage, a new webpage opens (that downloads kilobytes of data). The Silver Service Class provides a minimum bandwidth assurance for AF21 and CS2-marked flows.

The Silver Service Class favors short-lived, low-bandwidth TCP-based flows.

- Queue 2 — Bronze Service Class (PHB of CS1/AF11)

The switch uses the Bronze Service Class for longer-lived TCP-based flows, such as file transfers, e-mail, or noncritical Operation, Administration, and Maintenance (OAM) traffic. The Bronze Service Class provides a minimum bandwidth assurance for AF11 and CS1-marked flows. As a best practice, use the Bronze Service Class for noncritical OAM traffic with the CS1 DSCP marking.

Queue 1 and 0 — Standard (PHB of CS0/DF) and Custom Service Classes

The switch uses the Standard and Custom Service Classes for best-effort services. Delays, loss, or jitter guarantees for these service classes are not specified. However, the Standard Service Class has more forwarding resources than the custom service classes.

Internal QoS level

The internal QoS level or effective QoS level is a key element in the switch QoS architecture. The internal QoS level specifies the kind of treatment a packet receives. The switch classifies every packet that enters and assigns it an internal QoS level.

Internal QoS levels map to the queues on a port. For example, for an access port the internal QoS level is derived from the port QoS level. For Layer 3 trusted (core) ports, the system honors incoming DSCP or type of service (TOS) bits. The system assigns the internal QoS level using the ingress DSCP to QoS level map.



Important

Remarking packets with an ACL filter does not change the internal QoS level of the packets. You must add the **permit internal-qos [value]** statement to the ACL filter. For more information, see [Internal QoS Level and Remarking](#) on page 3698.

Ingress Mappings

The system uses ingress maps to translate incoming packet QoS markings to the internal QoS level. The system uses the internal QoS level to classify packets.

Ingress mappings include

- 802.1p to (internal) QoS level
- DSCP to (internal) QoS level

The following logical table shows how the system performs ingress mappings for data packets and for control packets not destined for the Control Processor (CP).

Table 228: Data packet ingress mapping

DSCP	Layer 2 trusted	Layer 3 trusted (DiffServ enabled and Access-diffserv disabled)	IP packet	Routed packet	Ingress tagged	Internal QoS
x	No	x	No	x	x	Use port QoS
x	Yes	x	No	x	No	Use port QoS
x	Yes	x	No	x	Yes	Use ingress p-bits mapping
0x1B	x	x	Yes	x	x	4
0x23	x	x	Yes	x	x	5
0x29	x	x	Yes	x	x	5
0x2F	x	x	Yes	x	x	6
x	No	No	x	x	x	Use port QoS
x	No	Yes	Yes	x	x	Use ingress DSCP mapping
x	Yes	No	Yes	x	No	Use port QoS
x	Yes	No	Yes	x	Yes	Use ingress p-bits mapping
x	Yes	Yes	Yes	No	No	Use ingress DSCP mapping

Table 228: Data packet ingress mapping (continued)

DSCP	Layer 2 trusted	Layer 3 trusted (DiffServ enabled and Access-diffserv disabled)	IP packet	Routed packet	Ingress tagged	Internal QoS
x	Yes	Yes	Yes	No	Yes	Use ingress p-bits mapping
x	Yes	Yes	Yes	Yes	Yes	Use ingress DSCP mapping

**Important**

On a tagged port that is Layer-2 trusted, Layer-3 trusted and DiffServ enabled, all multicast packets honor the ingress DSCP value.

On the 7830 Series, an NNI port is always considered Layer-2 trusted regardless of the configuration on the NNI interface.

The QoS level for control packets destined for the CPU is assigned internally to ensure timely packet processing and scaling numbers. You cannot configure the QoS level for these control packets. The system assigns the highest QoS-level to time-critical protocols.

The following table shows ingress IEEE 802.1p to QoS level mappings.

Table 229: Default ingress 802.1p to QoS mappings

Ingress IEEE 802.1p	PHB	QoS Level	Network Service Class (NSC)
0	CS0/DF	1	Standard
1	Custom	0	Custom
2	CS1/AF11	2	Bronze
3	CS2/AF21	3	Silver
4	CS3/AF31	4	Gold
5	CS4/AF41	5	Platinum
6	CS5/EF	6	Premium/EF
7	CS6/CS7	7	Network/Critical

The following table shows DSCP to internal QoS level mappings.

Table 230: Default ingress DSCP to QoS mapping

Ingress				Internal QoS level	PHB level
DSCP (decimal)	DSCP (binary)	DSCP (hexadecimal)	TOS (hexadecimal)		
00	000000	00	00	1	CS0/DF
00	000000	00	00	1	DF
01	000001	01	04	1	CS0
02	000010	02	08	1	CS0
03	000011	03	0C	1	CS0
04	000100	04	10	1	CS0
05	000101	05	14	1	CS0
06	000110	06	18	1	CS0
07	000111	07	1C	1	CS0
08	001000	08	20	2	CS1
09	001001	09	24	1	CS0
10	001010	0A	28	2	AF11
11	001011	0B	2C	1	CS0
12	001100	0C	30	2	CS1
13	001101	0D	34	1	CS0
14	001110	0E	38	2	CS1
15	001111	0F	3C	1	CS0
16	010000	10	40	3	CS2
17	010001	11	44	1	CS0
18	010010	12	48	3	AF21
19	010011	13	4C	1	CS0
20	010100	14	50	3	CS2
21	010101	15	54	1	CS0
22	010110	16	58	3	CS2
23	010111	17	5C	1	CS0
24	011000	18	60	4	CS3
25	011001	19	64	1	CS0
26	011010	1A	68	4	AF31
27	011011	1B	6C	4	CS3
28	011100	1C	70	4	CS3

Table 230: Default ingress DSCP to QoS mapping (continued)

Ingress				Internal QoS level	PHB level
DSCP (decimal)	DSCP (binary)	DSCP (hexadecimal)	TOS (hexadecimal)		
29	011101	1D	74	1	CS0
30	011110	1E	78	4	CS3
31	011111	1F	7C	1	CS0
32	100000	20	80	5	CS4
33	100001	21	84	1	CS0
34	100010	22	88	5	AF41
35	100011	23	8C	5	CS4
36	100100	24	90	5	CS4
37	100101	25	94	1	CS0
38	100110	26	98	5	CS4
39	100111	27	9C	1	CS0
40	101000	28	A0	6	CS5
41	101001	29	A4	5	CS4
42	101010	2A	A8	1	CS0
43	101011	2B	AC	1	CS0
44	101100	2C	B0	1	CS0
45	101101	2D	B4	1	CS0
46	101110	2E	B8	6	EF
47	101111	2F	BC	6	CS5
48	110000	30	C0	7	CS6
49	110001	31	C4	1	CS0
50	110010	32	C8	1	CS0
51	110011	33	CC	1	CS0
52	110100	34	D0	1	CS0
53	110101	35	D4	1	CS0
54	110110	36	D8	1	CS0
55	110111	37	DC	1	CS0
56	111000	38	E0	7	CS7
57	111001	39	E4	1	CS0
58	111010	3A	E8	1	CS0
59	111011	3B	EC	1	CS0

Table 230: Default ingress DSCP to QoS mapping (continued)

Ingress				Internal QoS level	PHB level
DSCP (decimal)	DSCP (binary)	DSCP (hexadecimal)	TOS (hexadecimal)		
60	111100	3C	F0	1	CS0
61	111101	3D	F4	1	CS0
62	111110	3E	F8	1	CS0
63	111111	3F	FC	1	CS0

Egress Mappings

Egress mappings include:

- QoS level to IEEE 802.1p mappings
- QoS level to DSCP mappings

When a packet is forwarded by the switch, the software does the following:

- Always performs 802.1p remarking before the packet egresses.
- If the ingress port has **enable-diffserv** and **access-diffserv** enabled, then the IP packet is DSCP remarked before the packet egresses.

If the ingress port is not configured this way, the packets are not DSCP remarked.

The following table shows egress QoS level to IEEE 802.1p mappings.

Table 231: Default egress QoS level to IEEE 802.1p mappings

QoS level	PHB	Default 1p remarking on egress	Network Service Class (NSC)
0	Custom	1	Custom
1	CS0/DF	0	Standard
2	CS1/AF11	2	Bronze
3	CS2/AF21	3	Silver
4	CS3/AF31	4	Gold
5	CS4/AF41	5	Platinum
6	CS5/EF	6	Premium/EF
7	CS6/CS7	7	Network/Critical

The following table shows QoS level to DSCP mappings.

Table 232: Default egress QoS level to DSCP mappings

Egress			
QoS level	DSCP (binary)	DSCP (hexadecimal)	DSCP
0	000000	00	0
1	000000	00	0
2	001010	0A	10
3	010010	12	18
4	011010	1A	26
5	100010	22	34
6	101110	2E	46
7	101110	2E	46

Port-Based Rate Limiting, Policing, and Shaping

Table 233: Port-Based Rate Limiting, Policing, and Shaping product support

Feature	Product	Release introduced
Egress port shaper	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 233: Port-Based Rate Limiting, Policing, and Shaping product support (continued)

Feature	Product	Release introduced
Ingress policer and port rate limiter	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5
QoS ingress port-based rate limiter	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Not Supported
	7720 Series	Not Supported
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	Not Supported

The switch QoS implementation supports egress port-based shaping for bandwidth management and traffic control. Egress port-based shaping is the process by which the system delays and transmits packets to produce an even and predictable flow rate.

Each port has eight unicast and multicast queues, Class of Service (CoS) 0 to CoS 7. Traffic shaping exists on the egress CoS 6 and CoS 7, but you cannot change the configuration. CoS 6 and CoS 7 are strict priority queues, with traffic shaping for CoS 6 at 50 percent and CoS 7 to five percent of line rate.

Each feature is important to deliver DiffServ within a QoS network domain.

The switch supports an ingress flow-based policer for ACLs. For information, see [Ingress Policer and Port Rate Limiter](#) on page 3704.

Token Buckets

Tokens are a key concept in traffic control. A port-based rate limiter, policer, shaper, or an ingress flow-based policer calculates the number of packets that passed, and at what data rate. Each packet corresponds to a token, and the port-based rate limiter,

policer, shaper, or an ingress flow-based policer transmits or passes the packet if the token is available. For more information, see [Figure 201](#).

The token container is like a bucket. In this view, the bucket represents both the number of tokens that a port-rate limiter, policer, or shaper can use instantaneously (the depth of the bucket) and the rate at which the tokens replenish (how fast the bucket refills).

Each policer has two token buckets: one for the peak rate and the other for the service rate. The following figure shows the flow of tokens.

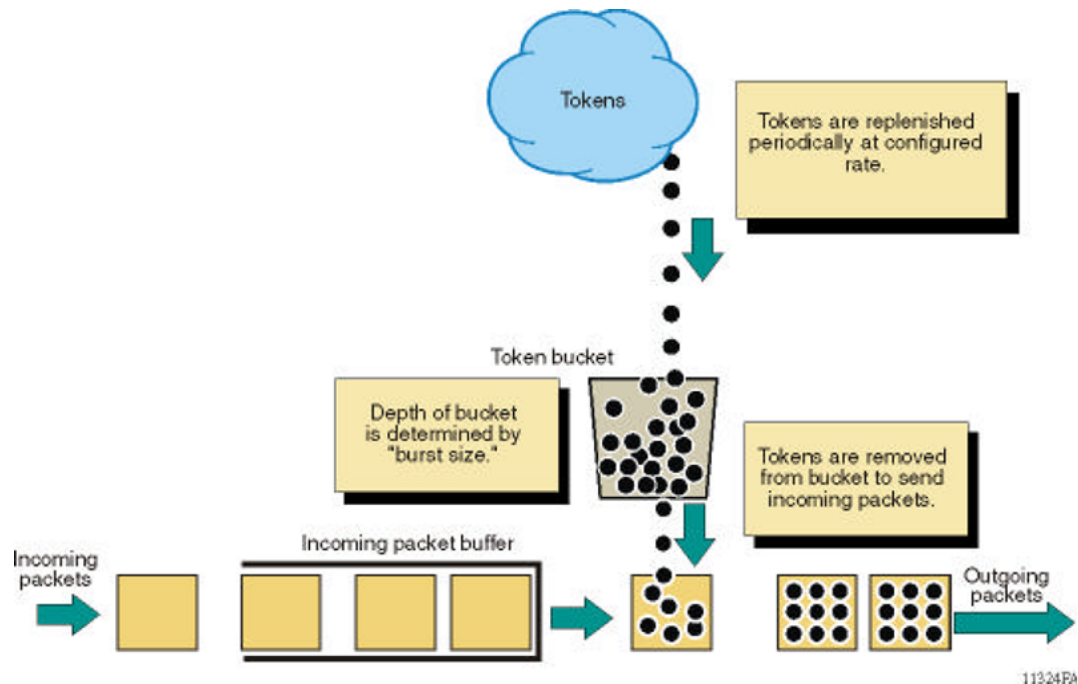


Figure 201: Token flow

Queuing

Queuing is a congestion-avoidance function that prioritizes packet delivery. Queuing ensures discriminate packet discard during network congestion, and can delay a packet in memory until the scheduled transmission.

You can use queuing to manage congestion. Congestion management involves the creation of queues, assignment of packets to the queues based on the classification of the packet, and scheduling of the packets in a queue for transmission.

The system schedules packets for transmission according to their assigned priority and the queuing mechanism configured for the interface. The scheduler determines the order of packet transmission by controlling how the system services queues with respect to each other. The switch uses 16 CPU queues (used by traffic destined to the CPU), and eight unicast and eight multicast queues for each port. The deepest queue does not go beyond 60,000 packets.

A scheduler services the eight queues for each port, using a combination of strict priority and round-robin. Queue zero through five use round robin, and queues six and seven drain completely, or up to certain rate limits.

There are eight priorities on each egress port. Each Weighted Round Robin (WRR) queue has a minimum guaranteed weight to ensure fair bandwidth allocation for each queue in case all QoS queues are utilized and there is congestion. With the default profile, Class of Service (CoS) 0 to CoS 5 are WRR, and the default weights are 5, 20, 30, 40, 50, and 50 respectively. CoS 6 and CoS 7 are strict priority queues. The switch subjects CoS 6 and CoS 7 to traffic shaping at 50 per cent and five per cent of line rate respectively.



Note

By default, rate limiting is enabled on queues 6 and 7. If you disable rate limiting on these queues on 7830 Series, no queue is considered a priority queue. Instead, all queues are managed using the Weighted Round Robin (WRR) scheduling algorithm with the following weights: 5, 20, 30, 40, 50, 50, 5.

All packets destined for the CPU are sent as unicast packets.

Front panel ports also have 8 queues. When a front panel port is oversubscribed with both unicast and multicast packets, the bandwidth is divided so that the unicast packets receive 80% of the bandwidth and multicast packets 20% of the bandwidth.

If there are no multicast packets using the reserved 20%, the bandwidth is evenly distributed between COS levels 0 to 6. This results in approximately 3.3% more bandwidth for each level. If a specific COS level is not oversubscribed, the unused bandwidth is evenly distributed between the other COS levels. The CLI provides commands to map ingressing packets to specific COS levels.



Note

Unicast packets will be tail dropped on the ingress card. Multicast packets will be dropped at egress card.

Among the WRR queues, the approximate bandwidth percentage is relative to the remaining bandwidth after subtracting the Strict Priority queues. The following tables provide the approximate bandwidth percentage based on queue congestion.

Table 234: Bandwidth for Queue 0 to 5 congestion

Queue	Approximate bandwidth percentage*
0	5%
1	10%
2	15%
3	20%
4	25%

Table 234: Bandwidth for Queue 0 to 5 congestion (continued)

Queue	Approximate bandwidth percentage*
5	25%
<p>* The approximate value assumes that all queues 0 to 5 are congested and similar packet sizes. Any queue that does not require bandwidth results in adding more bandwidth to the WRR queues.</p> <p>Percentage is relative to the total bandwidth minus the bandwidth consumed by Cos 6 and 7 (strict queues).</p>	

Table 235: Bandwidth for Queue 0 to 7 congestion

Queue	Approximate bandwidth percentage*
0	2.25%
1	4.5%
2	6.75%
3	9%
4	11.25%
5	11.25%
6	50%
7	5%
<p>* The approximate value assumes that all queues 0 to 7 are congested and similar packet sizes.</p>	

CPU Queues

The following table outlines the CPU traffic and which queue it uses.

Table 236: CPU traffic queues

Queue description	Queue ID	Queue pps, burst size	Traffic type
Other	0	1400 pps, 100 kbps	Others
BC_Other	1	1400 pps, 100 kbps	Broadcast
HI_CPU looper	2	4000 pps, 100 kbps	IP_COS01
HI_CPU	3	4000 pps, 100 kbps	IP_COS23, Multicast Data
HI_CPU	4	4000 pps, 100 kbps	IP_COS4
HI_CPU	5	4000 pps, 100 kbps	IP_COS5, IGMP, PIM Unicast
HI_CPU	6	4000 pps, 100 kbps	ARP, RARP, ND Multicast and Unicast, MLD
HI_CPU	7	4000 pps, 100 kbps	IP_COS6, Telnet, SSH

Table 236: CPU traffic queues (continued)

Queue description	Queue ID	Queue pps, burst size	Traffic type
mgmt	8	4000 pps, 200 kbps	IP_COS7, OSPF
Low rate looper	9	4000 pps, 6000 kbps	OSPF Multicast, PIM Multicast, RIP_v1, RIP_v2
Low rate 1	10	4000 pps, 12000 kbps	ISIS
Low latency	11	4000 pps, 1000 kbps	ISIS Hello, LLDP, EAP, IST Control
Low latency	12	500 pps, 32 kbps	VRRP, CFM
Low latency	13	256 pps, 32 kbps	SLPP
Low latency	14	100 pps, 32 kbps	BPDU
Low latency	15	250 pps, 32 kbps	LACP, VLACP, TUNI-extract

Queue Profiles

Table 237: Queue Profiles product support

Feature	Product	Release introduced
QoS per queue rate limiting	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

This section identifies optional ways to customize the egress queues and scheduling depending on your need to override the default configuration. You can also enable egress queue rate limiting, if desired.

Use a queue profile to apply configured egress queue parameters and modify each queue individually. You can use the queue profile to configure a minimum weight for

the queue and to enable rate limiting for the queue. The queue profile applies to all ports in the switch.



Note

If you enable rate limiting for all queues, the scheduler treats all queues as strict priority queues. If all queues are strict priority, the scheduler services the highest priority queue first until the maximum bandwidth is met, and then it services the next highest priority queue. Queue 0 is the lowest priority queue, which means that when over-subscribed, the lower priority queues are serviced last, or not at all.

The switch supports six queue profiles. The default queue profile, with the name default and ID 1, is automatically created during system startup and cannot be deleted.



Note

The egress queues with rate limiting enabled must be contiguous. For example, you can configure queues 3–6, but you cannot configure 3 and 6.

After you make a configuration change to a queue profile, you must apply the profile before the changes take effect.

Configuration Considerations

If you modify the QoS configuration for a port that is a member of MultiLink Trunking (MLT), all ports in the MLT inherit the same configuration. If you remove the port from the MLT, it keeps the QoS configuration it inherited from the MLT.

QoS Support for 10 GbE Interface in 1GbE mode

If you use QoS with a 10 gigabit Ethernet (GbE) interface and re-purpose the interface as a 1 GbE interface, you must make the necessary configuration changes to accommodate the new link speed.

Check your rate limiting and shaping configuration, if you choose to change the port link speed from 10 GbE to 1 GbE.

Review the following commands to ensure proper configuration for the port speed you use.

Command	Description
qos if-rate-limiting [port {slot/port}] rate <1000-400000000>	Configures ingress port rate limiting in kbps. Note: The range can vary depending on your hardware platform.
rate-limit broadcast <1-65535>	Configures ingress port broadcast rate limiting in packets/second.

Command	Description
rate-limit multicast <1-65535>	Configures ingress port multicast rate limiting in packets/second.
rate-limit unknown-unicast <1-65535>	Configures ingress port unknown unicast rate limiting in packets/second.
qos if-shaper [port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}] shape-rate WORD<1000-1000000000>	Specifies the shaping rate in Kb/s. If you try to configure a limit that is too high for the port speed, the switch displays the following message: Error: port slot/port, The QoS Egress shaper rate can not exceed the port capability.

802.1p and 802.1Q Best Practices

In a network, to map the 802.1p user priority bits, use 802.1Q-tagged encapsulation on customer-premises equipment (CPE). You require encapsulation because the switch does not provide classification when it operates in bridging mode.

To ensure consistent Layer 2 QoS boundaries within the service provider network, you must use 802.1Q encapsulation to connect CPE directly to the switch access node. If you do not require packet classification, use Ethernet Routing Switch 5600 to connect to the access node. In this case, configure the traffic classification functions in the Ethernet Routing Switch 5600.

At the egress access node, packets are examined to determine if their IEEE 802.1p or DSCP values must be re-marked before leaving the network. Upon examination, if the packet is a tagged packet, the IEEE 802.1p tag is configured based on the QoS level-to-IEEE 802.1p-bit mapping. For bridged packets, the DSCP is re-marked based on the QoS level.

Network Congestion and QoS Design

When you provide QoS in a network, one of the major elements you must consider is congestion, and the traffic management behavior during congestion. Congestion in a network is caused by many different conditions and events, including node failures, link outages, broadcast storms, and user traffic bursts.

At a high level, three main types or stages of congestion exist:

1. No congestion
2. Bursty congestion
3. Severe congestion

In a noncongested network, QoS actions ensure that delay-sensitive applications, such as real-time voice and video traffic, are sent before lower-priority traffic. The prioritization of delay-sensitive traffic is essential to minimize delay and reduce or eliminate jitter, which has a detrimental impact on these applications.

A network can experience momentary bursts of congestion for various reasons, such as network failures, rerouting, and broadcast storms. The switch has sufficient capacity to handle bursts of congestion in a seamless and transparent manner. If the burst is not sustained, the traffic management and buffering process on the switch allows all the traffic to pass without loss.

Severe congestion is defined as a condition where the network or certain elements of the network experience a prolonged period of sustained congestion. Under such congestion conditions, congestion thresholds are reached, buffers overflow, and a substantial amount of traffic is lost.

When you perform traffic engineering and link capacity analysis for a network, the standard design rule is to design the network links and trunks for a maximum average-peak utilization of no more than 80%. This value means that the network peaks to up to 100% capacity, but the average-peak utilization does not exceed 80%. The network is expected to handle momentary peaks above 100% capacity.

Layer 2 and Layer 3 Trusted and Untrusted Ports

You can configure interface module ports as trusted or untrusted at both Layer 2 (802.1p) or Layer 3 (DSCP) for ingress packet classification.

The switch provides eight internal QoS levels. These eight levels, numbered zero to seven, map to the queues through

- the ingress 8021p to (internal) QoS mapping table
- the ingress DSCP to (internal) QoS mapping table

To configure a port as trusted or untrusted, use the commands and the parameter values as shown in the following tables:

Layer 2 Trusted	Layer 2 Untrusted
802.1p-override	802.1p-override
disable	enable

Layer 3 Trusted		Layer 3 Untrusted	
enable-diffserv	access-diffserv *	enable-diffserv	access-diffserv *
enable	disable	disable	disable
		disable	enable
		enable	enable **

* Configure **access-diffserv** as either a core or access port. If enabled, this command specifies an access port and overrides incoming DSCP bits. If disabled, it specifies a core port that honors and services incoming DSCP bits.

** If the ingress port has **enable-diffserv** and **access-diffserv** enabled, then the packet is DSCP remarked at egress.

Layer 2 Untrusted and Layer 3 Untrusted

To configure a port as Layer 2 untrusted and Layer 3 untrusted, refer to the previous tables and assign the parameter values accordingly.

For more information, see [Table 228](#) on page 2992.

Layer 2 Untrusted and Layer 3 Trusted

To configure a port as Layer 2 untrusted and Layer 3 trusted, refer to the previous tables and assign the parameter values accordingly.

Use these configuration options to classify packet QoS through the DSCP parameter for all IP packets, whether tagged or untagged. Use this configuration when another QoS or DiffServ enabled and configured switch marks the IP packets at the edge. These already-marked packets arrive Layer 3 trusted, and the switch continues with the trust (DiffServ core port operation). For tagged packets, the system does not examine the 802.1p bits. For non-IP packets, this configuration causes classification by port QoS settings.

**Note**

For IP switched and tagged packets, use the 802.1p bits to derive the internal QoS. For untagged or routed packets, use the DSCP to derive the internal QoS.

For more information, see [Table 228](#) on page 2992.

Layer 2 Trusted and Layer 3 Trusted

To configure a port as Layer 2 trusted and Layer 3 trusted, refer to the previous tables and assign the parameter values accordingly.

Use these configuration options to classify packet QoS through DSCP for all IP packets, and through 802.1p for all tagged non IP packets. If it is an IP packet, DSCP is used. If it is a tagged non IP packet, 802.1p bits are used. If it is an untagged non IP packet, the port QoS is used.

For more information, see [Table 228](#) on page 2992.

Layer 2 Trusted and Layer 3 Untrusted

To configure a port as Layer 2 trusted and Layer 3 untrusted, refer to the previous tables and assign the parameter values accordingly.

Use these configuration options to classify packet QoS through 802.1p for all tagged packets, and port QoS levels for all untagged (IP or non-IP) packets. If the packet is an IP packet, the system does not modify or examine the DSCP parameter bits.

For more information, see [Table 228](#) on page 2992.

DiffServ Disabled

If you disable the DiffServ parameter, the system ignores the Layer 3 DSCP parameter. For more information, see [Table 228](#) on page 2992.

Broadcast, Unknown Unicast, and Multicast Traffic Bandwidth Limiters

The switch supports port-level bandwidth limiters for ingress broadcast, unknown unicast, and multicast traffic. The system drops traffic that violates the bandwidth limit. Enable this feature by configuring the rate limit, in packets per second (pps), on an individual port basis.

To see the number of packets a port drops because they exceeded the rate limit, use the following commands:

- **monitor ports error main**
- **show interfaces gigabitethernet error**

Configure Storm Control to determine the actions taken on the port after it drops packets for exceeding the rate limit. For more information, see [Storm Control](#) on page 3007.

Storm Control

Table 238: Storm Control product support

Feature	Product	Release introduced
Storm Control	4220 Series	Fabric Engine 9.3
	5320 Series	Fabric Engine 9.3 Not Supported on 5320-16P-2MXT-2X
	5420 Series	Fabric Engine 9.3
	5520 Series	Fabric Engine 9.3
	5720 Series	Fabric Engine 9.3
	7520 Series	Fabric Engine 9.3
	7720 Series	Fabric Engine 9.3
	7830 Series	Not Supported
	VSP 4900 Series	Not Applicable
	VSP 7400 Series	VOSS 9.3

A broadcast, unknown unicast, and multicast (BUM) traffic storm occurs when packets flood the LAN, which creates excessive traffic and degrades network performance. Use Storm Control to limit the amount of BUM ingress traffic on a specified physical interface.

Storm Control monitors ingress port traffic and records the number of packets dropped because they exceed the configured rate limit. The switch CPU then takes the configured action on the port to enforce traffic control policies and ensure network stability.

To use this feature on rate-limited ports, configure the following items:

- Configure a global polling interval—Periodically check the hardware counters on a port for dropped packets.
- Configure actions per port—Specify if the CPU shuts down the port or sends an SNMP trap.
- Configure a threshold per port— Specify the delta of dropped packets between polling intervals that triggers the shutdown action.

QoS and VoIP

Table 239: QoS and VoIP product support

Feature	Product	Release introduced
Automatic QoS	4220 Series	Not Supported Use Auto-sense with IP Phone support instead.
	5320 Series	Not Supported Use Auto-sense with IP Phone support instead.
	5420 Series	Not Supported Use Auto-sense with IP Phone support instead.
	5520 Series	Not Supported Use Auto-sense with IP Phone support instead.
	5720 Series	Not Supported Use Auto-sense with IP Phone support instead.
	7520 Series	Not Supported Use Auto-sense with IP Phone support instead.
	7720 Series	Not Supported Use Auto-sense with IP Phone support instead.
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

VoIP traffic requires low latency and jitter.

If you use edge routers, configure ingress ports as core ports to treat VoIP traffic appropriately. In this case, the system trusts QoS markings that apply to VoIP traffic, and the system does not re-mark QoS settings. However, if this configuration is not sufficient, you can also apply filters, route policies, or re-mark traffic.

QoS Re-marking on a Transparent Port UNI

A Transparent Port UNI port is normally configured as a Layer 2 trusted port. The T-UNI port honors incoming customer 802.1p bits and derives an internal QoS level. The 802.1p bit marking of the Backbone VLAN (B-VLAN) is derived from the internal QoS level. If the T-UNI port is set as a Layer 2 untrusted port, a best-effort queue is assigned. Customer packet headers are not modified.

The T-UNI port QoS configurations are:

- DiffServ = disable
- Layer3Trusted = access (for EDM configuration)
- access-diffserv enable (for CLI configuration)

QoS Considerations When a Port is Associated with a T-UNI I-SID

- You cannot configure `access-diffserv` and `enable diffserv` on a T-UNI port.
- When a port is associated with a T-UNI I-SID, the T-UNI QoS configuration automatically takes effect.
- When the port is removed from the T-UNI I-SID, the default port QoS configuration takes effect.

QoS Considerations When an MLT is Associated with a T-UNI I-SID

- When an MLT, static or LACP, is added to a T-UNI I-SID, the T-UNI QoS configuration take effect on all the ports of the MLT.
- When an MLT, static or LACP, is removed from a T-UNI I-SID, the port default QoS configuration is configured on all the member ports of the MLT.
- If a port is added dynamically to a T-UNI MLT, static or LACP, the port inherits the QoS properties of the T-UNI MLT ports.
- If a port is dynamically removed from a T-UNI MLT, static or LACP, the port retains the QoS configuration inherited from the MLT.

QoS and Channelization

You can enable or disable channelization on a channelization-capable port. Enabling or disabling channelization on a port resets the port QoS configuration to default values. For more information on channelization, see [Channelization](#) on page 704 and [Enable Channelization](#) on page 724.

QoS Examples and Best Practices

The sections that follow present QoS network scenarios for bridged and routed traffic over the core network.

Bridged Traffic

If you bridge traffic over the core network, you keep customer VLANs separate (similar to a Virtual Private Network). Normally, a service provider implements VLAN bridging (Layer 2) and no routing. In this case, the 802.1p-bit marking determines the QoS level

assigned to each packet. If DiffServ is active on core ports, the level of service received is based on the highest of the DiffServ or 802.1p settings.

The following cases provide sample QoS design guidelines you can use to provide and maintain high service quality in a network.

If you configure a core port, you assume that, for all incoming traffic, the QoS value is properly marked. All core switch ports simply read and forward packets; they are not re-marked or reclassified. All initial QoS markings are performed at the customer device or on the edge devices.

The following figure illustrates the actions performed on three different bridged traffic flows (that is VoIP, video conference, and email) at access and core ports throughout the network.

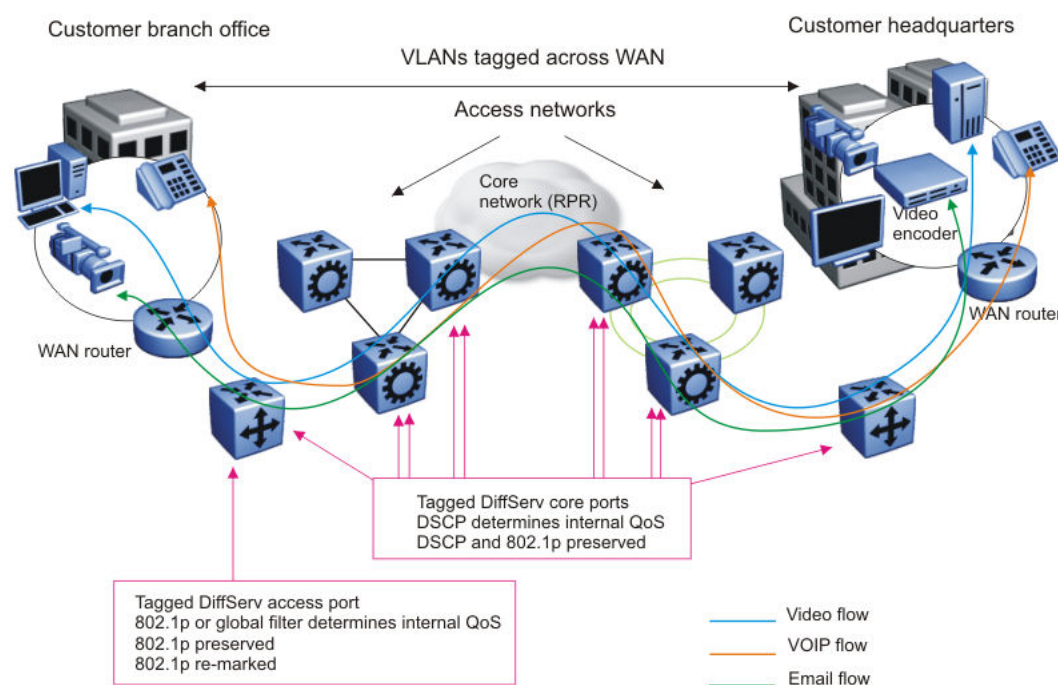


Figure 202: Trusted bridged traffic

For bridged, untrusted traffic, if you configure the port to access, mark and prioritize traffic on the access node using global filters. Reclassify the traffic to ensure it complies with the class of service specified in the SLA.

For Resilient Packet Ring (RPR) interworking, you can assume that, for all incoming traffic, the QoS configuration is properly marked by the access nodes. The core switch ports, configured as core or trunk ports, perform the RPR interworking. These ports preserve the DSCP marking and re-mark the 802.1p bit to match the 802.1p bit of the RPR. The following figure shows the actions performed on three different traffic flows (VoIP, video conference, and email) over an RPR core network.

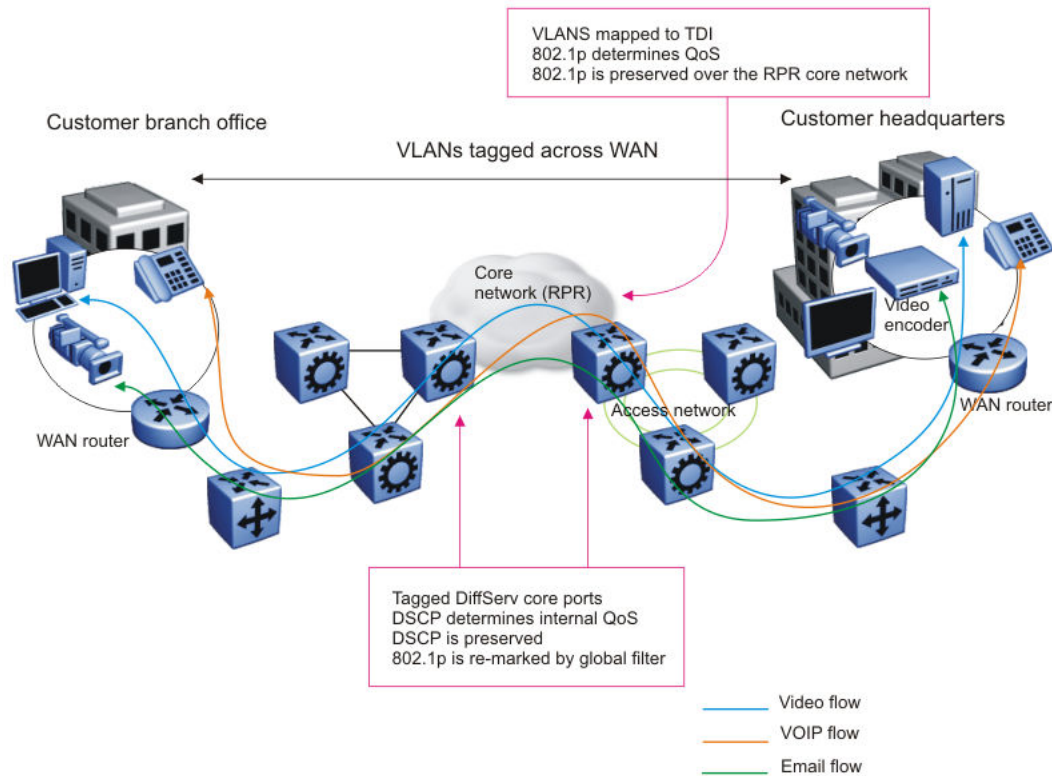


Figure 203: RPR QoS internetworking

Routed Traffic

If you route traffic over the core network, VLANs are not kept separate.

If you configure the port to core, you assume that, for all incoming traffic, the QoS configuration is properly marked. All core switch ports simply read and forward packets. The switch does not re-mark or classify the packets. The customer device or the edge devices perform all initial QoS markings.

The following figure shows the actions performed on three different routed traffic flows (that is VoIP, video conference, and email) at access and core ports throughout the network.

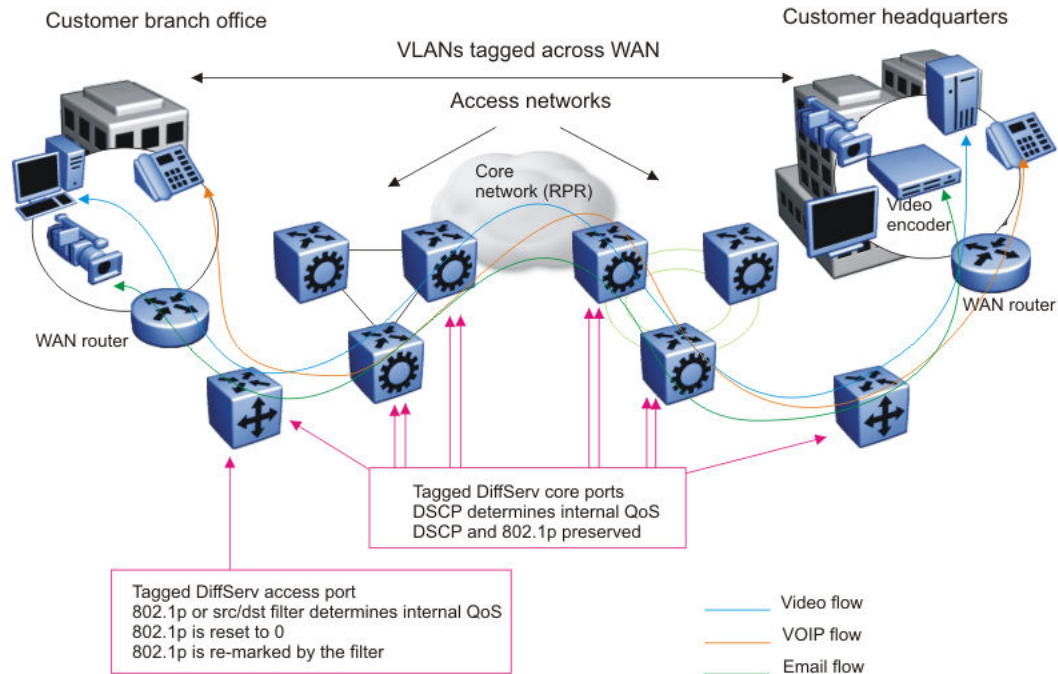


Figure 204: Trusted routed traffic

For routed, untrusted traffic, in an access node, packets that enter through a tagged or untagged access port exit through a tagged or untagged core port.

Basic DiffServ configuration using CLI

Use Differentiated Services (DiffServ) to provide appropriate Quality of Service (QoS) to specific traffic types.

Enabling DiffServ on a port

Enable DiffServ so that the system provides DiffServ-based QoS on the port. By default, DiffServ is enabled.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable DiffServ:

```
enable-diffserv [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]] [enable]
```

3. Disable DiffServ:

```
no enable-diffserv [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]] [enable]
```

Variable definitions

Use the data in the following table to use the **enable-diffserv** command.

Variable	Value
<i>enable</i>	Enables DiffServ for the specified port. The default is enabled.
<i>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configuring Layer 3 trusted or untrusted ports

Configure a port as trusted or untrusted to determine the Layer 3 QoS actions the switch performs. A trusted (core) port honors incoming Differentiated Services Code Point (DSCP) markings. An untrusted (access) port overrides DSCP markings. The default configuration is trusted.

Before You Begin

Enable DiffServ.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the port as an access port, use one of the following options:

```
no enable-diffserv [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]] [enable]
```

OR configure both parameters:

```
enable-diffserv [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]] [enable]
```

```
access-diffserv [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]] [enable]
```

3. Configure the port as a core port:

```
no access-diffserv [port {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]] [enable]
```

Variable definitions

Use the data in the following table to use the **access-diffserv** commands.

Variable	Value
<i>enable</i>	If enabled, specifies an access port and overrides incoming DSCP bits. If disabled, specifies a core port that honors and services incoming DSCP bits.
<i>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configuring Layer 2 trusted or untrusted ports

Configure a port as trusted or untrusted to determine the Layer 2 QoS actions the switch performs. A trusted port (override disabled) honors incoming 802.1p bit markings. An untrusted port (override enabled) overrides 802.1p bit markings.

Before You Begin

Enable DiffServ.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the port as Layer 2 untrusted:

```
qos 802.1p-override [enable]
```

3. Configure the port as Layer 2 trusted:

```
no qos 802.1p-override [enable]
```

Variable Definitions

Use the data in the following table to use the `qos 802.1p-override` command.

Variable	Value
<i>enable</i>	If you use this variable, the port overrides incoming 802.1p bits; if you do not use this variable, the port honors and services incoming 802.1p bits. The default is disable (Layer 2 trusted).

Viewing the port 802.1p override status

Use this procedure to view the port 802.1p override status. The system displays the port and 801.1p override status.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. View the port 802.1p override status:

```
show qos 802.1p-override
```

Example

```
Switch:1#show qos 802.1p-override
=====
Port 802.1p-Override Status
=====
PORT      802.1P OVERRIDE
-----
1/1       DISABLED
```

1/2	DISABLED
1/3	DISABLED
1/4	DISABLED
1/5	DISABLED
1/6	DISABLED
1/7	DISABLED
1/8	DISABLED
1/9	DISABLED
1/10	DISABLED
1/11	DISABLED
1/12	DISABLED
1/13	DISABLED
1/14	DISABLED
1/15	DISABLED
1/16	DISABLED

Configuring the port QoS level

Configure the port QoS level to assign a default QoS level for all traffic if the packet does not match an access control list (ACL) that re-marks the packet. If you configure port QoS levels, Layer 2 and Layer 3 traffic from the same port use the same QoS level. The default value is 1.

About This Task

For VoIP traffic, as a best practice, use QoS level 6.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the port QoS level:

```
qos level [port {slot/port[sub-port]]} <0-6>
```

Variable definitions

Use the data in the following table to use the **qos level** command.

Variable	Value
<0-6>	Specifies the default QoS level for the port traffic. The system reserves QoS level 7 for network control traffic. The default is 1.
<code>port {slot/port[/sub-port]}</code>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Basic DiffServ configuration using EDM

Use DiffServ to implement classification and mapping functions at the network boundary or access points to regulate packet behavior. You can configure a port as a trusted (core) or an untrusted (access) port at both Layer 2 and Layer 3.

You can also perform many of the procedures in this section on the Interface tab for the selected port. The procedures in this section show only one configuration method.

Enable DiffServ for a Port

About This Task

Enable DiffServ so that the switch provides DiffServ-based Quality of Service (QoS) on the port.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **Port QoS Config**.
3. In the row for the port, double-click the cell in the **DiffServ** column.
4. Select **true**.
5. Select **Apply**.

Port QoS Config *Field Descriptions*

Use the data in the following table to use the **Port QoS Config** tab.

Name	Description
Index	Specifies an index value that uniquely identifies a port.
DiffServ	Specifies whether DiffServ is enabled (true) or disabled (false) on the port. The default is true. This variable works in conjunction with Layer3Trust. The DiffServ variable is a global parameter that affects QoS DSCP operations. If the DiffServ parameter is false (DiffServ disabled), the system does not use the DSCP parameter for classification or modify it. If this variable is true, it activates the Layer3Trust parameter.
Layer3Trust	Configures the Layer 3 trusted port as an access or core port. The default is core. Core configures the port to a trusted state and access configures the port to an untrusted state. The DiffServ parameter determines the operation of this variable. If DiffServ is false, Layer3Trust has no effect; no modification of the DSCP or TOS bits occurs. If DiffServ is true, the core and access configuration take effect.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled (true) or disabled (false) on the port. The default is false. This variable primarily affects tagged packet treatment. If Layer2Override8021p is false, the port trusts the 802.1p-bits portion of a Q-tagged packet. The port trusts the 802.1p-bits marking regardless of the port setting (tagged or untagged); however, if the discard tagged packets parameter (DiscardTaggedFrames) on an untagged port is true, the system discards the packet. If Layer2Override8021p is true, the port does not trust the 802.1p bit marking. In this case, the QoS operation depends on other parameters, such as the port QoS level.
QosLevel	Specifies the QoS level to use when the system processes packets carried on this port. Values range from level 0–6 (the system reserves 7 for network control traffic). The default is 1.

Configure Layer 3 Trusted or Untrusted Ports

About This Task

Configure a port as trusted or untrusted to determine the Layer 3 QoS actions the switch performs. A trusted port honors incoming DSCP markings. An untrusted port overrides DSCP markings. The default is trusted.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **Port QoS Config**.
3. In the row for the port, double-click the cell in the **Layer3Trust** column.

4. Select **core** (trusted) or **access** (untrusted) as the port setting.
5. Select **Apply**.

Configure Layer 2 Trusted or Untrusted Ports

About This Task

Configure a port as trusted or untrusted to determine the Layer 2 QoS actions the switch performs. A trusted port (override false) honors incoming 802.1p bit markings. An untrusted port (override true) overrides 802.1p bit markings.

Procedure

1. In the navigation pane, expand **Configuration > QoS**.
2. Select **Port QoS Config**.
3. In the row for the port, double-click the cell in the **Layer2 Override 8021p** column.
4. To configure the port as a Layer 2 untrusted port, select **true**. To configure it as a Layer 2 trusted port, select **false**.
By default, all ports are Layer 2 trusted (Layer2 Override 8021p is false).
5. Select **Apply**.

Configure the Port QoS Level

About This Task

Use the default port QoS level to assign a default QoS level for all traffic, if the packet does not match an access control list (ACL) to remark the packet.

Procedure

1. In the navigation pane, expand **Configuration > QoS**.
2. Select **Port QoS Config**.
3. In the row for the port, double-click the cell in the **QosLevel** column.
4. Select the new level.
5. Select **Apply**.

QoS Configuration Using CLI

Use the procedures in this section to configure Quality of Service (QoS) on the switch.

Configure Broadcast, Unknown Unicast, and Multicast Bandwidth Limiting

Configure broadcast, unknown unicast, and multicast bandwidth limiting to limit the amount of ingress traffic on a port. The switch drops traffic that violates the bandwidth limit.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the bandwidth limiter to limit the amount of ingress traffic on a port:

```
rate-limit [port {slot/port[-slot/port]}[,...][slot/all][all]]  
{broadcast | multicast | unknown-unicast <1-65535>}
```

3. Verify the configuration:

```
show interfaces gigabitEthernet rate-limit [<1-4059> | {slot/port[/  
sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

Example

View the bandwidth limiter configuration for port 1/15:

```
Switch:1>show interfaces gigabitEthernet rate-limit 1/15
```

Port Rate(PPS) Limit							
PORT	TYPE	PPS-RATE	ENABLED/DISABLED	SHUTDOWN	THRESHOLD	SECS SINCE SHUTDOWN	TRAP-INTERVAL
1/15	BROADCAST	900	ENABLED	90	1	0	10
1/15	MULTICAST	900	ENABLED	90	1	0	10
1/15	UNKNOWN-UC	900	ENABLED	90	1	0	10

Variable Definitions

The following table defines parameters for the **rate-limit** command.

-

Variable	Value
<1-65535>	Specifies the bandwidth limit for broadcast, unknown unicast, and multicast traffic, in packets per second (pps).
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure Storm Control

Before You Begin

Configure rate limiting on the port.

About This Task

Configure Storm Control to determine the actions taken on the port after it drops packets for exceeding the rate limit.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Configure the interval at which to check hardware counters for dropped packets:


```
rate-limit poll-interval <3-300>
```
3. Verify the polling configuration:


```
show rate-limit
```
4. Display the drop counter:


```
show interfaces gigabitethernet error [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}
```

5. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

6. Configure the action to take after Storm Control detects dropped packets on the port:

- Shut down the port: `rate-limit action shutdown [<60-3600>]`
- Configure the delta of dropped packets per polling interval: `rate-limit action threshold <1-4294967295>`
- Specify the SNMP trap notification frequency: `rate-limit action trap-interval <0-1000>`

7. Verify the port configuration:

```
show interfaces gigabitethernet rate-limit
```

Example

View the number of packets dropped on port 1/15 because they exceeded the configured rate limit.

```
Switch:1>enable
Switch:1#show interfaces gigabitethernet error 1/15
=====
Port Ethernet Error
=====
PORT  ERROR  ERROR  FRAMES  TOO    LINK    CARRIER  CARRIER  SQETEST  IN    RATE-LIMIT
NUM   ALIGN  FCS    LONG    SHORT  FAILURE  SENSE     ERROR     ERRORS   DISCARD  DROP-COUNTER
-----
1/15  0      0      0       0      2       0        0        0        7800   780
```

Configure Storm Control on port 1/15.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#rate-limit poll-interval 10
Switch:1(config)#interface gigabitEthernet 1/15
Switch:1(config-if)#rate-limit action shutdown 90
Switch:1(config-if)#rate-limit action trap-interval 10
```

Verify the configuration.

```
Switch:1>show rate-limit
=====
RATE LIMIT
=====
Rate Limit Poll Interval (seconds):10
Switch:1>show interfaces gigabitEthernet rate-limit 1/15
=====
Port Rate(PPS) Limit
=====
PORT  TYPE          PPS-RATE  ENABLED/DISABLED  SHUTDOWN  THRESHOLD  SHUTDOWN  SECS SINCE  TRAP-INTERVAL
```

1/15	BROADCAST	900	ENABLED	90	1	0	10
1/15	MULTICAST	900	ENABLED	90	1	0	10
1/15	UNKNOWN-UC	900	ENABLED	90	1	0	10

Variable Definitions

The following table defines parameters for the **rate-limit poll-interval** command.

Variable	Value
<3-300>	Specifies the interval, in seconds, at which this feature checks the hardware counters for packets dropped because they exceeded the configured port rate limit. The default is 3 seconds.

The following table defines parameters for the **rate-limit action** command.

Variable	Value
shutdown [<60-3600>]	Configures the port to shut down, and specifies how long, in seconds, the port remains shut down, after Storm Control detects dropped packets. If you do not specify the time value, the port remains shut down until you manually re-enable it. The default is to keep the port up, which is the no rate-limit action shutdown variation of the command.
threshold <1-4294967295>	Specifies a threshold value for the delta of dropped packets per polling interval that triggers the shut down. The default is 1.
trap-interval <0-1000>	Specifies the frequency, in polling cycles, of SNMP trap generation. Configure the value to 0 to disable traps and not generate log entries. The default is disabled.

Viewing the Port-Based Shaper Information

About This Task

Use this procedure to view the port-based shaper information. The system displays the port, egress rate limit in Kbps, and the rate limit status.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the ingress port-rate limit information:
show qos shaper interface gigabitEthernet [{slot/port[-slot/port]}
[,...][slot/all][all]]

Example

```
Switch:1#show qos shaper interface gigabitEthernet 1/1
=====
Port Egress Rate-Limiting (Shape)
=====
-----
PORT      EGRESS RATE-LIMIT (kbps)    ENABLED/DISABLED
-----
1/1       0                          DISABLED
```

Configure the Port-Based Shaper

Use port-based shaping to rate-limit all outgoing traffic to a specific rate.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure port-based shaping:

```
qos if-shaper [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]] shape-rate WORD<1000-1000000000>
```

Variable Definitions

The following table defines parameters for the **qos if-shaper** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Specifies the slot and port number to which to apply shaping. This variable is optional.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>shape-rate WORD<1000-100000000></code>	<p>Specifies the shaping rate in Kb/s. If you try to configure a limit that is too high for the port speed, the switch displays the following message:</p> <p>Error: port slot/port, The QoS Egress shaper rate can not exceed the port capability.</p> <p>The default is 0, which means shaping is disabled on the port.</p>

Configure Ingress Mappings

You can modify the ingress mappings to change traffic priorities. However, as a best practice, use the default mappings.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure 802.1p bit to QoS ingress mappings:

```
qos ingressmap lp <0-7> <0-6>
```
3. Configure DSCP to QoS ingress mappings:

```
qos ingressmap ds <0-63> <0-6>
```
4. Ensure the configuration is correct:

```
show qos ingressmap [lp <0-7>]
```

```
show qos ingressmap [ds <0-63>]
```

Variable Definitions

Use the data in the following table to use the **qos ingressmap** command.

Variable	Value
<i>lp</i> <0-7> <0-6>	<p>Maps the IEEE 802.1p bit to QoS level. Each QoS level has a default IEEE 1P value:</p> <ul style="list-style-type: none"> • level 0—1 • level 1—0 • level 2—2 • level 3—3 • level 4—4 • level 5—5 • level 6—6 <p>The system reserves level 7 for Network Control. To use the default configuration, use the default option in the command:</p> <pre>default qos ingressmap lp</pre>
<i>ds</i> <0-63> <0-6>	<p>Maps the DS byte to QoS level. The system reserves level 7 for Network Control. To use the default configuration, use the default option in the command:</p> <pre>default qos ingressmap ds</pre>

Configure Egress Mappings

You can modify the egress mappings to change traffic priorities. However, as a best practice, use the default mappings.

Procedure

1. Enter Global Configuration mode:


```
enable
```

```
configure terminal
```
2. Configure QoS to 802.1p bit egress mappings:


```
qos egressmap lp <0-6> <0-7>
```
3. Configure QoS to DSCP egress mappings:


```
qos egressmap ds <0-7> WORD<1-6>
```
4. Ensure the configuration is correct:


```
show qos egressmap [lp <0-7>]
```

```
show qos egressmap [ds <0-7>]
```


Variable Definitions

Use the data in the following table to use the **qos egressmap** command.

Variable	Value
<i>1p</i> <0-6> <0-7>	<p>Maps the QoS level to IEEE 802.1p bit. Each QoS level has a default IEEE 1P value:</p> <ul style="list-style-type: none"> • level 0—1 • level 1—0 • level 2—2 • level 3—3 • level 4—4 • level 5—5 • level 6—6 <p>The system reserves level 7 for Network Control. To use the default configuration, use the default option in the command:</p> <pre>default qos egressmap 1p</pre>
<i>ds</i> <0-7> <i>WORD</i> <1-6>	<p>Maps the QoS level to DS byte. You can specify the DSCP in either hexadecimal, binary, or decimal format. To use the default configuration, use the default option in the command:</p> <pre>default qos egressmap ds</pre>

View Port Egress CoS Queue Statistics

View the port egress CoS queue statistics. The system displays the statistics of the forwarded packets and bytes, and the dropped packets and bytes.

About This Task

If you disable rate limiting on queue 6, bandwidth is shared based on the weights from queues 0 through 6. Based on equal weights assigned for queues 4 through 6, equal amounts of traffic egress queues 4 through 6. The system displays Queue 6 to provide much lower throughput than 50% of port bandwidth.



Note

If there is egress congestion when you enable the **boot config flags flow-control-mode** command, then the system does not drop the packets at egress, and does not increment the counters under the **show qos cosq-stats interface {slot/port[/sub-port]} [-slot/port[/sub-port]] [, ...] [slot/all] [all]** command output. Instead, the system drops the packets at the ingress port and the system counts them under the **INDISCARD** column of **show interfaces gigabitEthernet error {slot/port[/sub-port]} [-slot/port[/sub-port]] [, ...] [slot/all] [all]** command output.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View the port egress CoS queue statistics:
`show qos cosq-stats interface <PT_PORT>`



Note

The show command output varies based on your hardware platform. The show command output displays Out Packets and Out Bytes per interface, which shows the number of unicast packets sent out on each queue for an egress port.

Variable Definitions

Use the data in the following table to use the **show qos cosq-stats interface <PT_PORT>** command.

Variable	Value
<PT_PORT>	PT indicates the slot number; PORT indicates the port number.

Clear Port Egress CoS Queue Statistics

Clear the port egress CoS queue statistics in the hardware.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Clear the port egress CoS queue statistics:
`clear qos cosq-stats interface <PT_PORT>`

Variable Definitions

Use the data in the following table to use the **clear qos cosq-stats interface <PT_PORT>** command.

Variable	Definition
<PT_PORT>	PT indicates the slot number; PORT indicates the port number.

View CPU Queue Statistics

View the statistics of the forwarded packets and bytes, and the dropped packets and bytes for the traffic sent toward CP. The queue assignment is based on the protocol types, not on the internal CoS value. These statistics are useful for debugging purposes.

When a neighbor transitions to the STALE state, to initiate Neighbor Unreachability detection (NUD), a duplicate copy of the traffic destined to this neighbor is sent to the switch Control Processor (CP) on a low priority queue (queue 0). The original packet is forwarded to this neighbor. When NUD is initiated, the hardware records are updated and the traffic is no longer sent to the CP. When a high rate of such traffic is sent to CP, the switch can drop some of these packets due to the in built CP rate limiting feature, which protects the CP from DOS attacks.

Use the command **show qos cosq-stats cpu-port** to view drop statistics on the CPU queue. This design does not result in loss of traffic.

Use the command **ipv6 nd reachable-time <0-3600000>** to increase the default value of 3000 milliseconds which in turn delays the scenario of data path sending STALE neighbor destined packets to the CP.

About This Task



Note

If you disable rate limiting on queue 6, bandwidth is shared based on the weights from queues 0 through 6. Based on equal weights assigned for queues 4 through 6, equal amounts of traffic egress queues 4 through 6. The system displays Queue 6 to provide much lower throughput than 50% of port bandwidth.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View the CPU queue statistics:
`show qos cosq-stats cpu-port`



Note

The show command output varies based on your hardware platform. The output displays Out Packets and Out Bytes per interface, which shows the number of unicast packets sent out on each queue for an egress port.

Clear CPU Queue Statistics

Clear the CPU queue statistics.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Clear CPU queue statistics:
`clear qos cosq-stats cpu-port`

Configure an Egress QoS Queue Profile

Configure a queue profile to apply the configured egress queue parameters to queues and ports.

About This Task

After you make a configuration change to a queue profile, you must apply the profile before the changes take effect.



Note

The switch supports six queue profiles. The default queue profile, with the name default and ID 1, is automatically created during system startup and cannot be deleted.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the minimum weight for a specific queue:

```
qos queue-profile queue <1-6> <0-7> min-weight <1-100>
```

3. Enable rate limiting on a weighted queue:

```
qos queue-profile queue <1-6> <0-7> rate-limit-enable
```

4. Add a queue-profile port member:

```
qos queue-profile <1-6> member add {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}
```

5. Add a queue-profile name:

```
qos queue-profile <1-6> name WORD<0-64>
```

6. Apply the queue profile:

```
qos queue-profile <1-6> apply
```

7. Verify the egress queue configuration:

```
show qos queue-profile [<1-6> queue <0-7>|all]
```

8. (Optional) Configure the default values for an egress queue:

- Configure the default minimum weight using one of the following commands:

```
default qos queue-profile queue <1-6> <0-7> min-weight
```

```
no qos queue-profile queue <1-6> <0-7> min-weight
```

- Configure the default rate limiting on a weighted queue using one of the following commands:

```
default qos queue-profile queue <1-6> <0-7> rate-limit-enable
```

```
no qos queue-profile queue <1-6> <0-7> rate-limit-enable
```

**Note**

By default, rate limiting is enabled on queues 6 and 7. If you disable rate limiting on these queues on 7830 Series, no queue is considered a priority queue. Instead, all queues are managed using the Weighted Round Robin (WRR) scheduling algorithm with the following weights: 5, 20, 30, 40, 50, 50, 5.

Examples

Configure the queue profile for queue 6 to use a weight of 20.

```
Switch:1(config)#qos queue-profile 6
Switch:1(config)#qos queue-profile queue 6 1 min-weight 20
Switch:1(config)#qos queue-profile 6 apply
```

View the queue profile configuration.

```
Switch:1#show qos queue-profile
```

Qos Queue Profile		
Profile ID	Profile Name	Profile Port List
1	default	1/1-1/42,2/1-2/2
6		profile-6

The value in the **Weight Applied** column is the maximum bandwidth for Strict Priority queues (Rate Limit Enabled) and the weight in the WRR scheduling algorithm applied per queue. Strict Priority queues are serviced first and limited to the configured percentage of bandwidth. WRR queues (no rate limit applied) have no upper boundary, other than the delta between the available bandwidth per port and the one consumed by Strict Priority queues. For WRR queues, there is no association with the bandwidth percentage servicing this queue; the bandwidth that can be served is relative to the available bandwidth at a certain moment and the load of other queues.

```
Switch:1(config)#show qos queue-profile 1 queue 1
```

Qos Queue Profile Table		
-------------------------	--	--

```

=====
Profile Profile Queue Weight  Weight      Rate-limit Rate-limit
ID      Name   ID    Applied Configured Applied   Configured
-----
1      default 1      0      20      ENABLE   DISABLE
Switch:1(config)#show qos queue-profile 1 queue all

=====
                        Qos Queue Profile Table
=====

Profile Profile      Queue Weight  Weight      Rate-limit Rate-limit
ID      Name        ID    Applied Configured Applied   Configured
-----
1      default      0      10      10      DISABLE   DISABLE
1      default      1      20      20      DISABLE   DISABLE
1      default      2      30      30      DISABLE   DISABLE
1      default      3      40      40      DISABLE   DISABLE
1      default      4      50      50      DISABLE   DISABLE
1      default      5      50      50      DISABLE   DISABLE
1      default      6      50      50      DISABLE   ENABLE
1      default      7      5       5       ENABLE    ENABLE

```

Variable Definitions

The following table defines parameters for the **qos queue-profile queue** command.

Variable	Value
<1-6>	Specifies the queue profile ID. Note: The switch supports six queue profiles. The default queue is 1.
<0-7>	Specifies the egress queue to configure.
<i>min-weight</i> <1-100>	Configures the queue weight for weighted round robin (WRR), or the rate-limit in percentage of the link rate for queue shaping enabled on the queue. For rate-limit-enabled queues, 50 indicates 50% bandwidth. The overall rate-limit-enabled queues cannot sum higher than 100 min-weight. If rate-limiting is not enabled, this parameter configures the weight in the WRR scheduler. These values are not associated with bandwidth percentage. The following list identifies the default minimum weight for each queue: <ul style="list-style-type: none"> • Queue 0 — 5 • Queue 1 — 20 • Queue 2 — 30 • Queue 3 — 40 • Queue 4 — 50 • Queue 5 — 50 • Queue 6 — Rate limited to 50% of configured shaper rate • Queue 7 — Rate limited to 5% of configured shaper rate

Variable	Value
<i>member add</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Specifies a port member of the queue-profile to add or remove.
<i>name</i> WORD<0-64>	Specifies a profile name.
<i>rate-limit-enable</i>	Enables rate limiting on the queue. By default, rate limiting is enabled for queues 6 and 7 only; it is disabled for queues 0 through 5.

The following table defines parameters for the **show qos queue-profile** command.

Variable	Value
<1-6>	Specifies the queue profile ID. If you do not include a queue profile ID, the command output displays all configured profiles. Note: The switch supports six queue profiles. The default queue is 1.
<0-7>	Specifies the egress queue. Displays configuration settings of the specified egress queue.
<i>all</i>	The command output displays the configuration settings of all eight egress queues of the queue profile.

View Logical Interface CoS Queue Statistics

View the QoS CoS queue statistics for IS-IS logical interfaces. These statistics are useful for debugging purposes.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the logical interface queue statistics:
show qos cosq-stats logical-intf [isis <1-255>]

Clear Logical Interface CoS Queue Statistics

Clear the QoS CoS queue statistics for IS-IS logical interfaces.

Procedure

1. Enter Privileged EXEC mode:
enable

2. Clear the logical interface queue statistics:

```
clear qos cosq-stats logical-intf [isis <1-255>]
```

QoS Configuration Using EDM

Configure Quality of Service (QoS) to allocate network resources where you need them most.

Configure Port-Based Shaping

Configure egress port-based shaping to bind the maximum rate at which traffic leaves the port.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. From **EgressRateLimitState**, select **enable**.
6. In **EgressRateLimit**, type an egress rate limit in kilobits per second (Kb/s).
7. Select **Apply**.

Graph Stat Rate Limit Statistics for a Port

View stat rate limit statistics to view the total dropped packets and bytes.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **Stat Rate Limit** tab.
5. Select one or more values.
6. Select the type of graph to create.

Stat Rate Limit Field Descriptions

Use the data in the following table to use the **Stat Rate Limit** tab.

Name	Description
DropPktRate	Indicates the drop packet rate.
DropByteRate	Indicates the drop byte rate.
DropTotalBytes	Indicates the total bytes dropped.
DropTotalPkts	Indicates the total packets dropped.

Modify Ingress 802.1p to QoS Mappings

About This Task

Modify the ingress mappings to change traffic priorities. As a best practice, use the default mappings.

Do not change the default values. If you change the values, make sure that the values are consistent on all other devices in the network. Inconsistent mapping of table values can result in unpredictable service levels.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **Mapping Tables**.
3. Select the **Ingress 8021p to QoS** tab.
4. Double-click a QosLevel to change the value.
5. Select **Apply**.

Ingress 8021p To QoS Field Descriptions

Use the data in the following table to use the **Ingress 8021p To QoS** tab.

Name	Description
Inleee8021P	Specifies the value of the IEEE 802.1p bit of the incoming packet.
QosLevel	Specifies the equivalent egress QoS level (0–7).

Modify Ingress DSCP to QoS Mappings

About This Task

Modify the ingress Differentiated Services Code Point (DSCP) to QoS mappings to change traffic priorities. As a best practice, use the default mappings. Changes to the mapping table take effect after you restart the system.

Do not change the default values. If you change the values, make sure that the values are consistent on all other devices in the network. Inconsistent mapping of table values can result in unpredictable service levels.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **Mapping Tables**.
3. Select the **Ingress Dscp To QoS** tab.
4. Double-click a QosLevel to change the value.
5. Select **Apply**.

Ingress Dscp To QoS Field Descriptions

Use the data in the following table to use the **Ingress Dscp To QoS** tab.

Name	Description
InDscp	Specifies the value of the DiffServ codepoint (in decimal format) in the IP header of the incoming packet.
InDscpBinaryFormat	Specifies the value of the DiffServ codepoint (in binary format) in the IP header of the incoming packet.
QosLevel	Specifies the equivalent QoS level.

Modify Egress QoS to 802.1p Mappings

About This Task

Modify the egress mappings to change the mappings between the QoS levels and the IEEE 802.1p bits.

Do not change the default values. If you change the values, make sure that the values are consistent on all other devices in the network. Inconsistent mapping of table values can result in unpredictable service levels.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **Mapping Tables**.
3. Select the **Egress QoS to 8021p** tab.
4. Double-click **Outleeee8021P** to change the value.
5. Select **Apply**.

Egress QoS to 8021p Field Descriptions

Use the data in the following table to use the **Egress QoS to 8021p** tab.

Name	Description
QosLevel	Specifies the QoS level of the outgoing packet.
Outleeee8021P	Specifies the equivalent value of the IEEE 802.1p bit.

Modify Egress QoS to DSCP Mappings

About This Task

Modify the egress QoS to DSCP mappings to change traffic priorities. As a best practice, use the default mappings.

Do not change the default values. If you change the values, make sure that the values are consistent on all other devices in the network. Inconsistent mapping of table values can result in unpredictable service levels.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **Mapping Tables**.
3. Select the **Egress QoS To Dscp** tab.
4. Double-click **OutDscp** to change the value.
5. Select **Apply**.

Egress QoS To Dscp Field Descriptions

Use the data in the following table to use the **Egress QoS To Dscp** tab.

Name	Description
QosLevel	Specifies the QoS level of the outgoing packet.
OutDscp	Specifies the equivalent value of the DiffServ code point (in decimal format).
OutDscpBinaryFormat	Specifies the equivalent value of the DiffServ code point (in binary format).

View Port Egress CoS Queue Statistics

About This Task

Use the following procedure to retrieve the port egress CoS queue statistics. The system displays the statistics of the forwarded packets and bytes, and the dropped packets and bytes.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **CoS Queue Stats**.
3. Select the **Interface** tab.

Interface Field Descriptions

The following table describes the fields from the CoS Queue Stats **Interface** tab.

Name	Description
Index	Indicates the loopback port number from 192(1/1) to 241(1/50).
ClearStat	Clears the port egress statistics.
Que<0-7>OutPackets	Indicates the out packets by CoS queue number 0-7.

Name	Description
Que<0–7>OutBytes	Indicates the out bytes by CoS queue number 0–7.
Que<0–7>DropPackets	Indicates the drop packets by CoS queue number 0–7.
Que<0–7>DropBytes	Indicates the drop bytes by CoS queue number 0–7.

Clear CPU Statistics for the Chassis

Use the following procedure to clear the CPU statistics for the chassis.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **CoS Queue Stats**.
3. Select the **CPU-Stats-Clear** tab.
4. Select **CpuStatsClear**.
5. Select **Apply**.

CPU-Stats-Clear Field Descriptions

Use the data in the following table to use the **CPU-Stats-Clear** tab.

Name	Description
CpuStatsClear	Clears CPU QOS queue statistics.

View CPU Queue Statistics

Use the following procedure to retrieve the statistics of the forwarded packets and bytes, and the dropped packets and bytes for the traffic sent toward CP. The queue assignment is based on the protocol types, not on the internal CoS value. These statistics are useful for debugging purposes.

Procedure

1. In the navigation pane, expand **Configuration > QOS**.
2. Select **CoS Queue Stats**.
3. Select the **CPU-Port** tab.

CPU-Port Field Descriptions

Use the data in the following table to use the **CPU-Port** tab.

Name	Description
Index	Indicates the CoS queue number.
OutPackets	Indicates the out packets for the CPU port.

Name	Description
OutBytes	Indicates the out bytes for the CPU port.
DropPackets	Indicates the drop packets for the CPU port.
DropBytes	Indicates the drop bytes for the CPU port.

Configure an Egress QoS Queue Profile

Configure a queue profile to apply the configured egress queue parameters to queues and ports. You must apply the profile before the changes take effect.

About This Task

The switch supports six queue profiles. The default queue profile, with the name default and ID 1, is automatically created during system startup and cannot be deleted.

Procedure

1. In the navigation pane, expand **Configuration > QoS**.
2. Select **Queue Profile**.
3. Select the **Queue Profile** tab.
4. Select **Insert**.
5. In **Id**, type the queue profile value.
6. (Optional) In **Name**, specify a name for the queue profile.
7. To add a port to this queue, select the **PortList** ellipsis (...), choose a port or ports, and then select **Ok**.
8. Select **Apply**.
9. Select **Insert**.

Queue Profile Field Descriptions

Use the data in the following table to use the **Queue Profile** tab.

Field	Description
Id	Specifies the ID for the queue profile.
Name	Specifies the queue profile name.
Apply	Applies the queue profile.
PortList	Indicates the port members of the queue profile.

Edit Queue Profile Information

About This Task

Use the following procedure to edit queues of a queue profile, to configure a queue weight or enable rate limiting on the queue.



Note

After you make the configuration changes, you must apply the queue profile before the changes take effect.

Procedure

1. In the navigation pane, expand **Configuration > QoS**.
2. Select **Queue Profile**.
3. Update a queue to configure queue weight or rate limiting:
 - a. Select the **Queue** tab.
 - b. Edit **AdminWeight** and **AdminRateLimitStatus** by double-clicking on them, and then selecting or typing the new value.
 - c. Select **Apply**.
4. Apply the queue profile for the queue configuration to take effect:
 - a. Select the **Queue Profile** tab.
 - b. In **Apply**, double-click and select **true**.
 - c. Select **Apply**.
5. Select the **Queue** tab to verify updates to **OperWeight** and **OperRateLimitStatus**, for the respective queue.

Queue Field Descriptions

Use the data in the following table to use the **Queue** tab.

Field	Description
PId	Displays the queue profile ID.
Id	Displays the queue ID.

Field	Description
AdminWeight	<p>Configures the queue weight for weighted round robin (WRR), or the rate-limit in percentage of the link rate for queue shaping enabled on the queue.</p> <p>For rate-limit-enabled queues, 50 indicates 50% bandwidth. The overall rate-limit-enabled queues cannot sum higher than 100 min-weight.</p> <p>If rate-limiting is not enabled, this field configures the weight in the WRR scheduler. These values are not associated with bandwidth percentage.</p> <p>The following list identifies the default minimum weight for each queue:</p> <ul style="list-style-type: none"> • Queue 0 — 5 • Queue 1 — 20 • Queue 2 — 30 • Queue 3 — 40 • Queue 4 — 50 • Queue 5 — 50 • Queue 6 — Rate limited to 50% of configured shaper rate • Queue 7 — Rate limited to 5% of configured shaper rate
OperWeight	Displays the operational weight of the profile, described as a percentage.
AdminRateLimitStatus	<p>Enables rate limiting on the queue. By default, rate limiting is enabled (true) for queues 6 and 7 only; it is disabled (false) for queues 0 through 5.</p> <p>AuthorsNote: The following Note was added for 7830 in 9.3 for VOSS-33291.</p> <p>Note: By default, rate limiting is enabled on queues 6 and 7. If you disable rate limiting on these queues on 7830 Series, no queue is considered a priority queue. Instead, all queues are managed using the Weighted Round Robin (WRR) scheduling algorithm with the following weights: 5, 20, 30, 40, 50, 50, 5.</p>
OperRateLimitStatus	Displays the operational status of the queue rate limit.

Configure Port Rate Limits

Configure broadcast, unknown unicast, and multicast bandwidth limiting to limit the amount of ingress traffic on a port. The switch drops traffic that violates the bandwidth limit.

Procedure

1. On the Device Physical View, select a port or multiple ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Rate Limiting** tab.
5. In the row for the required traffic type, configure the bandwidth limit in **AllowedRatePps**.
6. In the row for the required traffic type, select true for **Enable**.
7. Select **Apply**.

Rate Limiting Field Descriptions

Use the data in the following table to use the **Rate Limiting** tab.

Name	Description
Index	Displays the port number.
TrafficType	Displays the type of traffic being rate limited, either broadcast, unknown unicast, or multicast traffic.
AllowedRatePps	Specifies the allowed traffic rate limit for the port in packets per second (pps). 1–65535 configures the limit in packets for each second. The default is 25.
Enable	Enables or disables rate limiting. The default is false (no rate limit).

Configure Rate Limits on an Extreme Integrated Application Hosting Port



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to configure the rate limit of broadcast or multicast packets and determine the total bandwidth limit on the Extreme Integrated Application Hosting (IAH) port.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **Rate Limiting** tab.
4. In the **AllowedRatePps** column, enter a time duration for the specific Extreme Integrated Application Hosting port.
5. In the **Enable** column, select **true** to enable rate limiting for the specific IAH port.
6. Select **Apply**.

Rate Limiting *Field Descriptions*

Use the data in the following table to use the **Rate Limiting** tab.

Name	Description
Index	Specifies the Extreme Integrated Application Hosting (IAH) port.
TrafficType	Displays the type of traffic being rate limited, either broadcast, unknown unicast, or multicast traffic.
AllowedRatePps	Specifies the allowed traffic rate limit for the IAH port in packets per second (pps).
Enable	Enables or disables rate limiting for the specific IAH port. The default is false (disabled).

Configure the Storm Control Poll Interval**About This Task**

Configure the interval at which to check hardware counters for dropped packets.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System Control** tab.
4. In **RateLimitActionPollInterval**, type the number of seconds.
5. Select **Apply**.

System Control *Field Descriptions*

Use the data in the following table to use the **System Control** tab.

Name	Description
TcpTimestampEnable	Enables or disables the TCP timestamp. The timestamp is enabled by default. The system displays the following warning message after you apply a new configuration: Warning: Existing TCP connections won't be affected. A config save and reboot is required to apply this configuration for all TCP connections.
KeepaliveTime	Specify the TCP keepalive time in seconds. The default is 60.
KeepaliveInterval	Specify the TCP keepalive interval in seconds. The default is 10.

Name	Description
KeepaliveProbes	Specify the TCP keepalive probes. The default is 5.
PrivExecPasswordEnable	Enables authentication to access Privileged EXEC CLI command mode. Authentication is disabled by default.
RateLimitActionPollInterval	Specifies the interval, in seconds, at which this feature checks the hardware counters for packets dropped because they exceeded the configured port rate limit. The default is 3 seconds.

Configure Storm Control Actions

About This Task

Configure the action to take after Storm Control detects dropped packets on the port.

Procedure

1. On the Device Physical View, select a port or multiple ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Storm Control** tab.
5. In **ShutdownTimeout**, specify if the port remains permanently shutdown, or only for a specific time period.
6. In **TrapInterval**, specify the number of polling cycles for trap generation.
7. In **Threshold**, specify the packet threshold.
8. Select **Apply**.

Storm Control Field Descriptions

Use the data in the following table to use the **Storm Control** tab.

Name	Description
Index	Displays the port information.
ShutdownTimeout	Configures the port to shut down, and specifies how long, in seconds, the port remains shut down, after Storm Control detects dropped packets. A value of 0 indicates the port does not shutdown. A value of 1 indicates the port remains shut down until you manually re-enable it. The default is to keep the port up.
TrapInterval	Specifies the frequency, in polling cycles, of SNMP trap generation. Configure the value to 0 to disable traps and not generate log entries. The default is disabled.
Threshold	Specifies a threshold value for the delta of dropped packets per polling interval that triggers the shut down. The default is 1.



RADIUS

[RADIUS Fundamentals](#) on page 3047

[RADSec](#) on page 3055

[RADIUS configuration using CLI](#) on page 3057

[RADSec Configuration Using CLI](#) on page 3079

[RADIUS configuration using Enterprise Device Manager](#) on page 3084

[RADSec Configuration Using EDM](#) on page 3105

Table 240: RADIUS product support

Feature	Product	Release introduced
RADIUS (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 240: RADIUS product support (continued)

Feature	Product	Release introduced
RADIUS, community-based users (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
RFC 5176 – Dynamic Authorization Extensions to RADIUS	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
RFC 5997 – RADIUS Reachability Server Status	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 240: RADIUS product support (continued)

Feature	Product	Release introduced
RADIUS Configuration using Fully Qualified Domain Name (FQDN)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0.2
	5420 Series	Fabric Engine 9.0.2
	5520 Series	Fabric Engine 9.0.2
	5720 Series	Fabric Engine 9.0.2
	7520 Series	Fabric Engine 9.0.2
	7720 Series	Fabric Engine 9.0.2
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.0.2
	VSP 7400 Series	VOSS 9.0.2

RADIUS Fundamentals

Remote Access Dial-In User Services (RADIUS) is a distributed client/server system that assists in securing networks against unauthorized access, allowing a number of communication servers and clients to authenticate users' identity through a central database. The database within the RADIUS server stores information about clients, users, passwords, and access privileges including the use of shared secret.

RADIUS is a fully open and standard protocol, defined by two Requests for Comments (RFC) (Authentication: RFC2865, Accounting: RFC2866). You use RADIUS authentication to get secure access to the system (console/Telnet/SSH/EDM), and RADIUS accounting to track the management sessions (CLI only).

If you configure RADIUS using a fully qualified domain name (FQDN), the switch resolves the name to an IP address. You must configure a DNS server in the network to resolve the FQDN to an IP address.

RADIUS Server Support for IPv6

RADIUS supports both IPv4 and IPv6 with no differences in functionality or configuration in all but the following case. When you add or update a RADIUS server in Enterprise Device Manager (EDM) you must specify if the address type is an IPv4 or an IPv6 address.

How RADIUS Works

A RADIUS application has two components:

<ul style="list-style-type: none">• RADIUS server	A computer equipped with server software (for example, a UNIX workstation) that is located at a central office or campus. The server has authentication and access information in a form that is compatible with the client. Typically, the database in the RADIUS server stores client information, user information, password, and access privileges, including the use of a shared secret. A network can have one server for both authentication and accounting, or one server for each service.
<ul style="list-style-type: none">• RADIUS client	A device, router, or a remote access server, equipped with client software, that typically resides on the same local area network (LAN) segment as the server. The client is the network access point between the remote users and the server.

The two RADIUS processes are

- RADIUS authentication—Identifies remote users before you give them access to a central network site.
- RADIUS accounting—Performs data collection on the server during a remote user's dial-in session with the client.

Configuration of the RADIUS Server and Client

For more information about how to configure a RADIUS server, see the documentation that came with the server software.

The switch software supports BaySecure Access Control (BSAC) and the Merit Network servers. To use these servers, you must first obtain the software for the server you will use. Also, you must make changes to one or more configuration files for these servers.

RADIUS Authentication

You can use RADIUS authentication to use a remote server to authenticate logons. The RADIUS server also provides access authority. RADIUS assists network security and authorization by managing a database of users. The device uses this database to verify user names and passwords as well as information about the type of access priority available to the user.

When the RADIUS client sends an authentication request requesting additional information such as a SecurID number, it sends it as a challenge-response. Along with the challenge-response, it sends a reply-message attribute. The reply-message is a text string, such as `Please enter the next number on your SecurID card:.` The RFC defined maximum length of each reply-message attribute is 253 characters. If you have multiple instances of reply-message attributes that together form a large message that displays to the user, the maximum length is 2000 characters.

You can use additional user names to access the device, in addition to the six existing user names of ro, L1, L2, L3, rw, and rwa. The RADIUS server authenticates the user name and assigns one of the existing access priorities to that name. Unauthenticated user names are denied access to the device. You must add user names ro, L1, L2, L3, rw, and rwa to the RADIUS server if you enable authentication. Users not added to the server are denied access.

The limitation on the number of characters in a user name for users logging into CLI or EDM configured with RADIUS authentication is 64 characters.

**Note**

RADIUS server used-by snmp does not support authentication.

The following list shows the user configurable options of the RADIUS feature:

- up to 10 RADIUS servers in each device for fault tolerance (each server is assigned a priority and is contacted in that order).
- secret key for each server to authenticate the RADIUS client
- server UDP port
- maximum retries allowed
- time-out period for each attempt

**Note**

If you enable enhanced secure mode with the **boot config flags enhancedsecure-mode** command, you enable different access levels, along with stronger password complexity, length, and minimum change intervals. With enhanced secure mode enabled, the switch supports the following access levels for RADIUS authentication:

- Administrator
- Privilege
- Operator
- Auditor
- Security

The switch associates each user name with a certain role and appropriate authorization rights to view and configure commands. For more information on system access fundamentals and configuration, see [Managing the Switch](#) on page 246.

Silent Devices Surviving Reboot

When the switch reboots, it retains authenticated client information, including MAC and IP addresses. To maintain device authentication, the switch sends regular ARP and ping messages. PINGs dispatch if an IP address is configured on the MGMT VLAN, while ARPs are sent if an IP address is configured on the platform VLAN associated with the port.

Use of RADIUS to Modify User Access to CLI Commands

The switch provides CLI command access based on the configured access level of a user. However, you can use RADIUS to override CLI command access provided by the switch.

To override user access to CLI commands, you must configure the command-access-attribute on the switch and on the RADIUS server. (The switch uses decimal value 194 as the default for this parameter.) On the RADIUS server, you can then define the commands that the user can or cannot access.



Important

When you enable RADIUS on the switch and configure a RADIUS server to be used by CLI or EDM, the server authenticates the connection, whether it is FTP, HTTPs, SSH, or TELNET. However, in the event that the RADIUS server is unresponsive or is unreachable, the switch falls back to the local authentication, so that you can access the switch using your local login credentials.

Regardless of the RADIUS server configuration, you must configure the user's access on the switch based on the six platform access levels.

RADIUS Accounting

RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Session-IDs for each RADIUS account generate as 12-character strings. The first four characters in the string form a random number in hexadecimal format. The last eight characters in the string indicate the number of user sessions started since the last restart, in hexadecimal format.

The Network Address Server (NAS) IP address for a session is the address of the device interface to which the remote session is connected over the network. For a console session, modem session, and sessions running on debug ports, this value is set to 0.0.0.0, as is the case with RADIUS authentication.

The following table summarizes the events and associated accounting information logged at the RADIUS accounting server.

Table 241: Accounting Events and Logged Information

Event	Accounting information logged at server
Accounting is turned on at router	<ul style="list-style-type: none">Accounting on request: NAS IP address
Accounting is turned off at router	<ul style="list-style-type: none">Accounting off request: NAS IP address
User logs on	<ul style="list-style-type: none">Accounting start request: NAS IP addressSession IDUser name

Table 241: Accounting Events and Logged Information (continued)

Event	Accounting information logged at server
More than 40 CLI commands are executed	<ul style="list-style-type: none"> Accounting interim request: NAS IP address Session ID CLI commands User name
User logs off	<ul style="list-style-type: none"> Accounting stop request: NAS IP address Session ID Session duration User name Number of input octets for session Number of octets output for session Number of packets input for session Number of packets output for session CLI commands

When the device communicates with the RADIUS accounting server, the following actions occur:

1. If the server sends an invalid response, the response is silently discarded and the server does not make an attempt to resend the request.
2. User-specified number of attempts are made if the server does not respond within the user-configured timeout interval. If a server does not respond to any of the retries, requests are sent to the next priority server (if configured). You can configure up to 10 RADIUS servers for redundancy.

**Note**

RADIUS server used-by endpoint-tracking does not support accounting.

RADIUS Attribute 8 Framed-IP-Address

The switch uses RADIUS Attribute 8 Framed-IP-Address to export the IPv4 address of an EAP/NEAP client to the RADIUS server.

You must enable EAP and Node Alias globally and on an interface. For more information, see [Extensible Authentication Protocol over LAN](#) on page 1008 and [Node Alias](#) on page 713.

RFC 4675 RADIUS Attributes: Egress VLAN

Egress VLAN controls egress traffic. Egress VLAN supports two standard RADIUS attributes as defined in RFC 4675:

- Egress-VLANID
- Egress-VLAN-Name

RADIUS attributes control the 802.1Q tagging for traffic egressing a port where RADIUS authentication is performed for a connected EAP or NEAP client.

Egress VLANs are standard attributes, therefore the RADIUS server supports the attributes by default and offer the ability to configure the attributes. Each attribute has two parts:

1. Indicates if the frames on the VLAN egress must be tagged or untagged
2. Specifies the VLAN name or VLAN ID

The switch applies the VLAN received in the Egress-VLAN attributes to the port where the client is authenticated through RADIUS and then sets the tagging rules (tagged or untagged) accordingly.

The switch processes the Egress-VLAN attributes when decoding the RADIUS packet, therefore the switch adds the port to the VLANs first and then sets the proper tagging for the VLANs. You must create VLANs in advance on the switch.

In the MultiVlan operation mode, the EAP applies ingress hardware rules to ensure untagged traffic from each authenticated client goes into its own VLAN. The unauthenticated clients send traffic to the Guest VLAN, which matches the default VLAN ID.

For more information, see [VLAN RADIUS Attributes](#) on page 3133.

RADIUS Server Reachability

Configure up to 10 EAP RADIUS servers on the switch to manage fault tolerance. Each server is assigned a priority and is contacted in the priority order. If the first server is unavailable, the switch tries the second server, and so on, until the switch establishes a successful connection. Higher priority means lower integer value.

RADIUS server reachability prevents clients from trying to establish a connection with non reachable servers. RADIUS server reachability runs a periodic check in the background to identify the available servers. The switch is aware of the first available EAP RADIUS server without going through each of the servers and wait for time-outs.

Use RADIUS server reachability to configure the switch to use RADIUS requests to determine the reachability of the RADIUS server. The switch regularly performs the reachability test to determine if the switch should fail over to the secondary RADIUS server or activate the Fail Open VLAN, if configured on the switch.

Use one of the following modes to configure RADIUS reachability:

- **status-server mode:** Status-Server packets provide a standard-compliant alternative to configuring dummy RADIUS requests. You can configure the switch to send status-server packets when the keep-alive timer or the unreachable timer expires. In order to use status-server mode, the configured RADIUS servers must support RFC5997.
- **use-radius mode:** Configure user-radius mode if any of the RADIUS servers do not support RFC5997. In user-radius mode, the switch regularly generates a dummy RADIUS request with the user name `reachme` and password `reachme`. The switch interprets either Request Accept or Request Reject responses as a confirmation for server reachability, therefore it is not necessary to add the credentials on the server to test server reachability. You can configure the user name and password

for the dummy account through CLI. Use-radius is the default mode for RADIUS reachability.

You can configure the RADIUS reachability mode in either CLI or EDM.

**Note**

RADIUS server reachability is enabled on the switch and is not a configurable option. The reachability process starts when at least one RADIUS server used by EAP is configured, and RADIUS is enabled globally.

Based on the number of EAP RADIUS servers configured, the switch performs the following:

- If the highest priority EAP RADIUS server is reachable, the server status is updated to reachable and further authentication will use this server. As long as the highest priority EAP RADIUS server is reachable, the rest of the EAP RADIUS servers are not tested for reachability.
- If the highest priority EAP RADIUS server is not reachable, then the switch tests the rest of the EAP RADIUS servers for reachability. The servers are checked one by one for reachability based on their priority from highest to lowest. The first server that is reachable is used for authentication and the rest of the lower priority EAP RADIUS servers if any, are skipped from the reachability test.
- If all the EAP RADIUS servers are unreachable, then no further authentication occurs until the next successful reachability check.

The intervals between two consecutive reachability checks can be configured. The default values are as follows:

- one minute, if the last check result was unreachable
- three minutes, if the last check result was reachable

A server is marked as unreachable after a number of retries and time-outs.

RFC 3580 RADIUS Attributes: IEEE 802.1X Remote Authentication Dial In User Service

RFC 3580 provides support for EAP and NEAP clients for the following RADIUS attributes:

- Called-Station ID attribute: For IEEE 802.1X authenticators, the Called-Station ID stores the bridge or access point MAC address in upper case ASCII format, with octet values separated by a hyphen (-). For example: 00-10-A4-23-19-C0.

In IEEE 802.11, where the SSID is known, the SSID must be appended to the access point MAC address and separated from the MAC address with a colon (:). For example: 00-10-A4-23-19-C0:AP1.

- Calling-Station ID: For IEEE 802.1X authenticators, the Calling-Station ID is used to store the supplicant MAC address in upper case ASCII format, with octet values separated by a hyphen (-). For example: 00-10-A4-23-19-C0.

- **NAS-Port ID:** The NAS-Port ID is used to identify the IEEE 802.1X Authenticator port which authenticates the Supplicant. The NAS-Port-Id differs from the NAS-Port in that it is a string of variable length whereas the NAS-Port is a 4 octet value.

RFC 5176 — Dynamic Session Change

RFC 5176 allows you to dynamically change the following user session characteristics:

- You can disconnect an authenticated user on a port and remove all associated session context.

If the RADIUS server issues a disconnect command to the switch and the switch identifies a user (that satisfies all attributes of the RADIUS server request) on a port that has enabled RADIUS dynamic extensions commands, the switch performs the following actions:

- Notify the user of the disconnect by sending an 802.1x disconnect message to the client.
 - Remove all session context from the port.
 - Remove the port from the RADIUS-assigned VLAN, if applicable.
 - Send the disconnect response Disconnect-ACK to the RADIUS server if the user session is disconnected and all steps successfully performed.
 - Send the Disconnect-NAK response to the RADIUS server if the user session is not found or if the Network Access Server (NAS) cannot disconnect the session and discard the session context.
- You can use the Change of Authorization command to dynamically change the VLAN used by the RADIUS server.

If the RADIUS server issues a Change of Authorization command to the switch and the switch identifies a user (that satisfies all attributes of the RADIUS server request) on a port that has enabled RADIUS dynamic extensions commands, the switch performs the following actions:

- If the Change of Authorization command specifies a valid VLAN ID for a port, the port is removed from the VLAN specified by RADIUS and added to the VLAN specified in the request.
 - A CoA-ACK response is sent to the RADIUS server.
 - If the user session is not found or an error is encountered in processing the Change of Authorization command, then a CoA-NAK response is sent to the RADIUS server.
 - If the Change of Authorization request specifies a VLAN that is not port-based, a CoA-NAK response is sent to the RADIUS server.
- You can dynamically initiate client re-authentication.

Re-authenticate requests can be made with Change of Authorization or Disconnect packet IDs, but they must have the Re-authentication Request Vendor-Specific Attributes (VSA) set to True.

Dynamic session changes are directed to specific user sessions, as identified by RADIUS attributes.

To enable dynamic session changes, configure the following:

- You must enable EAP or Endpoint Tracking globally and at the port level.
- You must enable RADIUS dynamic extensions commands at the port level.

You can use the **show radius dynamic-server statistics** command to view statistics about dynamic session changes.

```
Switch:1#enable
Switch:1#show radius dynamic-server statistics

=====
                        RADIUS Dynamic Authorization Global Statistics
=====
Disconnects From Invalid Client Addresses:      0
CoAs From Invalid Client Addresses:              0
=====
```

RADSec

Table 242: RADSec product support

Feature	Product	Release introduced
RADIUS Security (RADSec)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2
	VSP 7400 Series	VOSS 8.2

Remote Access Dial-In User Services (RADIUS) Security (RADSec) provides secure communication between RADIUS peers using Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP), or Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP).

Certificates

RADSec peers use certificates to establish trust relationships. Certificates are specified in the RADSec profile on the switch, which can be the default profile or a profile that you configure.

- Manually install certificates, certificate chains, and private keys into designated folders for each RADIUS secure profile.

The client and secure server must use the same certificate authority (CA) certificate file, server certificate, certificate key file and password to establish a RADSec connection. The password is used for packet encryption and decryption.

Secure Mode

Configure the RADSec security mode to use the TLS protocol or DTLS protocol. Both TLS and DTLS modes support IPv4 addresses, but only TLS mode supports IPv6 addresses.



Important

To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the `radsecproxy.conf` configuration file:

```
tls default{
    ...
    TlsVersion TLS1_2
}
```

radsecproxy

RADSec uses radsecproxy to encrypt packets sent between RADIUS clients and a secure server. A radsecproxy process starts when RADIUS and secure-mode are enabled both globally and for the current server. Radsecproxy uses multiple instances on the switch, one for each security-configured RADIUS server. The maximum number of radsecproxy instances is 10, which is equivalent to the maximum number of RADIUS servers you can configure on the switch.



Note

As a best practice, use radsecproxy version 1.9.1 or later. Using an earlier version of radsecproxy can result in authentication failures.

The following events cause all radsecproxy processes to restart:

- A Segmented Management Instance interface is deleted, enabled, or disabled.
- An IPv4 address or route is changed.
- An IPv6 address or route is changed.
- RADIUS global enable and secure-enable change.

The following changes only affect a radsecproxy instance and result only in the associated process restarting:

- Enabling or disabling the RADIUS server state or secure state.
- Modifying the RADIUS server secure mode, associated profile, or secure log level.

If one of the global RADIUS flags (radius state, radius secure state) is disabled, all radsecproxy instances stop.

A RADSec profile configuration change causes all processes associated with the profile to restart.

A RADSec process automatically stops when the corresponding RADIUS server entry is deleted, or if RADIUS or secure mode is disabled for the server.

RADIUS configuration using CLI

You can configure Remote Access Dial-In User Services (RADIUS) to secure networks against unauthorized access, and allow communication servers and clients to authenticate users identity through a central database.

The database within the RADIUS server stores client information, user information, password, and access privileges, including the use of shared secret.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using CLI.

RADIUS is a fully open and standard protocol, defined by RFCs (Authentication: RFC2865, accounting RFC2866). With the switch, you use RADIUS authentication to secure access to the device (console/Telnet/SSH), and RADIUS accounting to track the management sessions for Command Line Interface (CLI) only.

RADIUS authentication allows the remote server to authenticate logons. RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Configure RADIUS Attributes

Configure RADIUS to authenticate user identity through a central database.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure RADIUS access priority:
`radius access-priority-attribute <192-240>`
3. Configure RADIUS accounting:
`radius accounting {attribute-value <192-240>|enable|include-cli-commands}`
4. Configure the RADIUS authentication info attribute value:
`radius auth-info-attr-value <0-255>`
5. Clear RADIUS statistics:
`radius clear-stat`
6. Configure the value of the CLI commands:
`radius cli-commands-attribute <192-240>`

7. Configure the value of the command access attribute:

```
radius command-access-attribute <192-240>
```

8. Configure the maximum number of servers allowed:

```
radius maxserver <1-10>
```

9. Configure the multicast address attribute:

```
radius mcast-addr-attr-value <0-255>
```

10. Enable RADSec globally:

```
radius secure-flag
```

11. Configure the RADSec profile:

```
radius secure-profile WORD<1-16> [ca-cert-file | cert-file | key-file  
| key-pwd]
```

Example

Configure RADIUS access priority:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#radius access-priority-attribute 192
```

Configure RADIUS accounting to include CLI commands:

```
Switch:1(config)#radius accounting include-cli-commands
```

Variable Definitions

The following table defines parameters for the **radius** command.

Variable	Value
<i>access-priority-attribute</i> <i><192-240></i>	Specifies the value of the access priority attribute. The default is 192.
<i>accounting {attribute-value</i> <i><192-240> enable </i> <i>include-cli-commands}</i>	Configures the accounting attribute value, enable accounting, or configure if accounting includes CLI commands. The default is false.
<i>auth-info-attr-value</i> <0-255>	Specifies the value of the authentication information attribute. The default is 91.
<i>clear-stat</i>	Clears RADIUS statistics.
<i>cli-cmd-count</i> <1-40>	Specifies how many CLI commands before the system sends a RADIUS accounting interim request. The default value is 40.
<i>cli-commands-attribute</i> <i><192-240></i>	Specifies the value of CLI commands attribute. The default is 195.
<i>cli-profile</i>	Enable RADIUS CLI profiling. CLI profiling grants or denies access to users being authenticated by way of the RADIUS server. You can add a set of CLI commands to the configuration on the RADIUS server, and you can specify the command-access more for these commands. The default is false.

Variable	Value
<i>command-access-attribute</i> <192-240>	Specifies the value of the command access attribute. The default is 194.
<i>enable</i>	Enable RADIUS authentication globally on the switch.
<i>maxserver</i> <1-10>	Specific to RADIUS authentication, configures the maximum number of servers allowed for the device. The default is 10.
<i>mcast-addr-attr-value</i> <0-255>	Specifies the value of the multicast address attribute. The default is 90.
<i>secure-flag</i>	Specifies whether RADIUS Security (RADSec) is globally enabled. The default is disabled.

Variable	Value
<i>secure-profile</i>	Specifies the RADSec profile name.
<pre>server host WORD<0-113> key WORD<0-32> [used-by {cli snmp web} [acct-enable] [acct-port <1-65536>] [enable] [port <1-65536>] [priority <1-10>] [retry <0-6> secure-enable secure-ocsp secure-log-level {critical debug error info warning} secure-mode {dtls tls} secure-profile WORD<1-16>] [timeout <1-60>]</pre>	<ul style="list-style-type: none"> • <i>host WORD<0-113></i> Creates a host server. <i>WORD<0-113></i> signifies an IP address or Fully Qualified Domain Name (FQDN). • <i>key WORD<0-32></i> Specifies a secret key in the range of 0-32 characters. • <i>used-by {cli eapol endpoint-tracking snmp web}</i> Specifies how the server functions. Configures the server for: <ul style="list-style-type: none"> ◦ cli authentication ◦ eapol authentication ◦ endpoint-tracking authentication ◦ snmp accounting ◦ web authentication • <i>acct-enable</i> Enables RADIUS accounting on this server. The system enables RADIUS accounting by default. • <i>acct-port <1-65536></i> Specifies a UDP port of the RADIUS accounting server. The default value is 1816. The UDP port value set for the client must match the UDP value set for the RADIUS server. • <i>enable</i> Enables the server. The default is true. • <i>port <1-65536></i> Specifies a UDP port of the RADIUS server. The default value is 1812. • <i>priority <1-10></i> Specifies the priority value for this server. The default is 10. • <i>retry <0-6></i> Specifies the maximum number of authentication retries. The default is 3. • <i>secure-enable</i> Enable secure mode on the server. • <i>secure-ocsp</i>

Variable	Value
	<p>Enable RADIUS secure Online Certificate Status Protocol (OCSP) checking for this server. The default is disabled.</p> <ul style="list-style-type: none"> <code>secure-log-level{critical debug error info warning}</code> <p>Specifies the RADIUS secure server log severity level.</p> <ul style="list-style-type: none"> <code>secure-mode{dtls tls}</code> <p>Specifies the protocol for establishing the secure connection with the server. IPv4 supports both dtls and tls modes. IPv6 only supports tls mode.</p> <ul style="list-style-type: none"> <code>secure-profileWORD<1-16></code> <p>Specifies the secure profile name.</p> <ul style="list-style-type: none"> <code>timeout <1-60></code> <p>Specifies the number of seconds before the authentication request times out. The default is 3.</p>

Configure RADIUS Profile

Use RADIUS CLI profiling to grant or deny CLI command access to users being authenticated by way of the RADIUS server. You can add a set of CLI commands to the configuration file on the radius server, and you can specify the command-access mode for these commands. The default is false.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable RADIUS CLI profiling:

```
radius cli-profile
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# radius cli-profile
```

Enable RADIUS Authentication

About This Task

Enable or disable RADIUS authentication globally on the device to allow further configuration to take place. Use the no option to disable RADIUS authentication globally. The default is false or disabled.

Procedure

- 1. Enter Global Configuration mode:
`enable`
`configure terminal`
- 2. Enable RADIUS authentication globally on the switch:
`radius enable`

Enable RADIUS Accounting

Before You Begin

- You must configure a RADIUS server before you can enable RADIUS accounting.

About This Task

Enable Remote Access Dial-in User Services (RADIUS) accounting to log all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Procedure

- 1. Enter Global Configuration mode:
`enable`
`configure terminal`
- 2. Enable RADIUS accounting globally:
`radius accounting enable`
- 3. Include or exclude CLI commands in RADIUS accounting updates:
`radius accounting include-cli-commands`
- 4. Specify the integer value of the CLI commands attribute:
`radius accounting attribute-value <192-240>`

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# radius accounting enable
Switch:1(config)# radius accounting include-cli-commands
```

Variable Definitions

The following table defines parameters for the **radius accounting** command.

Variable	Value
<i>enable</i>	Enable RADIUS globally.
<i>include-cli-commands</i>	Include CLI commands in RADIUS accounting updates.
<i>attribute-value</i> <i><192-240></i>	Specify the integer value of the CLI commands attribute.

Enable RADIUS-SNMP Accounting

Before You Begin

- You must configure a RADIUS server before you can enable RADIUS-SNMP accounting.

About This Task

Enable Remote Access Dial-in User Services (RADIUS) Simple Network Managing Protocol (SNMP) accounting globally. Use SNMP to remotely collect management data. An SNMP agent is a software process that monitors the UDP port 161 for SNMP messages. Each SNMP message sent to the agent contains a list of management objects.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Enable RADIUS Simple Network Management Protocol (SNMP) accounting globally:

```
radius-snmp acct-enable
```
- Set a timer to send a stop accounting message for RADIUS Simple Network Management Protocol (SNMP):

```
radius-snmp abort-session-timer <30-65535>
```
- Set the timer for re-authentication of the SNMP session:

```
radius-snmp re-auth-timer <30-65535>
```
- Specify the user name for SNMP access:

```
radius-snmp user WORD <0-20>
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# radius-snmp acct-enable
Switch:1(config)# radius-snmp abort-session-timer 30
```

Variable Definitions

The following table defines parameters for the **radius-snmp** command.

Variable	Value
<i>acct-enable</i>	Enables RADIUS accounting globally. You cannot enable RADIUS accounting before you configure a valid server. The system disables RADIUS accounting by default. The default is false. Use the no option to disable RADIUS accounting globally: no radius-snmp acct-enable
<i>abort-session-timer</i> <i><30-65535></i>	Set the timer, in seconds, to send a stop accounting message. The default is 180.

Variable	Value
<code>re-auth-timer <30-65535></code>	Sets timer for re-authentication of the SNMP session. The timer value ranges from 30 to 65535 seconds. The default is 180.
<code>user WORD <0-20></code>	Specifies the user name for SNMP access. WORD <0-20> specifies the user name in a range of 0 to 20 characters. The default is <code>snmp_user</code> .

Configure RADIUS Accounting Interim Request

About This Task

Configure RADIUS accounting interim requests to create a log whenever a user executes more than the number of CLI commands you specify.

If the packet size equals or exceeds 1.8 KB, an interim request packet is sent even if the configured limit is not reached. Therefore, the trigger to send out the interim request is either the configured value or a packet size greater than, or equal to 1.8 KB, whichever happens first.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure RADIUS accounting interim requests:
`radius cli-cmd-count <1-40>`
3. Include or exclude CLI commands in RADIUS accounting:
`radius accounting include-cli-commands`



Important

You must configure the **`radius accounting include-cli-commands`** command for accounting interim requests to function.

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# radius cli-cmd-count 30
Switch:1(config)# radius accounting include-cli-commands
```

Variable Definitions

The following table defines parameters for the **radius cli-cmd-count** command.

Variable	Value
<1-40>	Specifies how many CLI commands, from 1 to 40, before the system sends a RADIUS accounting interim request. The default value is 40.

Configure RADIUS Authentication and RADIUS Accounting Attributes

About This Task

Configure RADIUS authentication and RADIUS accounting attributes to determine the size of the packets received.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the RADIUS authentication attribute value:

```
radius command-access-attribute <192-240>
```
3. Configure the RADIUS accounting attribute value:

```
radius accounting attribute-value <192-240>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#radius command-access-attribute 192
Switch:1(config)#radius accounting attribute-value 192
```

Variable Definitions

The following table defines parameters for the **radius** command.

Variable	Value
<i>access-priority-attribute</i> <192-240>	Specifies the value of the access priority attribute. The default is 192.
<i>accounting {attribute-value <192-240> enable include-cli-commands}</i>	Configures the accounting attribute value, enable accounting, or configure if accounting includes CLI commands. The default is false.
<i>auth-info-attr-value</i> <0-255>	Specifies the value of the authentication information attribute. The default is 91.
<i>clear-stat</i>	Clears RADIUS statistics.
<i>cli-cmd-count</i> <1-40>	Specifies how many CLI commands before the system sends a RADIUS accounting interim request. The default value is 40.

Variable	Value
<i>cli-commands-attribute</i> <192-240>	Specifies the value of CLI commands attribute. The default is 195.
<i>cli-profile</i>	Enable RADIUS CLI profiling. CLI profiling grants or denies access to users being authenticated by way of the RADIUS server. You can add a set of CLI commands to the configuration on the RADIUS server, and you can specify the command-access more for these commands. The default is false.
<i>command-access-attribute</i> <192-240>	Specifies the value of the command access attribute. The default is 194.
<i>enable</i>	Enable RADIUS authentication globally on the switch.
<i>maxserver</i> <1-10>	Specific to RADIUS authentication, configures the maximum number of servers allowed for the device. The default is 10.
<i>mcast-addr-attr-value</i> <0-255>	Specifies the value of the multicast address attribute. The default is 90.
<i>secure-flag</i>	Specifies whether RADIUS Security (RADSec) is globally enabled. The default is disabled.

Variable	Value
<i>secure-profile</i>	Specifies the RADSec profile name.
<pre>server host WORD<0-113> key WORD<0-32> [used-by {cli snmp web} [acct-enable] [acct-port <1-65536>] [enable] [port <1-65536>] [priority <1-10>] [retry <0- 6> secure-enable secure-ocsp secure-log-level {critical debug error info warning} secure-mode {dtls tls} secure-profile WORD<1-16>] [timeout <1- 60>]</pre>	<ul style="list-style-type: none"> • <i>host WORD<0-113></i> Creates a host server. <i>WORD<0-113></i> signifies an IP address or Fully Qualified Domain Name (FQDN). • <i>key WORD<0-32></i> Specifies a secret key in the range of 0-32 characters. • <i>used-by {cli eapol endpoint-tracking snmp web}</i> Specifies how the server functions. Configures the server for: <ul style="list-style-type: none"> ◦ cli authentication ◦ eapol authentication ◦ endpoint-tracking authentication ◦ snmp accounting ◦ web authentication • <i>acct-enable</i> Enables RADIUS accounting on this server. The system enables RADIUS accounting by default. • <i>acct-port <1-65536></i> Specifies a UDP port of the RADIUS accounting server. The default value is 1816. The UDP port value set for the client must match the UDP value set for the RADIUS server. • <i>enable</i> Enables the server. The default is true. • <i>port <1-65536></i> Specifies a UDP port of the RADIUS server. The default value is 1812. • <i>priority <1-10></i> Specifies the priority value for this server. The default is 10. • <i>retry <0-6></i> Specifies the maximum number of authentication retries. The default is 3. • <i>secure-enable</i> Enable secure mode on the server. • <i>secure-ocsp</i>

Variable	Value
	<p>Enable RADIUS secure Online Certificate Status Protocol (OCSP) checking for this server. The default is disabled.</p> <ul style="list-style-type: none"> <code>secure-log-level{critical debug error info warning}</code> <p>Specifies the RADIUS secure server log severity level.</p> <ul style="list-style-type: none"> <code>secure-mode{dtls tls}</code> <p>Specifies the protocol for establishing the secure connection with the server. IPv4 supports both dtls and tls modes. IPv6 only supports tls mode.</p> <ul style="list-style-type: none"> <code>secure-profileWORD<1-16></code> <p>Specifies the secure profile name.</p> <ul style="list-style-type: none"> <code>timeout <1-60></code> <p>Specifies the number of seconds before the authentication request times out. The default is 3.</p>

Add a RADIUS Server

About This Task

Add a RADIUS server to provide RADIUS service on the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Add a RADIUS server:

```
radius server host WORD<0-113> key WORD<0-32> [acct-enable | acct-port
<1-65536> | enable | port <1-65536> | priority <1-10> | retry <0-6>
| secure-enable |secure-log-level <critical | debug | error | info
| warning> | secure-mode <dtls | tls> |secure-profile WORD<1-16> |
timeout <1-180> | used-by {cli|eapol|endpoint-tracking|snmp|web}]
```

Examples

Add a RADIUS server using an IPv6 address:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#radius server host 4717:0000:0000:0000:0000:0000:7933:0001 key testkey1
used-by snmp port 12 retry 5 timeout 10 enable
```

Add a RADIUS server using a Fully Qualified Domain Name (FQDN):

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#radius server host host.example.com key testkey1 used-by snmp port 12
retry 5 timeout 10 enable
```

Variable Definitions

The following table defines parameters for the **radius server host** command.

Variable	Value
<i>host</i> <i>WORD</i> <0-113>	Configures a host server. <i>WORD</i> <0-113> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.
<i>acct-enable</i>	Enables RADIUS accounting on this server. The system enables RADIUS accounting by default.
<i>acct-port</i> <1-65536>	Configures the UDP port of the RADIUS accounting server. The default value is 1813. Important: The UDP port value set for the client must match the UDP value set for the RADIUS server.
<i>enable</i>	Enables the RADIUS server. The default is true.
<i>key</i> <i>WORD</i> <0-32>	Configures the secret key of the authentication client.
<i>port</i> <1-65536>	Configures the UDP port of the RADIUS authentication server. The default value is 1812.
<i>priority</i> <1-10>	Configures the priority value for this server. The default is 10.
<i>retry</i> <0-6>	Configures the number of authentication retries the server will accept. The default is 1.
<i>secure-enable</i>	Enable RADIUS Security (RADSec).
<i>secure-log-level</i> < <i>critical</i> <i>debug</i> <i>error</i> <i>info</i> <i>warning</i> >	Specifies the log severity level. The default is error.

Variable	Value
<code>secure-mode <dtls tls></code>	<p>Specifies the protocol for establishing the secure connection with the server. The possible values are:</p> <ul style="list-style-type: none"> • dtls - Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP) • tls - Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP) <p>The default is TLS.</p> <p>Important: To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the <code>radsecproxy.conf</code> configuration file:</p> <pre> tls default{ ... TlsVersion TLS1_2 } </pre>
<code>secure-ocsp</code>	<p>Enable RADIUS Online Certificate Status Protocol (OCSP) checking.</p> <p>The default is disabled.</p>
<code>secure-profile WORD<1-16></code>	<p>Configures the secure profile for the server.</p>
<code>timeout <1-180></code>	<p>Configures the number of seconds before the authentication request times out. The default is 8.</p>
<code>used-by {cli eapol endpoint-tracking snmp web}</code>	<p>Configures how the server functions:</p> <ul style="list-style-type: none"> • cli—configure the server for CLI authentication. • eapol—configure the server for EAPoL authentication. • endpoint-tracking—configure the server for Endpoint Tracking authentication. • snmp—configure the server for SNMP accounting. • web—configure the server for HTTP(s) authentication. <p>The default is cli.</p>

Modify RADIUS Server Settings

About This Task

Change a specified RADIUS server value without having to delete the server and recreate it again.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Modify a RADIUS server:

```
radius server host WORD <0-255> used-by {cli|eapol|endpoint-tracking|
snmp|web} [acct-enable | acct-port <1-65536> | enable | key WORD<0-32>
| port <1-65536> | priority <1-10> | retry <0-6> | secure-enable |
secure-log-level <critical | debug | error | info | warning> | secure-
mode <dtls | tls> |secure-profile WORD<1-16> | timeout <1-180>]
```

Example

Modify a RADIUS server:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#radius server host 4717:0000:0000:0000:0000:7933:0001 used-by snmp
port 12 retry 5 timeout 10 enable
```

Variable Definitions

The following table defines parameters for the **radius server host** command.

Variable	Value
<i>host WORD<0-113></i>	Configures a host server. <i>WORD<0-113></i> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.
<i>acct-enable</i>	Enables RADIUS accounting on this server. The system enables RADIUS accounting by default.
<i>acct-port <1-65536></i>	Configures the UDP port of the RADIUS accounting server. The default value is 1813. Important: The UDP port value set for the client must match the UDP value set for the RADIUS server.
<i>enable</i>	Enables the RADIUS server. The default is true.
<i>key WORD<0-32></i>	Configures the secret key of the authentication client.
<i>port <1-65536></i>	Configures the UDP port of the RADIUS authentication server. The default value is 1812.
<i>priority <1-10></i>	Configures the priority value for this server. The default is 10.
<i>retry <0-6></i>	Configures the number of authentication retries the server will accept. The default is 1.
<i>secure-enable</i>	Enable RADIUS Security (RADSec).

Variable	Value
<code>secure-log-level</code> <critical debug error info warning>	Specifies the log severity level. The default is error.
<code>secure-mode</code> <dtls tls>	<p>Specifies the protocol for establishing the secure connection with the server. The possible values are:</p> <ul style="list-style-type: none"> • dtls - Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP) • tls - Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP) <p>The default is TLS.</p> <p>Important: To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the <code>radsecproxy.conf</code> configuration file:</p> <pre>tls default{ .. TlsVersion TLS1_2 }</pre>
<code>secure-ocsp</code>	<p>Enable RADIUS Online Certificate Status Protocol (OCSP) checking.</p> <p>The default is disabled.</p>
<code>secure-profile</code> WORD<1-16>	Configures the secure profile for the server.
<code>timeout</code> <1-180>	Configures the number of seconds before the authentication request times out. The default is 8.
<code>used-by</code> {cli eapol endpoint-tracking snmp web}	<p>Configures how the server functions:</p> <ul style="list-style-type: none"> • cli—configure the server for CLI authentication. • eapol—configure the server for EAPoL authentication. • endpoint-tracking—configure the server for Endpoint Tracking authentication. • snmp—configure the server for SNMP accounting. • web—configure the server for HTTP(s) authentication. <p>The default is cli.</p>

View RADIUS Information

View the global status of RADIUS information to ensure you configured the RADIUS feature according to the needs of the network.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. View the global status of RADIUS information:

```
show radius
```

Example

```
Switch:1>show radius
      acct-attribute-value : 193
      acct-enable         : false
acct-include-cli-commands : false
access-priority-attribute : 192
      auth-info-attr-value : 91
command-access-attribute  : 194
      cli-commands-attribute : 195
      cli-cmd-count        : 40
      cli-profile-enable    : false
      enable               : false
      igap-passwd-attr     : standard
igap-timeout-log-fsize    : 512
      maxserver            : 10
mcast-addr-attr-value     : 90
supported-vendor-ids      : 1584, 562, 1916
      secure-flag          : false
```

View RADIUS Server Information

If your system is configured with a RADIUS server you can display the RADIUS server information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. To view the RADIUS server information enter the following command:

```
show radius-server
```



Note

If no RADIUS server is configured, the system displays the following message:

```
no RADIUS server configured
```

Example

```
Switch:1>show radius-server
=====
                        Radius Server Entries
=====
```

ACCT	NAME	USEDDBY	SECRET	PORT	PRIOR	RETRY	TIMEOUT	ENABLED	ENABLED	MODE	INSTANCE	PROFILE	LOG-LEVEL	OCSP	PORT
192.0.2.14	cli	*****	1812	10	1	8	true	true	tls	0	radsecp	error	false	1813	
192.0.2.15	cli	*****	1812	10	1	8	true	false	tls	-1	default	error	false	1813	

View RADIUS Server Statistics

About This Task

You cannot collect the following network statistics from a console port: the number of input and output packets, and the number of input and output bytes. All other statistics from console ports are available to assist with debugging.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display RADIUS server statistics:
`show radius-server statistics`
3. Clear server statistics:
`clear radius statistics`

Example

```
Switch:1#show radius-server statistics
```

```
Responses with invalid server address: 0
```

```
Radius Server(UsedBy) : 192.0.2.58 (cli)
```

```
-----  
Access Requests : 52  
Access Accepts : 0  
Access Rejects : 0  
Bad Responses : 52  
Client Retries : 52  
Pending Requests : 0  
Acct On Requests : 1  
Acct Off Requests : 0  
Acct Start Requests : 47  
Acct Stop Requests : 46  
Acct Interim Requests : 0  
Acct Bad Responses : 94  
Acct Pending Requests : 0  
Acct Client Retries : 94  
Access Challenges : 0  
Round-trip Time :  
Nas Ip Address : 192.0.2.32
```

```
Radius Server(UsedBy) : 192.0.2.58 (snmp)
```

```
-----  
Access Requests : 0  
Access Accepts : 0  
Access Rejects : 0  
Bad Responses : 0  
Client Retries : 0  
Pending Requests : 0  
Acct On Requests : 0  
Acct Off Requests : 0  
Acct Start Requests : 0  
Acct Stop Requests : 0  
Acct Interim Requests : 0  
Acct Bad Responses : 0  
Acct Pending Requests : 0  
Acct Client Retries : 0  
Access Challenges : 0  
Round-trip Time :
```



```
Nas Ip Address : 192.0.2.32
--More-- (q = quit)
```

Configure RADIUS Server Reachability

About This Task

Use this procedure to configure the RADIUS server reachability settings.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the RADIUS keep-alive timer:
`radius reachability keep-alive-timer <30-600>`
3. Configure the RADIUS reachability mode:
`radius reachability mode <status-server | use-radius>`
4. If you selected `use-radius` as the RADIUS reachability mode, configure the RADIUS request user name and password:
`radius reachability username WORD<1-16> password WORD<1-16>`
5. Configure the RADIUS reachability unreachable timer:
`radius reachability unreachable-timer <30-600>`

Example

Configure values for RADIUS server reachability:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

Configure the RADIUS keep-alive timer:

```
Switch:1(config)#radius reachability keep-alive-timer 30
```

Configure the RADIUS reachability mode:

```
Switch:1(config)#radius reachability mode status-server
```

Configure the RADIUS reachability unreachable timer:

```
Switch:1(config)#radius reachability unreachable-timer 30
```

Variable Definitions

The following table defines parameters for the **radius reachability** command.

Variable	Value
keep-alive-timer <30-600>	Specifies, in seconds, the interval between checks when radius server is reachable. The default is 180 seconds.
mode <status-server use-radius>	Specifies status-server mode or use-radius mode. Status-server mode provides a standard-compliant method for RADIUS reachability. Use-radius mode requires the configuration of dummy packets that are sent to RADIUS servers. The default is use-radius mode.
password <i>WORD</i> <1-16>	Configures the RADIUS request password. The default is extremenetworks.
unreachable-timer <30-600>	Specifies, in seconds, the interval between checks when radius server is unreachable. The default is 60 seconds.
username <i>WORD</i> <1-16>	Configures the RADIUS request user name. The default is extremenetworks.

Display RADIUS Server Reachability

About This Task

Use this procedure to display the RADIUS server reachability configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the RADIUS server reachability configuration:

```
show radius reachability
```

Example

Display the RADIUS server reachability configuration.

```
Switch:1#show radius reachability
EAP RADIUS reachability mode      : use-radius
EAP RADIUS reachability status    : RADIUS is disabled globally
EAP RADIUS reachable server       : none
Time until next check             : N/A
RADIUS username                   : reachme
RADIUS password                   : reachme
RADIUS keep-alive-timer           : 180
RADIUS unreachable-timer          : 60
```

Display RADIUS SNMP Configurations

Display current RADIUS SNMP configurations.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display the current RADIUS server SNMP configurations:

```
show radius snmp
```

Example

```
Switch:1>show radius snmp
      abort-session-timer : 180
      acct-enable : false
      user : snmp_user
      enable : false
      re-auth-timer : 180
```

Display RADIUS Dynamic Server Information

About This Task

Use the following procedure to display information about the configuration of RADIUS dynamic session changes.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about dynamic change configuration:

```
show radius dynamic-server [client | statistics]
```

Examples

```
Switch:1>show radius dynamic-server
```

```
=====
                        RADIUS Dynamic Authorization General Info
=====
CLIENT      RESOLVED      UDP      CLIENT      SECRET
ADDRESS      ADDRESS      PORT      ENABLED      KEY
-----
192.0.2.15           1026      Disabled
192.0.2.16           1027      Disabled
host.example.com     3799      Enabled
=====
```

All 2 out of 2 Total Num of RADIUS Dynamic Authorization clients displayed

Display dynamic server statistics:

```
Switch:1>show radius dynamic-server statistics
```

```
=====
                        RADIUS Dynamic Authorization Global Statistics
=====
Disconnects From Invalid Client Addresses:      0
CoAs From Invalid Client Addresses:              0
=====
```

Variable Definitions

The following table defines parameters for the **show radius dynamic-server** command.

Variable	Value
client	Display RADIUS Dynamic Authorization configuration values for the specified client.
statistics	Display statistics for RADIUS Dynamic Authorization clients.

Configure a RADIUS Dynamic Client

Before You Begin

You must enable EAPOL globally and at the port level.

About This Task

Configure a client to enable processing of dynamic session changes. You can configure up to eight clients.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure client values:
`radius dynamic-server client WORD<0-255> [port <1024-65535>] [secret WORD<0-32>] enable`

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#radius dynamic-server client 192.0.2.15 port 1025 secret 123 enable
```

Variable Definitions

The following table defines parameters for the **radius dynamic-server** command.

Variable	Value
client WORD<0-255>	Specifies the address or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.
enable	Enable processing of dynamic session changes.
port <1024-65535>	Specifies the port value.
secret WORD<0-32>	Specifies a value for secret key.

RADSec Configuration Using CLI

This section contains procedures for configuring RADSec using Command Line Interface (CLI).

Configure a RADIUS Secure Profile using Manual Certificates

Configure RADIUS secure profiles with certificate information, certificate key information, and password information, which enables RADSec peers to establish connections.

About This Task

A new profile directory is created for each new profile in the `/intflash/.radsec/profile/radsec` directory. Profile configuration file "profile_info.cfg" is available in `/intflash/.radsec/profile` directory.



Note

All the files (certificates and keys) must be in .pem format and you must copy them to the `/intflash` directory.

You can configure a maximum of 10 RADSec profiles.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure a name for the RADIUS profile:
`radius secure-profile WORD<1-16>`
3. Configure the full file path of the certificate authority (CA) certificate for the RADIUS secure profile:
`radius secure-profile WORD<1-16> ca-cert-file WORD<0-128>`
`[WORD<1-128>]`
4. Configure the full path of the server certificate for the RADIUS secure profile:
`radius secure-profile WORD<1-16> cert-file WORD<0-128> [WORD<1-128>]`
5. Configure the full path of the private key file for the RADIUS secure profile:
`radius secure-profile WORD<1-16> key-file WORD<0-128> [WORD<1-128>]`
6. Configure the private key password for the RADIUS secure profile:
`radius secure-profile WORD<1-16> key-pwd WORD<0-255>`

Variable Definitions

The following table defines parameters for the **radius secure-profile** command.

Variable	Value
WORD<1-16>	Specifies the RADIUS secure profile name.
CA-cert-file WORD<0-128>	Specifies the full file path of the certificate authority (CA) certificate.
cert-file WORD<0-128>	Specifies the full file path of the server certificate.
key-file WORD<0-128>	Specifies the full file path of the private key file.
WORD<1-128>	Specifies the file name.
key-pwd WORD<0-255>	Specifies the private key password.

Display RADIUS Secure Profile information

About This Task

Use the following procedure to display information about the configuration of RADIUS secure profiles.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display information about RADIUS secure profile configuration:
`show radius secure-profile`

Example

```
Switch:1>show radius secure-profile
=====
                        Secure RADIUS profile
=====
Profile default:
  Name      : default
  RootCert  : n/a
  Cert      : n/a
  Key       : n/a
  Password  : *****
```

Associate a RADSec Profile with a RADIUS Server

Before You Begin

Add a RADIUS Server. For more information, see [Add a RADIUS Server](#) on page 3068.

About This Task

RADSec peers use certificates to establish trust relationships. You must associate a RADSec profile with a RADIUS server in order to use certificates.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Associate a RADSec profile with a RADIUS server:

```
radius server host WORD<0-113> used-by {cli | eapol | endpoint-  
tracking | snmp | web} secure-profile WORD<1-16>
```

Variable Definitions

The following table defines parameters for the **radius server host** command.

Variable	Value
<i>WORD<0-255></i>	Configures a host server. <i>WORD<0-255></i> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.
<i>used-by</i>	Specifies how the server functions. Configures the server for one of the following: <ul style="list-style-type: none">• cli authentication• eapol authentication• endpoint-tracking authentication• snmp accounting• web authentication
<i>secure-enable</i>	Enables RADSec on the RADIUS server.

Configure RADSec Secure Mode

Configure the secure mode for RADSec as either Transport Layer Security (TLS) protocol or Datagram Transport Layer Security (DTLS) protocol.

Before You Begin

To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the `radsecproxy.conf` configuration file:

```
tls default{  
    ...  
    TlsVersion TLS1_2  
}
```

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the secure mode:

```
radius server host WORD<0-113> used-by {cli | eapol | endpoint-
tracking | snmp | web} secure-mode {tls | dtls}
```

Variable Definitions

The following table defines parameters for the **radius server host** command.

Variable	Value
<i>WORD<0-113></i>	Configures a host server. <i>WORD<0-113></i> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.
<i>used-by</i>	Specifies how the server functions. Configures the server for one of the following: <ul style="list-style-type: none"> cli authentication eapol authentication endpoint-tracking authentication snmp accounting web authentication
<i>secure-mode</i>	Specifies the RADSec security mode. Possible values are: <ul style="list-style-type: none"> tls - Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP) dtls - Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP) The default is tls.

Enable RADSec

About This Task

You must enable RADSec globally and at the RADIUS server level in order for RADSec to be enabled for a RADIUS server.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable RADSec globally:

```
radius secure-flag
```

3. Enable RADSec on the RADIUS server:

```
radius server host WORD<0-113> used-by {cli | eapol | endpoint-
tracking | snmp | web} secure-enable
```


Change the Logging Level for RADSec

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Change the logging level:

```
radius server host WORD<0-113> used-by {cli|eapol|endpoint-tracking|snmp|web} secure-log-level <critical | debug | error | info | warning>
```

Variable Definitions

The following table defines parameters for the **radius server host WORD<0-255> used-by {cli|eapol|endpoint-tracking|snmp|web} secure-log-level** command.

Variable	Value
<code><critical debug error info warning></code>	Specifies the RADIUS secure server log severity level. The default is error.
<code>used-by {cli eapol endpoint-tracking snmp web}</code>	Configures how the server functions: <ul style="list-style-type: none"> • cli - configures the server for CLI authentication. • eapol - configures the server for Extensible Authentication Protocol over LAN (EAPoL) authentication • endpoint-tracking - configures the server for Endpoint Tracking authentication • snmp - configures the server for Simple Network Management Protocol (SNMP) accounting • web - configures the server for web authentication
<code>WORD<0-113></code>	Configures a host server. <code>WORD<0-113></code> specifies the IPv4 address, IPv6 address, or fully qualified domain name (FQDN). If you use an FQDN, you must also configure the switch to use DNS.

Enable OCSP Checking

About This Task

Perform the following procedure to enable RADIUS Online Certificate Status Protocol (OCSP) checking for this server.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable RADIUS OCSP checking:
`radius server host WORD<0-113> used-by cli secure-ocsp`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#radius server host 192.0.2.24 used-by cli secure-ocsp
```

Display Log Information for Secure RADIUS Instance

About This Task

Use the following procedure to display log information for secure instance of RADIUS.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display log information for a secure instance of RADIUS:
`show radius secure-server log <0-9>.`

Variable Definitions

The following table defines parameters for the **show radius secure-server** command.

Variable	Value
<code>log<0-9></code>	Specifies the log for a secure RADIUS instance.

RADIUS configuration using Enterprise Device Manager

You can configure Remote Access Dial-In User Services (RADIUS) to assist in securing networks against unauthorized access, and allow communication servers and clients to authenticate the identity of users through a central database.

The database within the RADIUS server stores client information, user information, password, and access privileges, including the use of shared secret.

RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all but the following case. When adding a RADIUS server in Enterprise Device Manager (EDM) or modifying a RADIUS configuration in EDM, you must specify if the address type is an IPv4 or an IPv6 address.

RADIUS is a fully open and standard protocol, defined by RFCs (Authentication: RFC2865, accounting RFC2866). With the switch, you use RADIUS authentication to

secure access to the device (console/Telnet/SSH), and RADIUS accounting to track the management sessions for Command Line Interface (CLI) only.

RADIUS authentication allows the remote server to authenticate logons. RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Enable RADIUS Authentication

About This Task

Enable RADIUS authentication globally to enable all features and functions of RADIUS to operate with the RADIUS server.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Global** tab.
4. Select **Enable**.
5. In **MaxNumberServer**, type a value for the maximum number of servers.
6. In **AccessPriorityAttrValue**, type an access policy value (by default, this value is 192).
7. Configure the rest of the parameters in the RADIUS global tab.
8. Select **Apply**.

RADIUS Global Field Descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The default is 192.
AcctEnable	Enables RADIUS accounting.
AcctAttriValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.

Name	Description
McastAttributeValue	Specifies the value of the Mcast attribute. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
CliCmdCount	Gives the value for the CLI command count. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.
SupportedVendorIds	Shows the vendor IDs that the software supports for RADIUS.
RadiusReachability	Specifies the mode for RADIUS reachability. Status-server mode provides a standard-compliant method for RADIUS reachability. Use-radius mode requires the configuration of dummy packets that are sent to RADIUS servers. The default is use-radius mode.
SecureEnable	Enable RADIUS Security (RADSec).
UserName	Specifies the user name for RADIUS server reachability. The default is extremenetworks.
Password	Specifies the password for RADIUS server reachability. The default is extremenetworks.
Confirm Password	Confirms the password for RADIUS server reachability.
Unreachable Timer	Specifies, in seconds, the interval between checks when radius server is unreachable. The default is 60 seconds.
Keep Alive Timer	Specifies, in seconds, the interval between checks when radius server is reachable. The default is 180 seconds.

Enable RADIUS Accounting

Before You Begin

- You must set up a RADIUS server and add it to the configuration file of the device before you can enable RADIUS accounting on the device. Otherwise, the system displays an error message.

About This Task

Enable RADIUS accounting to log all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Global** tab.
4. Select **AcctEnable**.
5. In **AcctAttrValue**, type an access policy value (by default, this value is 193).
6. Select **Apply**.

RADIUS Global Field Descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The default is 192.
AcctEnable	Enables RADIUS accounting.
AcctAttrValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
CliCmdCount	Gives the value for the CLI command count. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.
SupportedVendorIds	Shows the vendor IDs that the software supports for RADIUS.

Name	Description
RadiusReachability	Specifies the mode for RADIUS reachability. Status-server mode provides a standard-compliant method for RADIUS reachability. Use-radius mode requires the configuration of dummy packets that are sent to RADIUS servers. The default is use-radius mode.
SecureEnable	Enable RADIUS Security (RADSec).
UserName	Specifies the user name for RADIUS server reachability. The default is extremenetworks.
Password	Specifies the password for RADIUS server reachability. The default is extremenetworks.
Confirm Password	Confirms the password for RADIUS server reachability.
Unreachable Timer	Specifies, in seconds, the interval between checks when radius server is unreachable. The default is 60 seconds.
Keep Alive Timer	Specifies, in seconds, the interval between checks when radius server is reachable. The default is 180 seconds.

Disable RADIUS Accounting

Before You Begin

- You cannot globally disable RADIUS accounting unless a server entry exists.

About This Task

Disabling RADIUS accounting removes the accounting function from the RADIUS server.

Procedure

- In the navigation pane, expand **Configuration > Security > Control Path**.
- Select **RADIUS**.
- Select the **RADIUS Global** tab.
- Disable RADIUS accounting by clearing **AcctEnable**.
- Select **Apply**.

Enable RADIUS Accounting Interim Request

About This Task

Enable the RADIUS accounting interim request feature to create a log whenever more than the specified number of CLI commands are executed.

Procedure

- In the navigation pane, expand **Configuration > Security > Control Path**.
- Select **RADIUS**.

3. Select the **RADIUS Global** tab.
4. Type the number of CLI commands in **CliCmdCount**.
5. Select **Apply**.

RADIUS Global Field Descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The default is 192.
AcctEnable	Enables RADIUS accounting.
AcctAttriValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
CliCmdCount	Gives the value for the CLI command count. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.
SupportedVendorIds	Shows the vendor IDs that the software supports for RADIUS.
RadiusReachability	Specifies the mode for RADIUS reachability. Status-server mode provides a standard-compliant method for RADIUS reachability. Use-radius mode requires the configuration of dummy packets that are sent to RADIUS servers. The default is use-radius mode.
SecureEnable	Enable RADIUS Security (RADSec).

Name	Description
UserName	Specifies the user name for RADIUS server reachability. The default is extremenetworks.
Password	Specifies the password for RADIUS server reachability. The default is extremenetworks.
Confirm Password	Confirms the password for RADIUS server reachability.
Unreachable Timer	Specifies, in seconds, the interval between checks when radius server is unreachable. The default is 60 seconds.
Keep Alive Timer	Specifies, in seconds, the interval between checks when radius server is reachable. The default is 180 seconds.

Add a RADIUS Server

Before You Begin

If you specify an FQDN for the server, you must configure DNS for address resolution.

About This Task

Add a RADIUS server to provide RADIUS service on the switch.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Servers** tab.
4. Select **Insert**.
5. Select a value for **AddressType**.
6. In **Address**, type the address or FQDN.
7. In **UsedBy**, select the RADIUS server purpose.
8. (Optional) Configure optional parameters as required.
9. Select **Insert**.

RADIUS Servers Field Descriptions

Use the data in the following table to use the **RADIUS Servers** tab.

Name	Description
AddressType	Specifies the address type as either an IPv4 address, IPv6 address, or dns for a Fully Qualified Domain Name (FQDN).
Address	Specifies the IP address or FQDN of the RADIUS server.

Name	Description
UsedBy	<p>Configures how the server functions:</p> <ul style="list-style-type: none"> • cli—configure the server for CLI authentication. • eapol—configure the server for EAPoL authentication. • endpointTracking—configure the server for Endpoint Tracking authentication. • snmp—configure the server for SNMP accounting. • web—configure the server for HTTP(s) authentication. <p>The default is cli.</p>
Priority	Specifies the priority of each server, or the order of servers to send authentication. The default is 10.
TimeOut	Specifies the time interval in seconds before the client retransmits the packet. The default is 8.
Enable	Enables or disables authentication on the server. The default is true.
MaxRetries	Specifies the maximum number of retransmissions allowed. The default is 1.
UdpPort	<p>Specifies the UDP port that the client uses to send requests to the server. The default value is 1812.</p> <p>The UDP port value set for the client must match the UDP value set for the RADIUS server.</p>
SecretKey	Specifies the RADIUS server secret key, which is the password used by the client to be validated by the server.
AcctEnable	Enables or disable RADIUS accounting. The default is true.
AcctUdpPort	<p>Specifies the UDP port of the RADIUS accounting server. The default value is 1813.</p> <p>The UDP port value configured for the client must match the UDP value configured for the RADIUS server.</p>
SecureEnable	Enable RADIUS Security (RADSec). The default is disabled.
SecureMode	<p>Specifies the RADSec security mode. Possible values are:</p> <ul style="list-style-type: none"> • tls - Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP) • dtls - Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP) <p>The default is tls.</p> <p>Important:</p> <p>To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the <code>radsecproxy.conf</code> configuration file:</p> <pre> tls default{ ... TlsVersion TLS1_2 } </pre>

Name	Description
SecureProfile	Specifies the name of the secure profile.
SecureLogLevel	Specifies the log severity level. Possible values are: <ul style="list-style-type: none">• critical• error• warning• info• debug The default is error.
SecureOCSP	Enables or disable RADIUS Online Certificate Status Protocol (OCSP) checking for this server. The default is disabled.
ResolvedAddressType	Specifies the type of IP address of the RADIUS Server host. It can be IPv4 or IPv6.
ResolvedAddress	Specifies the RADIUS Server resolved IP address associated with the Fully Qualified Domain Name.

Show RADIUS Server Statistics

About This Task

Use the server statistics feature to display the number of input and output packets and the number of input and output bytes. Statistics from console ports are available to assist with debugging.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Servers Stats** tab.

RADIUS Server Stats Field Descriptions

Use the data in the following table to use the **RADIUS Server Stats** tab.

Name	Description
AddressType	Specifies the type of IP address. RADIUS supports IPv4 addresses only.
Address	Shows the IP address of the RADIUS server.
Used by	Identifies the client.
AccessRequests	Shows the number of access-response packets sent to the server; does not include retransmissions.
AccessAccepts	Shows the number of access-accept packets, valid or invalid, received from the server.
AccessRejects	Shows the number of access-reject packets, valid or invalid, received from the server.

Name	Description
BadResponses	Shows the number of invalid access-response packets received from the server.
PendingRequests	Shows the access-request packets sent to the server that have not yet received a response or that have timed out.
ClientRetries	Shows the number of authentication retransmissions to the server.
AcctOnRequests	Shows the number of accounting on requests sent to the server.
AcctOffRequests	Shows the number of accounting off requests sent to the server.
AcctStartRequests	Shows the number of accounting start requests sent to the server.
AcctStopRequests	Shows the number of accounting stop requests sent to the server.
AcctInterimRequests	<p>Number of Accounting Interim requests sent to the server.</p> <p>Important: The AcctInterimRequests counter increments only if you select AcctIncludeCli from the RADIUS Global tab.</p>
AcctBadResponses	Shows the number of Invalid responses discarded from the server.
AcctPendingRequests	Shows the number of requests waiting to be sent to the server.
AcctClientRetries	Shows the number of retries made to this server.
RoundTripTime	Shows the time difference between the instance when a RADIUS request is sent and the corresponding response is received.
AccessChallenges	Shows the number of RADIUS access-challenges packets sent to this server. This does not include retransmission.
NasIpAddress	Shows the RADIUS client NAS Identifier for this server.

Reauthenticate the RADIUS SNMP Server Session

About This Task

Specify the number of challenges that you want the RADIUS SNMP server to send to authenticate a given session.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS SNMP** tab.
4. Select **Enable**.

5. In **ReauthenticateTimer**, type a value to specify the interval between RADIUS SNMP server reauthentications.

**Important**

To end the RADIUS SNMP server session, type a value for the **AbortSessionTimer**, and then select **Enable**.

6. (Optional) Select **AcctEnable**.
7. Select **Apply**.

RADIUS SNMP Field Descriptions

Use the data in the following table to use the **RADIUS SNMP** tab.

Name	Description
Enable	Enables or disables timer authentication on the server. The default is true.
AbortSessionTimer	Specifies the allowable time, in seconds, before aborting the RADIUS SNMP server session (30 to 65535). The default is 180.
ReAuthenticateTimer	Specifies the time, in seconds, between reauthentications of the RADIUS SNMP server (30 to 65535). The default is 180.
AcctEnable	Enables or disables the RADIUS SNMP session timer.
UserName	Specifies the user name for the RADIUS SNMP accounting.

Configure RADIUS SNMP

About This Task

Configure RADIUS SNMP parameters for authentication and session times.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS SNMP** tab.
4. Select **Enable** to enable RADIUS SNMP.
5. In **AbortSessionTimer**, type the period after which the session expires in seconds.
6. In **ReAuthenticateTimer**, type the period of time the system waits before reauthenticating in seconds.
7. Select **AcctEnable** to enable RADIUS accounting for SNMP.
8. In **UserName**, type the RADIUS SNMP user name.
9. Select **Apply**.

RADIUS SNMP *Field Descriptions*

Use the data in the following table to use the **RADIUS SNMP** tab.

Name	Description
Enable	Enables or disables timer authentication on the server. The default is true.
AbortSessionTimer	Specifies the allowable time, in seconds, before aborting the RADIUS SNMP server session (30 to 65535). The default is 180.
ReAuthenticateTimer	Specifies the time, in seconds, between reauthentications of the RADIUS SNMP server (30 to 65535). The default is 180.
AcctEnable	Enables or disables the RADIUS SNMP session timer.
UserName	Specifies the user name for the RADIUS SNMP accounting.

Modify a RADIUS Configuration

About This Task

Use this procedure to modify an existing RADIUS configuration or single function such as retransmissions and RADIUS accounting.

RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all except the following case. When modifying a RADIUS configuration in Enterprise Device Manager (EDM), you must specify if the address type is an IPv4 or an IPv6 address.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Servers** tab.
4. In the row and field to modify, type the information or use the lists to make a selection. Access the lists by double-clicking in a field.
5. Select **Apply**.

RADIUS Servers *Field Descriptions*

Use the data in the following table to use the **RADIUS Servers** tab.

Name	Description
AddressType	Specifies the address type as either an IPv4 address, IPv6 address, or dns for a Fully Qualified Domain Name (FQDN).
Address	Specifies the IP address or FQDN of the RADIUS server.

Name	Description
UsedBy	<p>Configures how the server functions:</p> <ul style="list-style-type: none"> • cli—configure the server for CLI authentication. • eapol—configure the server for EAPoL authentication. • endpointTracking—configure the server for Endpoint Tracking authentication. • snmp—configure the server for SNMP accounting. • web—configure the server for HTTP(s) authentication. <p>The default is cli.</p>
Priority	Specifies the priority of each server, or the order of servers to send authentication. The default is 10.
TimeOut	Specifies the time interval in seconds before the client retransmits the packet. The default is 8.
Enable	Enables or disables authentication on the server. The default is true.
MaxRetries	Specifies the maximum number of retransmissions allowed. The default is 1.
UdpPort	<p>Specifies the UDP port that the client uses to send requests to the server. The default value is 1812.</p> <p>The UDP port value set for the client must match the UDP value set for the RADIUS server.</p>
SecretKey	Specifies the RADIUS server secret key, which is the password used by the client to be validated by the server.
AcctEnable	Enables or disable RADIUS accounting. The default is true.
AcctUdpPort	<p>Specifies the UDP port of the RADIUS accounting server. The default value is 1813.</p> <p>The UDP port value configured for the client must match the UDP value configured for the RADIUS server.</p>
SecureEnable	Enable RADIUS Security (RADSec). The default is disabled.
SecureMode	<p>Specifies the RADSec security mode. Possible values are:</p> <ul style="list-style-type: none"> • tls - Transport Layer Security (TLS) encryption over Transmission Control Protocol (TCP) • dtls - Datagram Transport Layer Security (DTLS) encryption over User Datagram Protocol (UDP) <p>The default is tls.</p> <p>Important:</p> <p>To avoid TLS handshake issues if the switch and RADsec proxy server run different versions of OpenSSL, manually force TLS version 2 negotiation through the RADsec proxy by adding the following text to the <code>radsecproxy.conf</code> configuration file:</p> <pre> tls default{ ... TlsVersion TLS1_2 } </pre>

Name	Description
SecureProfile	Specifies the name of the secure profile.
SecureLogLevel	Specifies the log severity level. Possible values are: <ul style="list-style-type: none">• critical• error• warning• info• debug The default is error.
SecureOCSP	Enables or disable RADIUS Online Certificate Status Protocol (OCSP) checking for this server. The default is disabled.
ResolvedAddressType	Specifies the type of IP address of the RADIUS Server host. It can be IPv4 or IPv6.
ResolvedAddress	Specifies the RADIUS Server resolved IP address associated with the Fully Qualified Domain Name.

Delete a RADIUS Configuration

About This Task

Delete an existing RADIUS configuration.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Servers** tab.
4. Identify the configuration to delete by selecting anywhere in the row.
5. Select **Delete**.

Configure RADIUS Server Reachability

About This Task

Use this procedure to configure the RADIUS server reachability settings.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **RADIUS Global** tab.
4. In **RadiusReachability**, select either **useStatusServerPackets** or **useDummyRadiusRequests**.
5. In **UserName**, type the reachability user name.
6. In **Password**, type the reachability password.
7. In **Confirm Password**, retype the reachability password.

8. In **Unreachable Timer**, type the interval in seconds between checks when the RADIUS server is unreachable.
9. In **KeepAlive Timer**, type the interval in seconds between checks when the RADIUS server is reachable.
10. Select the **Apply**.

RADIUS Global Field Descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The default is 192.
AcctEnable	Enables RADIUS accounting.
AcctAttrValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
CliCmdCount	Gives the value for the CLI command count. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.
SupportedVendorIds	Shows the vendor IDs that the software supports for RADIUS.

Name	Description
RadiusReachability	Specifies the mode for RADIUS reachability. Status-server mode provides a standard-compliant method for RADIUS reachability. Use-radius mode requires the configuration of dummy packets that are sent to RADIUS servers. The default is use-radius mode.
SecureEnable	Enable RADIUS Security (RADSec).
UserName	Specifies the user name for RADIUS server reachability. The default is extremenetworks.
Password	Specifies the password for RADIUS server reachability. The default is extremenetworks.
Confirm Password	Confirms the password for RADIUS server reachability.
Unreachable Timer	Specifies, in seconds, the interval between checks when radius server is unreachable. The default is 60 seconds.
Keep Alive Timer	Specifies, in seconds, the interval between checks when radius server is reachable. The default is 180 seconds.

Configure a RADIUS Dynamic-Server Client

Before You Begin

- You must enable EAPOL globally and at the port level.
- You must enable RADIUS dynamic extensions command processing at the port level.
- If you specify an FQDN, you must configure DNS for address resolution.

About This Task

Use this procedure to configure a client to process dynamic session changes. You can configure up to eight clients.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS CoA**.
3. Select **Global**.
4. Select **Insert**.
5. In **AddressType**, select **ipv4** or **dns**.
6. In **Address**, type a client address or FQDN.
7. (Optional) In **UdpPort**, type a value for the client port.
8. (Optional) In **Secret**, type a value for the secret key.
9. Select **Enable** to enable dynamic session processing for the client.
10. Select **Insert**.

Global Field Descriptions

Use the data in the following table to use the **Global** tab.

Name	Description
AddressType	Specifies the type of address contained in the corresponding instance of a RADIUS Dynamic Authorization Client configured in this entry. The address can be either an IPv4 address or dns for a Fully Qualified Domain Name (FQDN).
Address	Specifies the internet address or FQDN of a RADIUS Dynamic Authorization Client configured in this entry.
UdpPort	Specifies the UDP port number the server/NAS listens on for requests from the RADIUS Dynamic Authorization Client configured in this entry. The default is 3799.
Secret	Specifies the secret key shared between the RADIUS Dynamic Authorization Client and Server. Note that when this object is retrieved, the value always displays as a zero-length octet string.
Enable	Enable processing of dynamic session changes. The default is disabled.
ResolvedAddressType	Specifies the type of address that DNS resolves for the RADIUS Dynamic Authorization Client. The address can be an IPv4 address.
ResolvedAddress	Specifies the RADIUS Dynamic Authorization Client resolved IP address associated with the Fully Qualified Domain Name.

Display RADIUS CoA Disconnect Statistics

About This Task

Use this procedure to display RADIUS CoA Disconnect statistics.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS CoA**.
3. Select the **Disconnect Stats** tab.

Disconnect Stats *Field Descriptions*

Use the data in the following table to use the **Disconnect Stats** tab.

Name	Description
DisconRequests	Specifies the number of RADIUS disconnect requests received from this Dynamic Authorization Client. This also includes the RADIUS disconnect requests that have a Service-Type attribute with a value of Authorize Only.
DisconAuthOnlyRequests	Specifies the number of RADIUS disconnect requests that include a Service-Type attribute with a value of Authorize Only received from this Dynamic Authorization Client.
DupDisconRequests	Specifies the number of duplicate RADIUS Disconnect-Request packets received from this Dynamic Authorization Client.
DisconAcks	Specifies the number of RADIUS Disconnect-ACK packets sent to this Dynamic Authorization Client.
DisconNaks	Specifies the number of RADIUS Disconnect-NAK packets sent to this Dynamic Authorization Client. This includes the RADIUS Disconnect-NAK packets sent with a Service-Type attribute with a value of Authorize Only and the RADIUS Disconnect-NAK packets sent because no session context was found.
DisconNakAuthOnlyRequests	Specifies the number of RADIUS Disconnect-NAK packets that include a Service-Type attribute with a value of Authorize Only sent to this Dynamic Authorization Client.
DisconNakSessNoContext	Specifies the number of RADIUS Disconnect-NAK packets sent to this Dynamic Authorization Client because no session context was found.
DisconUserSessRemoved	Specifies the number of user sessions removed for the disconnect requests received from this Dynamic Authorization Client. Depending on site-specific policies, a single disconnect request can remove multiple user sessions. In cases where this Dynamic Authorization Server has no knowledge of the number of user sessions that are affected by a single request, each such disconnect request counts as a single affected user session only.

Name	Description
MalformedDisconRequests	Specifies the number of malformed RADIUS Disconnect-Request packets received from this Dynamic Authorization Client. Bad authenticators and unknown types are not included as malformed disconnect requests
DisconBadAuthenticators	Specifies the number of RADIUS Disconnect-Request packets that contain an invalid Authenticator field received from this Dynamic Authorization Client.
DisconPacketsDropped	Specifies the number of incoming disconnect requests from this Dynamic Authorization Client silently discarded by the server application for some reason other than malformed packets, bad authenticators, or unknown types.
UnknownTypes	Specifies the number of incoming requests that have an invalid ID type.

Display RADIUS CoA Reauthenticate Statistics

About This Task

Use this procedure to display RADIUS CoA Reauthenticate statistics.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS CoA**.
3. Select the **Reauthenticate Stats** tab.

Reauthenticate Stats Field Descriptions

Use the data in the following table to use the **Reauthenticate Stats** tab.

Name	Description
Requests	Specifies the number of RADIUS Reauthentication-Requests received from this Dynamic Authorization Client. This also includes the Reauthentication requests that have a Service-Type attribute with a value of Authorize Only.
AuthOnlyRequests	Specifies the number of RADIUS Reauthentication-Requests that include a Service-Type attribute with value Authorize Only received from this Dynamic Authorization Client.

Name	Description
DupRequests	Specifies the number of duplicate RADIUS Reauthentication-Request packets received from this Dynamic Authorization Client.
Acks	Specifies the number of incoming Reauthentication packets from this Dynamic Authorization Client silently discarded by the server application for some reason other than malformed, bad authenticators, or unknown types.
Nacks	Specifies the number of RADIUS Reauthentication-NAK packets sent to this Dynamic Authorization Client. This includes the RADIUS Reauthentication-NAK packets sent with a Service-Type attribute value of Authorize Only, and the RADIUS Reauthentication-NAK packets sent because no session context was found.
NacksAuthOnlyRequests	Specifies the number of RADIUS Reauthentication-NAK packets that include a Service-Type attribute with a value of Authorize Only sent to this Dynamic Authorization Client.
NacksNoSess	Specifies the number of RADIUS Reauthentication-NAK packets sent to this Dynamic Authorization Client because no session context was found.
SessReauthenticated	Specifies the number of user sessions reauthenticated for the Reauthentication-Requests received from this Dynamic Authorization Client. Depending on site-specific policies, a single Reauthentication-Request can change the authorization of multiple user sessions. In cases where the Dynamic Authorization Server has no knowledge of the number of user sessions that are affected by a single request, each CoA-Request counts as a single affected user session.
Malformed	Specifies the number of malformed RADIUS Reauthentication-Request packets received from the Dynamic Authorization Client. Bad authenticators and unknown types are not included as malformed Reauthentication-Requests.

Name	Description
Dropped	Specifies the number of incoming Reauthentication packets from the Dynamic Authorization Client silently discarded by the server application for some reason other than malformed, bad authenticators, or unknown types.
BadAuths	Specifies the number of RADIUS Reauthentication-Request packets that contained an invalid Authenticator field received from the Dynamic Authorization Client.
UnknownTypes	Specifies the number of incoming requests that have an invalid ID type.

Display RADIUS CoA Statistics

About This Task

Use this procedure to display information about RADIUS CoA statistics.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS CoA**.
3. Select the **CoA Stats** tab.

CoA Stats Field Descriptions

Use the data in the following table to use the **CoA Stats** tab.

Name	Description
AddressType	Specifies the RADIUS Dynamic Authorization Client IP address type.
Address	Specifies the RADIUS Dynamic Authorization Client IP address.
DisconRequests	Specifies the number of RADIUS Disconnect-Requests received from this Dynamic Authorization Client. This also includes the Disconnect requests that have a Service-Type attribute with a value of Authorize Only.
DisconAuthOnlyRequests	Specifies the number of RADIUS Disconnect-Requests that include a Service-Type attribute with value Authorize Only received from this Dynamic Authorization Client.
DupDisconRequests	Specifies the number of duplicate RADIUS Disconnect-Request packets received from this Dynamic Authorization Client.

Name	Description
DisconAcks	Specifies the number of RADIUS Disconnect-ACK packets sent to this Dynamic Authorization Client.
DisconNaks	Specifies the number of RADIUS Disconnect-ACK packets sent to this Dynamic Authorization Client. This includes the RADIUS Disconnect-ACK packets sent with a Service-Type attribute value of Authorize Only, and the RADIUS Disconnect-ACK packets sent because no session context was found.
DisconNakAuthOnlyRequests	Specifies the number of RADIUS Disconnect-NAK packets that include a Service-Type attribute with a value of Authorize Only sent to this Dynamic Authorization Client.
DisconNakSessNoContext	Specifies the number of RADIUS Disconnect-NAK packets sent to this Dynamic Authorization Client because no session context was found.
DisconUserSessRemoved	Specifies the number of user sessions removed for the Disconnect-Requests received from this Dynamic Authorization Client. Depending on site-specific policies, a single Disconnect-Request can change the authorization of multiple user sessions. In cases where the Dynamic Authorization Server has no knowledge of the number of user sessions that are affected by a single request, each Disconnect-Request counts as a single affected user session.
MalformedDisconRequests	Specifies the number of malformed RADIUS Disconnect-Request packets received from the Dynamic Authorization Client. Bad authenticators and unknown types are not included as malformed Disconnect-Requests.
UnknownTypes	Specifies the number of incoming requests that have an invalid ID type.

RADSec Configuration Using EDM

This section contains procedures to configure RADSec with Enterprise Device Manager (EDM).

Configure a RADIUS Secure Profile using Manual Certificates

Configure RADIUS secure profiles with certificate information, certificate key information, and password information, which enables RADSec peers to establish connections.

About This Task

A new profile directory is created for each new profile in the `/intflash/.radsec/profile/radsec` directory. Profile configuration file "profile_info.cfg" is available in `/intflash/.radsec/profile` directory.



Note

All the files (certificates and keys) must be in .pem format and you must copy them to the `/intflash` directory.

You can configure a maximum of 10 RADSec profiles.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **RADIUS**.
3. Select the **Secure Certificate Profile** tab.
4. Select **Insert**.
5. For **Name**, enter the name for the profile.
6. For **RootCert**, enter the full path for the profile root certificate.
7. For **Cert**, enter the full path for the profile certificate.
8. For **Key**, enter the full path for the profile private key.
9. For **Password**, enter the password for the profile private key.
10. For **RootCertDestFile**, enter the file name to use when installing the root certificate.
11. For **CertDestFile**, enter the file name to use when installing the certificate.
12. For **KeyDestFile**, enter the file name to use when installing the key.
13. Select **Insert**.

Secure Certificate Profile Field Descriptions

Use the data in the following table to use the **Secure Certificate Profile** tab.

Name	Description
Name	Specifies the profile name.
RootCert	Specifies the full path of the profile root certificate. This parameter applies to certificates you install and manage manually.
Cert	Specifies the full path of the profile certificate. This parameter applies to certificates you install and manage manually.

Name	Description
Key	Specifies the full path of the profile private key. This parameter applies to certificates you install and manage manually.
Password	Specifies the password used to decrypt the profile private key.
RootCertDestFile	Specifies the file name to use when installing the root certificate. This parameter applies to certificates you install and manage manually.
CertDestFile	Specifies the file name to use when installing the certificate. This parameter applies to certificates you install and manage manually.
KeyDestFile	Specifies the file name to use when installing the key. This parameter applies to certificates you install and manage manually.



RADIUS Attributes

[Vendor Specific Attributes](#) on page 3108

[IETF Attributes](#) on page 3133

The following topics list RADIUS attributes for EAP.

Vendor Specific Attributes

RADIUS Vendor Specific Frame Format

The following figure shows the frame format for RADIUS vendor specific attributes (VSA).




Figure 205: RADIUS vendor-specific frame format

FA-VLAN-ISID

On EAP-enabled Flex-UNI ports, the context for the authenticating MAC is the association using the FA-VLAN-ISID attribute. This attribute provides one VLAN:ISID binding.

- defined under vendor IDs 562 and 1584 and uses the value 171
- I-SID values are 0-15999999
- VLAN values are 0 - 4096 to create untagged S-UNIs in MHSA mode and MAC-BasedS-UNIs in MHMV mode
- RADIUS server provides up to 94 bindings
- for MHMV with multiple bindings, only the last binding is created

FA-Client-Type

This attribute is used in packets transmitted to the RADIUS server to inform that the current authenticated MAC is associated to a FA Client of type x. This attribute is defined under Nortel vendor id (562) and uses the value 182.

Only one FA Client per port is supported.

The following values are defined for FA Client type:

- 0: FA Element Not Found
- 1: FA Element Type Other
- 2: FA Server
- 3: FA Proxy
- 4: FA Server No Authentication
- 5: FA Proxy No Authentication
- 6: FA Client Wireless AP Type 1 [clients direct network attachment]
- 7: FA Client Wireless AP Type 2 [clients tunneled to controller]
- 8: FA Client Switch
- 9: FA Client Router
- 10: FA Client IP Phone
- 11: FA Client Camera
- 12: FA Client IP Video
- 13: FA Client Security Device
- 14: FA Client Virtual Switch
- 15: FA Client Server Endpoint
- 63: FA Client Unknown

Any integer value provided by the FA is sent to the RADIUS server without any check from the EAP application.

Extreme-Dynamic-MHSA

On EAP-enabled Flex UNI ports, this attribute changes the current interface EAPoL operating mode from MHMV to MHSA.

- defined under Extreme Networks vendor ID 1916 and uses the value 250
- parsed only on ports with Flex-UNI enabled
- value contained in this attribute must be 1

Extreme-Dynamic-ACL

See the following sections:

- [Introduction](#) on page 3110
- [Examples](#) on page 3112
- [List Parameter](#) on page 3112
- [List Examples](#) on page 3113
- [Shortened Syntax](#) on page 3114
- [Shortened Syntax Example](#) on page 3116

Introduction

This attribute assigns a dynamic ACL for an EAP-enabled port. The dynamic behavior of the ACL depends on the EAP port state (MHMV or MHSA).

This VSA is defined under Extreme Networks vendor ID 1916 and uses the value 251. The VSA must configure the following keywords in order:



Note

This requirement does not apply when you configure the default-action.

1. **CLIENT**—Each group of Extreme-Dynamic-ACL messages must begin with a **CLIENT** **<client_name>** command to identify the EAP client.
2. **acl**—Create an ACL for the client. The port type is mandatory (outPort or inPort) but the name is optional. Packet type (pktType) supports only nonipv6.
3. **ace**—Create ACEs to add rules to the ACL. Separate each rule by an ampersand (&). You must specify an ACE action and if the ACE is a security or QoS ACE.

The following tables identify the support differences between CLI and VSA, and show how to transform a CLI command to VSA.

Table 243: Support differences

Filter component	CLI	VSA
ACL type	inVlan inPort outPort inVsn	inPort outPort
Packet type	IPv4 IPv6	IPv4 (nonipv6)

Table 243: Support differences (continued)

Filter component	CLI	VSA
Default actions	permit deny control-packet-action policer	permit deny control-packet-action
ACE type	ARP Ethernet IP IPv6 protocol policer	ARP Ethernet IP protocol

Table 244: CLI to VSA transformation

CLI	VSA	Notes
<code>filter acl <1-2048> type { inPort outPort } [name <name>]</code>	<code>acl type { inPort outPort } [name <name>]</code>	ACL creation If you do not provide a name, then the implicit ACL name is EAP-MHMOV-<port> or EAP-MHSA-<port>.
<code>filter acl set <1-2048> {default- action {permit deny [control-packet- action]]} global- action { monitor-dst- mlt monitor-dst- ports} }</code>	<code>acl set {default-action {permit deny [control- packet-action]]} global-action { monitor-dst-mlt monitor-dst-ports} }</code>	ACL global configuration
<code>filter acl ace <1-2048> <1-2000> [name WORD<0-32>]</code>	<code>ace <number> qos sec</code>	ACE creation The ACE name is automatically created as HostName-Rule<ACE ID>.
<code>filter acl ace {arp ethernet ip protocol } <1-2048> <1-2000> <options></code>	<code>ace <number> qos sec {arp ethernet ip protocol } <options></code>	ACE qualifiers

Examples

The following examples provide a RADIUS configuration, followed by the VSA syntax for the same configuration. This example is for MAC 0a:0a:0a:0a:0a:0a on port 1/1 and EAP is in MHMV mode.



Note

Count is automatically added to the ACE even if you do not specify it explicitly.

```
filter acl 1 type inPort
filter acl port 1 1/1

filter acl ace 1 1 name RadiusGuest-Rule01
filter acl ace ethernet 1 1 src-mac eq 0a:0a:0a:0a:0a:0a
filter acl ace ethernet 1 1 ether-type eq 0x800
filter acl ace ip 1 1 ip-protocol-type eq 17
filter acl ace protocol 1 1 dst-port eq 53
filter acl ace 1 1 action permit count
filter acl ace 1 1 enable

filter acl ace 1 2 name RadiusGuest-Rule02
filter acl ace ethernet 1 2 src-mac eq 0a:0a:0a:0a:0a:0a
filter acl ace ethernet 1 2 ether-type eq 0x800
filter acl ace ip 1 2 dst-ip mask 192.0.2.1 24
filter acl ace 1 2 action permit count
filter acl ace 1 2 enable

filter acl ace 1 3 name RadiusGuest-Rule03
filter acl ace ethernet 1 3 src-mac eq 0a:0a:0a:0a:0a:0a
filter acl ace 1 3 action deny count
filter acl ace 1 3 enable
```

The RADIUS VSA does not specify the MAC nor the port number because they are already known at the EAP level.

```
Extreme-Dynamic-ACL = "CLIENT RadiusGuest",
Extreme-Dynamic-ACL += "acl inPort",
Extreme-Dynamic-ACL += "ace 1 sec ethernet ether-type eq 0x800 & ip ip-protocol-type eq 17 & protocol dst-port eq 53 action permit count",
Extreme-Dynamic-ACL += "ace 2 sec ethernet ether-type eq 0x800 & ip dst-ip mask 192.0.2.1 24 action permit count",
Extreme-Dynamic-ACL += "ace 3 sec action deny count"
```

The following ACE syntax example provides the ability to remark DSCP value for IP traffic (0x800):

```
Extreme-Dynamic-ACL = "ace 1 qos action permit internal-qos 5 remark-dot1p 5 remark-dscp phbaf41 & ethernet ether-type eq 0x800",
Extreme-Dynamic-ACL += "acl set default-action permit"
```

For more information and additional examples, see [RADIUS Dynamic User-Based Policies](#) on page 1023.

List Parameter

You can use the list parameter with the Extreme-Dynamic-ACL RADIUS VSA to increase the number of ACEs sent in a RADIUS response packet, which can only be 4,096 bytes long. Use the list parameter to group ACE match values like ports and masks when the rest of the ACE match criteria is common.

For example, if you need to create 10 ACLs for different destination ports, rather than create 10 individual VSAs with almost identical syntax, you can create one VSA using the list parameter to group the destination ports, separating each port in the list with a comma. For illustration purposes, see [List Examples](#).

Each Extreme-Dynamic-ACL can use only one list parameter.

List Examples

The following examples provide both the list and expanded syntax for the same configuration.

In the following table, the expanded syntax configures three ACEs, each for a specific destination port, that are then combined into one syntax statement by using the list parameter to group the different destination ports.

Table 245: Port list comparison

Expanded syntax	List syntax
Extreme-Dynamic-ACL += 'ace 1 ethernet ether-type equal ip & ip ip-protocol-type eq tcp & protocol dst-port equal 37 & action permit count', Extreme-Dynamic-ACL += 'ace 2 ethernet ether-type equal ip & ip ip-protocol-type eq tcp & protocol dst-port equal 38 & action permit count', Extreme-Dynamic-ACL += 'ace 3 ethernet ether-type equal ip & ip ip-protocol-type eq tcp & protocol dst-port equal 1427 & action permit count',	Extreme-Dynamic-ACL += 'ace 1 ethernet ether-type equal ip & ip ip-protocol-type eq tcp & protocol dst-port equal list 37, 38, 1427 & action permit count',

In the following table, the expanded syntax configures six ACEs, each for a specific source mask, that are then combined into one syntax statement by using the list parameter to group the different source masks.

Table 246: Mask list comparison

Expanded	List
<pre>Extreme-Dynamic-ACL += 'ace 1 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask 9100 0xFFFFC & ip dst mask 172.16.0.0 12 & action permit count', Extreme-Dynamic-ACL += 'ace 2 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask 9220 0xFFFFE & ip dst mask 172.16.0.0 12 & action permit count', Extreme-Dynamic-ACL += 'ace 3 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask 9290 0xFFFFE & ip dst mask 172.16.0.0 12 & action permit count', Extreme-Dynamic-ACL += 'ace 4 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask 49152 0xC000 & ip dst mask 172.16.0.0 12 & action permit count', Extreme-Dynamic-ACL += 'ace 5 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask 800 0xFFFF8 & ip dst mask 172.16.0.0 12 & action permit count', Extreme-Dynamic-ACL += 'ace 6 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask 808 0xFFFFE & ip dst mask 172.16.0.0 12 & action permit count',</pre>	<pre>Extreme-Dynamic-ACL += 'ace 1 ethernet ether-type equal ip & ip-protocol-type equal tcp & protocol source mask list 9100 0xFFFFC, 9220 0xFFFFE, 9290 0xFFFFE, 49152 0xC000, 800 0xFFFF8, 808 0xFFFFE & ip dst mask 172.16.0.0 12 & action permit count',</pre>

Shortened Syntax

The previous section explains how you can increase the number of ACEs the RADIUS server sends in a single RADIUS response. You can further reduce the character count of the syntax statements by using the same shortened command form used with [CLI Autocompletion](#) on page 187.

Use a shortened syntax to automatically fill in the unique parts of a command string rather than typing the entire command. Use the CLI to determine the appropriate shortened form for each command argument.

The same letter or letter combination can represent many arguments depending on the context of the preceding argument. For example, **r** is both **routed-only** and **routing-type** in the following CLI command completion examples:

```
Switch:1(config)#filter acl ace ip 1 1 r<TAB>
Switch:1(config)#filter acl ace ip 1 1 routed-only

#OR
```



```
Switch:1(config)#filter acl ace protocol 1 1 r<TAB>  
Switch:1(config)#filter acl ace protocol 1 1 routing-type
```

**Note**

Keep in mind that new releases can modify or add commands, which means the uniqueness for command completion can change. If another argument that begins with the letters **ro** is added under the **protocol** argument, **r** is no longer unique to routing-type.

Shortened Syntax Example

The following example illustrates how the list argument used with the CLI shortened syntax can reduce the character count. The full syntax is first reduced by using the list argument, and then further reduced with CLI shortened syntax.

Table 247: List and Shortened syntax example

Syntax	Character count
<pre>Extreme-Dynamic-ACL += "ace 13 ethernet ether-type equal 0x800 & ip ip-protocol- type equal tcp & protocol dst-port equal 5672 & action permit", Extreme-Dynamic-ACL += "ace 14 ethernet ether-type equal 0x800 & ip ip- protocol-type equal tcp & protocol dst- port equal 9734 & action permit", Extreme-Dynamic-ACL += "ace 15 ethernet ether-type equal 0x800 & ip ip- protocol-type equal tcp & protocol dst- port equal 1500 & action permit", Extreme-Dynamic-ACL += "ace 16 ethernet ether-type equal 0x800 & ip ip- protocol-type equal tcp & protocol dst- port equal 4000 & action permit", Extreme-Dynamic-ACL += "ace 17 ethernet ether-type equal 0x800 & ip ip- protocol-type equal tcp & protocol dst- port equal 3306 & action permit", Extreme-Dynamic-ACL += "ace 18 ethernet ether-type equal 0x800 & ip ip- protocol-type equal tcp & protocol dst- port equal 7000 & action permit", Extreme-Dynamic-ACL += "ace 19 ethernet ether-type equal 0x800 & ip ip- protocol-type equal tcp & protocol dst- port equal 5001 & action permit",</pre>	819
<pre>Extreme-Dynamic-ACL += "ace 13 ethernet ether-type equal 0x800 & ip ip-protocol- type equal tcp & protocol dst-port equal list 5672,9734,1500,4000,3306,7000,5001 & action permit",</pre>	152
<pre>Extreme-Dynamic-ACL += "ace 13 et e eq 0x800 & ip ip-p eq tcp & pr d eq list 5672,9734,1500,4000,3306,7000,5001 & action permit",</pre>	103

Extreme-Dynamic-Config

See the following sections:

- [Introduction](#) on page 3117
- [AN-ADVERTISEMENTS:100Half or AN-ADVERTISEMENTS:100H](#) on page 3118
- [BPDU](#) on page 3118

- [DAI](#) on page 3118
- [DHCP Snooping](#) on page 3118
- [IGMP Snooping](#) on page 3119
- [IGMPv3](#) on page 3119
- [IGMP Fast Leave](#) on page 3119
- [IPSG](#) on page 3119
- [Port Bounce](#) on page 3119
- [Re-authentication](#) or [Re-authentication:100](#) on page 3120
- [SLPP Guard](#) or [SLPP Guard:<xxxxx>](#) on page 3120
- [WOL](#) on page 3120
- [Session Timeout and Re-authentication VSA Behavior](#) on page 3120
- [Expected Behavior for DHCP Snooping and DAI Vendor Specific Attributes](#) on page 3121
- [Expected Behavior of IGMPv3 and IGMP Fast Leave Vendor Specific Attributes](#) on page 3123
- [Extreme-Dynamic-Config Operational Considerations](#) on page 3124
- [Extreme-Dynamic-Config Examples](#) on page 3126

Introduction

This VSA configures port- and VLAN-based attributes.

- Name: Extreme-Dyn-Config
- Value: 252
- Type: String
- Vendor: Extreme Networks
- Extreme Networks Vendor ID is 1916

The following features can be configured using Extreme-Dynamic-Config RADIUS VSA:

- VLAN-based features:
 - DHCP Snooping
 - Dynamic ARP Inspection (DAI)
 - IGMP Snooping
 - IGMP version 3 (IGMPv3)
 - IGMP Fast Leave
- Port-based features:
 - Custom Auto-Negotiation Advertisements (CANA)
 - Bridge Protocol Data Unit (BPDU) Guard
 - IP Source Guard (IPSG)
 - Port bounce
 - Re-authentication

- Simple Loop Prevention Protocol (SLPP) Guard
- Traffic Control (Wake on LAN - WoL)

**Note**

If you use the Extreme-Dynamic-Config RADIUS VSA, then Dynamic ARP Inspection and DHCP Snooping are only enabled on the default VLAN if IPSP is present in the RADIUS attributes.

AN-ADVERTISEMENTS:100Half or AN-ADVERTISEMENTS:100H

AN-ADVERTISEMENTS:100Half or AN-ADVERTISEMENTS:100H settings configure CANA speed and duplex to the following supported values:

- 10Half
- 100Full
- 100Half
- 100Full
- 1000Full

BPDUGUARD

This attribute enables BPDU Guard on the port where the client resides.

DAI

This attribute enables DAI on the VLAN received from the RADIUS server. For a Flex-UNI port, DAI is enabled on the platform VLAN associated with the I-SID received from the RADIUS server.

DAI is also enabled on the default VLAN of the port to prepare for IPSP, which requires DAI and DHCP Snooping enabled on all VLANs. If the RADIUS server does not return a VLAN of I-SID, DAI is enabled on the default VLAN. For Flex-UNI ports, DAI is enabled on the platform VLAN associated with the untagged I-SID.

**Note**

If you use the Extreme-Dynamic-Config RADIUS VSA, then DAI is only enabled on the default VLAN if IPSP is present in the RADIUS attributes.

DHCP Snooping

This attribute enables DHCP Snooping on the VLAN received from the RADIUS server. For a Flex-UNI port, DHCP Snooping is enabled on the platform VLAN associated with the I-SID received from the RADIUS server.

DHCP Snooping is also enabled on the default VLAN of that port to prepare for IPSP, which requires DAI and DHCP Snooping enabled on all VLANs. If the RADIUS server does not return a VLAN of I-SID, DHCP Snooping is enabled on the default VLAN. For

Flex-UNI ports, DHCP Snooping is enabled for the platform VLAN associated with the untagged I-SID.

**Note**

If you use the Extreme-Dynamic-Config RADIUS VSA, then DHCP Snooping is only enabled on the default VLAN if IPSPG is present in the RADIUS attributes.

IGMPSNOOP

This attribute enables IGMP Snooping on the VLAN received from the RADIUS server. For a Flex-UNI port, IGMP Snooping is enabled on the platform VLAN associated with the I-SID received from the RADIUS server.

IGMPV3

This attribute enables IGMP version 3 on the VLAN specified by the RADIUS server. If IGMP Snooping is enabled on that VLAN, the IGMP version settings provided by the RADIUS VSA apply to the same VLAN interface. For a Flex-UNI port, IGMPv3 is enabled on the platform VLAN associated with the I-SID received from the RADIUS server. Applying this attribute disables IGMP Fast Leave (if previously enabled) and enables IGMPv3 compatibility mode.

IGMPFAST

This attribute enables IGMP Fast Leave on the VLAN specified by the RADIUS server. If IGMP Snooping is enabled on that VLAN, the IGMP Fast Leave settings provided by the RADIUS VSA apply to the same VLAN interface. For a Flex-UNI port, IGMP Fast Leave is enabled on the platform VLAN associated with the I-SID received from the RADIUS server. With IGMPv2, Fast Leave is enabled; with IGMPv3, Fast Leave is enabled along with explicit host tracking.

IPSG

This attribute enables IPSG on the port where the client resides.

In order to apply IPSG, DHCP Snooping and DAI must be configured on the RADIUS server. DHCP Snooping and DAI must be enabled on all VLANs.

The following is an example of a log message that displays if a setting is not configured correctly:

```
GlobalRouter EAP WARNING Cannot apply Radius IP Source Guard attribute
on port 1/15 without DHCP Snooping and DAI attributes.
```

PORTBOUNCE

This attribute runs a port bounce on an EAP port including Power over Ethernet (PoE), if capable. Current EAP sessions are preserved during the port bounce process.

Use this attribute for on-demand re-authentication, for example, to use a power cycle to trigger a port bounce for a PC behind a phone, which is being rebooted, as in the following use cases:

- To ensure a PC behind a phone acquires an IP address from the VLAN or I-SID assigned by the RADIUS server, the PC triggers a new DHCP negotiation if it receives the port-bounce attribute in the RADIUS-Accept packet.
- To force a client to renew its DHCP IP address by using a RADIUS dynamic Change-of-Authorization (CoA) or Disconnect request.

CLI and EDM output that displays the interface state indicates if a port is down because of a RADIUS VSA port bounce.

REAUTH or REAUTH:100

This attribute enables EAPOL re-authentication on a port either manually using CLI or dynamically through RADIUS. The origin identifies how re-authentication was configured, either CONFIG or RADIUS.

SLPPGUARD or SLPPGUARD:<xxxxxx>

This attribute enables SLPP Guard on the port where the client resides. You can configure a timer to enable the port only after the timer expires. The default timer value is 60 seconds.

WOL

This attribute enables EAP traffic-control (Wake On LAN) on the port where the client resides.

Session Timeout and Re-authentication VSA Behavior

Session	REAUTH status	Cause
EAP Session	Without REAUTH VSA	On MAC age out. Note: Session timeouts if the client is not connected.
	With REAUTH VSA	On periodically timer and MAC age out.

Session	REAUTH status	Cause
NEAP Session	Without REAUTH VSA	On MAC age out.
	With REAUTH VSA	Useful for silent devices, session stays active despite MAC age out. Session is removed only by manual intervention or RADIUS reject/timeout. Note: The same command activates both EAP and NEAP re-authentication. If re-authentication is needed for EAP, NEAP re-authentication for silent devices is automatically activated.

Command	Level	Processing	Configuration	Prerequisites
Session-timeout	per session basis.	Changes the reauth interval for a particular session.	none.	Enable port level re-authentication.
VSA (REAUTH:300 or REAUTH)	per port basis.	Changes the port configuration.	Enable re-authentication and configure interval. Caution: Configure the re-authentication period value to 0 to prevent EAP or NEAP sessions from aging out. Preventing re-authentication can introduce a security risk.	none.

Expected Behavior for DHCP Snooping and DAI Vendor Specific Attributes

The following table shows the behavior for Dynamic Host Configuration Protocol Snooping (DHCP Snooping) and Dynamic ARP Inspection (DAI) Vendor Specific Attributes (VSA) for each authentication scenario, depending on Extensible

Authentication Protocol (EAP) Operational mode, Flex-UNI mode, and IP Source Guard (IPSG) VSA or RADIUS VLAN/I-SID.

Table 248: Expected Behavior for DHCP Snooping and DAI VSAs

EAP Operational Mode	IPSG VSA received	Flexi-UNI	RADIUS VLAN/I-SID	DHCP Snooping and DAI VSA received
Mutiple Host Single Authentication (MHSA)	Yes	Yes	Yes	DHCP, DAI, and IPSG enabled on all associated Platform VLANs.
MHSA	Yes	Yes	No	DHCP, DAI, and IPSG enabled on all associated Platform VLANs.
MHSA	Yes	No	Yes	DHCP, DAI, and IPSG enabled on all associated Platform VLANs.
MHSA	Yes	No	No	DHCP, DAI, and IPSG enabled on all associated Platform VLANs.
MHSA	No	Yes	Yes	DHCP and DAI enabled on Platform VLAN associated with RADIUS I-SIDs.
MHSA	No	Yes	No	DHCP and DAI enabled on all static Platform VLANs.
MHSA	No	No	Yes	DHCP and DAI enabled on received RADIUS VLAN.
MHSA	No	No	No	DHCP and DAI enabled on default VLAN.
Mutiple Host Multiple VLAN (MHMV)	Yes	Yes	Yes	DHCP, DAI, and IPSG enabled on Platform VLAN associated with RADIUS I-SID and Platform VLAN associated with untagged I-SID.
MHMOV	Yes	Yes	No	DHCP, DAI, and IPSG enabled on Platform VLAN associated with untagged I-SID.
MHMOV	Yes	No	Yes	DHCP, DAI, and IPSG enabled on RADIUS VLAN and default VLAN.
MHMOV	Yes	No	No	DHCP, DAI, and IPSG enabled on default VLAN.
MHMOV	No	Yes	Yes	DHCP and DAI enabled on Platform VLAN associated with RADIUS I-SID.

Table 248: Expected Behavior for DHCP Snooping and DAI VSAs (continued)

EAP Operational Mode	IPSG VSA received	Flexi-UNI	RADIUS VLAN/I-SID	DHCP Snooping and DAI VSA received
MH MV	No	Yes	No	DHCP and DAI enabled on Platform VLAN associated with untagged I-SID.
MH MV	No	No	Yes	DHCP and DAI enabled on received RADIUS VLAN.
MH MV	No	No	No	DHCP and DAI enabled on default VLAN.
Note: If no RADIUS VLAN or Platform VLAN is associated with RADIUS I-SID, DHCP Snooping and DAI are enabled on the default VLAN or Platform VLAN associated with static untagged I-SID.				

Expected Behavior of IGMPV3 and IGMPFAST Vendor Specific Attributes

The following table shows the expected behavior of IGMPV3 and IGMPFAST Vendor Specific Attributes for each authentication scenario, depending on Extensible Authentication Protocol (EAP) and the attributes returned by the RADIUS server.

Table 249: Expected Behavior of IGMPV3 and IGMPFAST Vendor Specific Attributes

Mode	Port type	Condition	Action
MH MV	Regular port	VLAN attribute received from RADIUS	Apply IGMP VSA on that VLAN
MH MV	Regular port	No VLAN attribute received	Apply IGMP VSA on the default VLAN
MH MV	Flex-UNI port	I-SID attribute received from RADIUS	Apply IGMP VSA on the associated platform VLAN
MH MV	Flex-UNI port	No I-SID attribute received	Apply IGMP VSA on the platform VLAN associated with the default untagged I-SID
MH SA	Regular port	One or more VLAN attributes received from RADIUS	Apply IGMP VSA on only one VLAN (usually the last one received)
MH SA	Regular port	No VLAN attribute received	Apply IGMP VSA on the default VLAN

Table 249: Expected Behavior of IGMPV3 and IGMPFAST Vendor Specific Attributes (continued)

Mode	Port type	Condition	Action
MHSA	Flex-UNI port	One or more I-SID attributes received from RADIUS	Apply IGMP VSA on all platform VLANs associated with those I-SIDs
MHSA	Flex-UNI port	No I-SID attribute received	Apply IGMP VSA on all VLANs that contain the client port

Extreme-Dynamic-Config Operational Considerations

Consider the following when you use port and VLAN based attributes.

General Considerations

Only attributes that can be configured manually can be configured dynamically using EAP.

Dynamic cleanup is supported. When the last client to authenticate using a dynamic attribute is removed, the following dynamic attributes are also removed:

- Dynamic ARP Inspection (DAI)
- DHCP Snooping
- IGMP Snooping
- IP Source Guard
- Re-authentication

However, the following attributes can only be removed by disabling EAP:


- SLPP Guard
- BPDU Guard
- Traffic Control (Wake on LAN)
- Custom Auto-Negotiation Advertisements

VLAN-Based Features

- DHCP Snooping Option 82 is not supported.
- IGMP Snooping is not supported on a DvR Leaf.
- If multiple client authentication is permitted in MHMV mode, you can apply RADIUS attributes incrementally as subsequent clients authenticate.

If a client authenticates with DHCP Snooping, DAI, and IP Source Guard attributes on the VLAN and a second client attempts to authenticate with the same attributes, the following occurs:

- If the second client uses the same VLAN as the first client, only IP Source Guard applies on the RADIUS configuration port.

- If the second client uses a different VLAN, DHCP Snooping and DAI apply on the VLAN and the IP Source Guard applies on the RADIUS configuration port.
-  **Note**
The following consideration only applies if IP Source Guard is received from the RADIUS server. Otherwise DHCP Snooping and Dynamic ARP Inspection (DAI) is enabled based on the behavior described in [Expected Behavior for DHCP Snooping and DAI Vendor Specific Attributes](#) on page 3121.

If you configure a port in Multiple Host Single Association (MHSA) mode and VLAN based attributes are received from the RADIUS server, features are enabled on the default VLAN and on all VLANs that contain the authentication port.

Port-Based Features

- Configuring Custom Auto-Negotiation Advertisements (CANA) on a port triggers a port bounce, which generates new client authentication.
- Disconnect requests execute a disconnect command for the client. If the PORTBOUNCE attribute is included, then it only performs a port bounce.
- Re-authenticate requests execute a re-authentication of the client. These requests can include all attributes.
- Change-of-Authorization (CoA) requests perform the change of authorized configuration for the supported attributes. If the PORTBOUNCE attribute is included, then it only performs a port bounce.
- On Flex-UNI ports, if the I-SID received from the RADIUS server does not have a platform VLAN associated with it, attributes are not applied. When a platform VLAN is associated with the I-SID, EAP re-authentication is generated to apply the attribute by bouncing a port, bouncing EAP on a port, or by using CoA Re-authenticate.
- IP Source Guard restrictions apply even if the feature is configured on the RADIUS server.
 - Maximum 10 entries per port
 - Maximum 1000 entries per server
 - DHCP Snooping and DAI must be enabled on all VLAN members of the RADIUS configured port.
- If you configure a Guest VLAN on a port and the RADIUS server returns IP Source Guard as a result of EAP or NEAP authentication, then you should manually remove static VLANs from that port. Alternatively, you can enable DHCP Snooping and DAI on static VLANs.
- If you configure a port with multiple platform VLANs and the RADIUS server returns IP Source Guard as a result of EAP/NEAP authentication, then you must manually configure DHCP Snooping and DAI on static platform VLANs.

The re-authentication flag and re-authentication period attributes origin can be either CONFIG or RADIUS. Different origins for re-authentication flag and re-authentication period attributes are valid:

- You can configure the re-authentication flag with or without a time interval in CLI or RADIUS VSA. If you do not specify a time interval when you enable re-authentication on a port from RADIUS, the re-authentication period origin does not change.
- If a RADIUS client specifies the same value as the one that already exists in static configuration through CLI, the origin remains as CONFIG.
- If you enable re-authentication through CLI and you configure a specific period using the *re-authentication-period <0, 60-65535>* command, the origin is CONFIG.

The following message displays to indicate that RADIUS clients use the configuration:

```
WARNING: Setting used by Radius Client. Are you sure you want to
continue? (y/n)?
```

If the re-authentication period attribute was configured with the re-authentication flag through RADIUS VSA, the origin is RADIUS.

When you change the re-authentication period attribute in CLI, the following message displays to indicate that the origin of this parameter is RADIUS.

```
WARNING: Current port reauth period has RADIUS origin. Are you sure
you want to continue? (y/n)?
```

Changing a parameter in CLI that was originally configured using RADIUS, changes the origin to CONFIG.

For more information, see [Extreme-Dynamic-Config](#) on page 3116.

Extreme-Dynamic-Config Examples

The following example enables DHCP Snooping and Dynamic ARP Inspection (DAI) on VLAN 100 and IP Source Guard on authentication port:

```
Auth-Type := Local
User-Password == "000000000001"
Service-Type = Administrative-User
Tunnel-Type = VLAN
Tunnel-Medium-Type = IEEE-802
Tunnel-Private-Group-Id = 100
Extreme-Dynamic-Config = "DHCP Snooping"
Extreme-Dynamic-Config += "DAI"
Extreme-Dynamic-Config += "IPSG"
```

The following example configures the Custom Auto-Negotiation Advertisements (CANAs) to 100Half on authentication port:

```
Auth-Type := Local
User-Password == "000000000008"
Service-Type = Administrative-User
Extreme-Dynamic-Config = "AN-ADVERTISEMENTS:100H"
```

The following example enables BPDU Guard on authentication port:

```
Auth-Type := Local
User-Password == "00000000000006"
Service-Type = Administrative-User
Extreme-Dynamic-Config = "BPDU"
```

The following example enables DHCP Snooping and DAI on the platform VLAN associated with I-SID 10000 and IP Source Guard on authentication port:

```
Auth-Type := Local
User-Password == "00000000000002"
Service-Type = Administrative-User
Fabric-Attach-ISID = "0:10000"
Extreme-Dynamic-Config = "DHCP Snooping"
Extreme-Dynamic-Config += "DAI"
Extreme-Dynamic-Config += "IPSG"
```

The following example enables IGMP Snooping on VLAN 100:

```
Auth-Type := Local
User-Password == "00000000000003"
Service-Type = Administrative-User
Tunnel-Type = VLAN
Tunnel-Medium-Type = IEEE-802
Tunnel-Private-Group-Id = 100
Extreme-Dynamic-Config = "IGMP Snooping"
```

The following example enables IGMP Snooping platform VLAN associated with I-SID 10000:

```
Auth-Type := Local
User-Password == "00000000000004"
Service-Type = Administrative-User
Fabric-Attach-ISID = "0:10000"
Extreme-Dynamic-Config = "IGMP Snooping"
```

The following example enables reauthentication on a port without a configured time interval:

```
Auth-Type := Local
User-Password == "00000000000007"
Service-Type = Administrative-User
Extreme-Dynamic-Config = "REAUTH"
```

The following example enables reauthentication on a port with a configured time interval:

```
Auth-Type := Local
User-Password == "00000000000007"
Service-Type = Administrative-User
Extreme-Dynamic-Config = "REAUTH:1230"
```

The following examples enable SLPP Guard on authentication port:

```
Auth-Type := Local
User-Password == "00000000000007"
Service-Type = Administrative-User
Extreme-Dynamic-Config = "SLPPGUARD"

Auth-Type := Local
User-Password == "00000000000007"
```

```
Service-Type = Administrative-User
Extreme-Dynamic-Config = "SLPPGUARD:12345"
```

The following example enables EAP traffic-control on authentication port:

```
Auth-Type := Local
User-Password == "0000000000005"
Service-Type = Administrative-User
Extreme-Dynamic-Config = "WOL"
```

The following example triggers a PORTBOUNCE on the port that authenticates the user:

```
Auth-Type := Local, User-Password == "0000000000001"
Extreme-Dynamic-Client-Assignments += "create=vlan, pv=300, vni=100300",
Extreme-Dynamic-Config += "PORTBOUNCE"
```

Extreme-Dynamic-Client-Assignments

The Extreme-Dynamic-Client-Assignments Vendor Specific Attribute (VSA) is a RADIUS VSA for dynamic VLAN and Private VLAN (PVLAN) creation. You can also configure VLAN parameters, such as VLAN name, I-SID to VLAN association, I-SID name, I-SID to dynamic VRF association, and IGMP Querier address.

- Name: Extreme-Dynamic-Client-Assignments
- Value: 253
- Type: String
- Vendor: Extreme Networks
- Extreme Networks Vendor ID is 1916

The Extreme-Dynamic-Client-Assignments VSA supports the following VLAN-based features:

- IGMP Snooping
- DHCP Snooping
- Dynamic ARP Inspection (DAI)



Note

You can use Extreme-Dynamic-Config RADIUS VSA to configure the VLAN-based features.

Note the following points regarding VLAN creation with the Extreme-Dynamic-Client-Assignments VSA:

- Platform VLAN manual configuration—This maps the Extreme-Dynamic-Client-Assignments Radius VSA without `create`. Only the options used to define the existing functionality of mapping clients to VLAN on regular ports or `vlan:isid` on flex-uni ports are used—Primary VLAN (PV), Virtual Network Identifier (VNI), and Egress VLAN (EV). The Secondary VLAN (SV), VLAN Name (VN), and VNIN attributes for dynamic VLAN are ignored.
- Regular VLAN configuration—This maps the VSA with the option `create=vlan`. The Multiple Spanning Tree Protocol (MSTP) instance is 0. When platform VLAN is dynamically created, all VLAN parameters are also dynamically applied; a static VLAN

setting is not allowed. Dynamic platform VLANs and dynamic I-SIDs names exist as long as EAP clients reference them. If clients do not use them, then they are deleted and are not saved in the configuration file.

- PVLAN configuration—This maps with the usage of Extreme-Dynamic-Client-Assignments Radius VSA with the `create=pvlan` option. The EAP and Private VLANs on regular ports are not supported. But they are supported on flex-uni ports. The MSTP instance is 0. When you dynamically create a private VLAN, all VLAN parameters are also dynamically created; a static VLAN setting is not allowed. Dynamic PVLANS and dynamic I-SIDs names exist as long as EAP clients reference them. If clients do not use them, then they are deleted and are not saved in the configuration file.

Use the information in the following tables and this string format to create a dynamic VLAN:

```
create=vlan|pvlan|none, pv=Primary VLANID, sv=secondary VLANID,
vni=ISID, ev=EGRESS-VLAN-tag, vn=vlan-name, vnin=isid-name, mvni=ISID,
igmpqaddr=<IPv4 address>
```



Note

If `ev` is missing, it defaults to 0. You can also use `U` or `T` (case-sensitive). When `ev` is set to `U`, it is untagged or 0. When `ev` is set to `T`, it takes the value of `pv` or tagged. If `pv` is not specified, then an error occurs and the VSA is ignored.

Table 250: String Options for Dynamic VLAN Creation

Option	Description
<code>create=vlan pvlan none</code>	<p>If <code>create</code> is missing, the assumption is that a manually created VLAN exists. Note the following examples:</p> <ul style="list-style-type: none"> • <code>create=vlan</code>—dynamically creates a platform VLAN. • <code>create=pvlan</code>—dynamically creates a private VLAN. • <code>create=none</code>—dynamically uses an existing VLAN or PVLAN instead of creating a new one <p>This option is ignored on DvR Leafs.</p> <p>Note: <code>create</code> arguments are case sensitive.</p>
<code>pv=Primary VLANID</code>	<p>The platform VLAN that the client is assigned.</p> <p>This option is valid for any combination of the <code>create</code> command.</p>
<code>sv=Secondary VLANID</code>	<p>This option is only valid for a private VLAN and if the <code>create</code> option is used.</p>

Table 250: String Options for Dynamic VLAN Creation (continued)

Option	Description
vni=ISID	If you did not use <code>create</code> then you can use <code>vni</code> on flex-uni ports with <code>ev</code> to assign a client to a Switched UNI (S-UNI). The <code>vni</code> command also has the role of mapping the dynamic created VLAN to the VNI.
ev=EGRESS-VLAN-tag	Use this option on regular ports to tag or untag the egress for the PV. Use this option on flex-uni ports as <code>c-vid</code> in S-UNI creation.
vn=Vlan name	Valid only if you use <code>create</code> .
vnin=ISID name	Valid only if you use <code>create</code> .
mvni=ISID	Use this option to configure IP Multicast config-lite for Fabric Connect functionality. This option configures the I-SID value for the Layer 3 VSN VLAN and enables SPB Multicast on the VLAN. You can associate <code>mvni=ISID</code> to a dynamic or a static VLAN. For information about IP Multicast config-lite for Fabric Connect, see IP Multicast config-lite for Fabric Connect on page 1952.
igmpqaddr=<IPv4 address>	Use this option if the VLAN interface uses one of the following: <ul style="list-style-type: none"> IGMP Snooping—the switch configures the received IPv4 address as the IGMP Snooping Querier address and enables the IGMP Snooping Querier. Multicast Lite or Routed Multicast (Routed SPB)—the switch configures the received IPv4 address as the IGMP Routed SPB Querier address. The IGMP Querier address applies on the VLAN interface you specify in the <code>pv</code> argument.

Table 251: VSA Equivalency with Radius Attributes

Port Type	RADIUS Attribute	Extreme-Dynamic-Client-Assignments Radius VSA	Comment
Regular port	Tunnel Private GroupID	Without <code>create</code> . <code>pv=Primary VLANID</code>	This adds the port to the primary VLAN; tag is the port tag.
	Egress-VLANID	Without <code>create</code> .	Untagged: <code>ev = 0</code> .

Table 251: VSA Equivalency with Radius Attributes (continued)

Port Type	RADIUS Attribute	Extreme-Dynamic-Client-Assignments Radius VSA	Comment
		pv=Primary VLANID. ev= EGRESS-VLAN- tag	Tagged: ev = pv. This adds the port to the primary VLAN; The VLAN egress tag is dictated by the ev.
	Egress-VLAN-Name	Not Supported	
Flex-Uni ports	FA VLAN:ISID	Without create. vni=[ISID] ev= [EGRESS-VLAN- tag]	Untagged: ev = 0. Tagged: ev != 0. A S-UNI is created, either MAC-based or regular, depending on MHMV/MHSA setting; uses i-sid (vnid) and c-vid (ev) values.
	Egress-VLANID + FA VLAN:I-SID	Supported by the same combination for FA VLAN:ISID	
	Egress-VLANID + Tunnel Private GroupID + autoIsidOffset	Without create. pv=Primary VLANID. ev= EGRESS-VLAN- tag	Untagged: ev = 0. Tagged: ev != 0. A S-UNI is created, either MAC-based or regular, depending on MHMV/MHSA setting; uses auto configured i-sid from pv value and c-vid (ev) values.
	Egress-VLAN-Name + FA VLAN:I-SID >	Not supported	
	Egress-VLAN-Name + Tunnel Private GroupId + autoIsidOffset	Not supported	

The dynamic VLAN is deleted after you disconnect all of the clients across the Extensible Authentication Protocol (EAP) ports. The port is removed when the last client is disconnected and the saved I-SID name is restored.

You cannot delete a static VLAN if EAP ports are assigned to it. However, you can delete the VLAN if you have added EAP FlexUNI ports to it. This flushes all MAC addresses and deletes any Non-EAP (NEAP) sessions. The MAC address is re-learned in I-SID and a new RADIUS authentication can now create a dynamic VLAN. For EAP sessions, the session

moves to the re-authentication state, and the new RADIUS authentication can create a dynamic VLAN.

In Multiple Host Multiple VLAN (MHMV), if there are multiple assignments received for one session, the rule is to process only the last one and ignore the rest.

Operational Considerations

Consider the following when you use the Extreme-Dynamic-Client-Assignments Vendor Specific Attribute (VSA):

- Extreme-Dynamic-Client-Assignments takes precedence over FA-VLAN-ISID, EGRESS-VLAN-ID, EGRESS-VLAN-Name, and RAD_ATTR_TUNNEL_PRI_GROUP_ID.
- The RADIUS packet is limited to 4096 bytes, but can carry any VSA combination that does not exceed 4096 bytes. The Extreme-Dynamic-Client-Assignments VSA has a variable length and combinations can scale up to 94 VSAs in a packet. To reach 94 VSAs in a packet, each VSA should be about 40 bytes long.
- You cannot manually modify dynamic VLANs.
- You cannot use this VSA to dynamically create VLANs on a DvR Leaf node.
- Static VLAN configuration is only supported for IP Multicast config-lite for Fabric Connect functionality.
- MultiLink Trunk (MLT) and Split MultiLink Trunk (SMLT) ports are not supported.
- RADIUS Change-of-Authorization (CoA) and RADIUS Disconnect functions are not supported.

EAP-Port-Priority



Important

Configure these attributes only if you require dynamic port priority.

Port priority (vendor-specific) attributes:

- Vendor ID: value 562 (Nortel)



Note

Vendor ID 1916 (Extreme Networks) does not support port priority.

- Attribute Number: value 1, Port Priority
- Attribute Value: value 0 (zero) to 6 (this value indicates the port priority value assigned to the specified user)

The following list provides the switch Port Priority frame format:

- vendor specific type = 26
- length = 12
- vendor-id = 562
- string = vendor type = 1 + vendor length = 6 + attribute specific = priority

The following figure shows an example of the port priority frame format.

26	12	562	01	06	(0...6)
----	----	-----	----	----	---------

Figure 206: Port priority frame format

IETF Attributes

VLAN RADIUS Attributes

To control the tagging for traffic egressing a port, configure the preferred tagging option and one of the following attributes on the RADIUS server:

- EGRESS-VLAN-ID
- EGRESS-VLAN-NAME

VLAN attributes are defined by RFC 4675.

EGRESS-VLAN-ID

Functionality is combined with other RADIUS attributes and EAPoL configurations to create untagged S-UNIs.

This attribute supports the following combinations:

- untagged EGRESS-VLAN-ID + VLAN:ISID -> VLAN from S-UNI becomes untagged
- untagged EGRESS-VLAN-ID + vid[NEW_RADIUS_VID] + autolsidOffset -> create S-UNI for I-SID (VLAN+autolsidOffset)

Examples

From the RADIUS attributes, for untagged VLAN 15. The VLAN from S-UNI becomes untagged:

```
Egress-VLANID = "838860815" (untagged vlan 15)
Fabric-Attach-ISID += "15:23513"
```

The following example results in the value 0:1015 for creating an untagged S-UNI.

From the RADIUS attributes, for untagged VLAN 15 and RADIUS-assigned VLAN. Creates S-UNI for I-SID (VLAN+autolsidOffset):

```
Egress-VLANID = "838860815"
Tunnel-Private-Group-Id = "15"
```

From the EAP configuration:

```
eapol auto-isid-offset 1000
eapol auto-isid-offset enable
```

Considerations

Consider the following when using this attribute:

- Any combinations with EGRESS-VLAN-ID for a tagged VLAN are ignored.
- When EGRESS-VLAN-ID and VLAN:ISID binding attributes are used, the VLAN must have the same value.
- The RADIUS server can provide more than one VLAN:ISID binding but one must match the VLAN used in EGRESS-VLAN-ID to create an untagged S-UNI.
- Only one EGRESS-VLAN-ID attribute is used. If more than one is found in the packet, only the first one is used.
- If the untagged S-UNI is derived from EGRESS-VLAN-ID + VLAN-ID + VLAN attribute + auto-isid-offset, the packet must not contain any S-UNI bindings.

EGRESS-VLAN-NAME

Functionality is combined with other RADIUS attributes and EAPoL configurations to create untagged S-UNIs.

This attribute supports the following combinations:

- EGRESS-VLAN-ID + VLAN:I-SID
- EGRESS-VLAN-ID + VLAN + autolsidOffset

This attribute includes the following information:

- Egress-VLAN-Name = "1VLAN_NAME" for tagged VLAN
- Egress-VLAN-Name = "2VLAN_NAME" for untagged VLAN
- VLAN_NAME is the name of the VLAN on the switch

Example

The following example creates an untagged Switched UNI with I-SID 1000. VLAN 10 is named "Dummy_VLAN".

```
b0adaa41b830 Cleartext-Password := "b0adaa41b830"
Service-Type = Framed-User,
Framed-Protocol = PPP,
Auth-Type := Accept,
Fabric-Attach-ISID = "10:1000",
Egress-VLAN-name = "2Dummy_VLAN"
```

RADIUS Tunnel Attributes

The switch uses the RADIUS tunnel attributes to place a port into a particular VLAN to support dynamic VLAN switching based on authentication. The server must send these attributes together in the same RADIUS packet. If one attribute is missing, the switch ignores the others.

The RADIUS server indicates the desired VLAN by including the tunnel attribute within the Access-Accept message. RADIUS uses the following tunnel attributes for VLAN membership:

- RAD_ATTR_TUNNEL_TYPE(64)

Tunnel-Type: value 13, Tunnel-Type-VLAN

- RAD_ATTR_TUNNEL_MEDIUM_TYPE(65)

Tunnel-Medium-Type: value 6, Tunnel-Medium-Type-802

- RAD_ATTR_TUNNEL_PRI_GROUP_ID

Tunnel-Private-Group-ID: ASCII value 1-4094 (this value identifies the specified VLAN)

Tunnel attributes are defined by RFC 2868.

**Important**

Configure these attributes only if you require Dynamic VLAN membership.

The VLAN ID is 12 bits, uses a value from $\langle 1-4094 \rangle$, and is encoded as a string.

In addition, you can configure the RADIUS server to send a vendor-specific attribute (VSA) to configure port priority. You can assign the switch Supplicant port a QoS value from 0 to 6.

For more information, see [EAP-Port-Priority](#) on page 3132.



Representational State Transfer Configuration Protocol (RESTCONF)

[Representational State Transfer Configuration Protocol \(RESTCONF\) Fundamentals](#) on page 3136

[RESTCONF configuration using CLI](#) on page 3139

[RESTCONF Configuration using EDM](#) on page 3143

[Use Representational State Transfer Configuration Protocol \(RESTCONF\) to Configure a Switch](#) on page 3145

[RESTCONF Operational Behavior Examples](#) on page 3147

Table 252: Representational State Transfer Configuration Protocol (RESTCONF) product support

Feature	Product	Release introduced
Representational State Transfer Configuration Protocol (RESTCONF)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Representational State Transfer Configuration Protocol (RESTCONF) Fundamentals

Representational State Transfer Configuration Protocol (RESTCONF) is a next generation northbound interface that provides an additional way to configure and monitor the switch. RESTCONF is an HTTP-based protocol that provides a programmatic interface to access data defined in a YANG model using the datastore concepts defined in NETCONF. RESTCONF uses a client-server model. The server acts as an entry point to a datastore, a conceptual place to store and access information. Clients use HTTP or HTTPS to interface with the server to configure and monitor devices.

RESTCONF Client and Server

A typical RESTCONF interaction consists of an HTTP/HTTPS request sent by a RESTCONF client and an HTTP/HTTPS response sent by the server. The HTTP/HTTPS request and response contain a required set of expected HTTP headers and can also contain a request or response message body. The message body is encoded in JSON.

An HTTP request consists of the HTTP method (such as GET or POST) identifier, resource identifier, HTTP protocol version, HTTP headers, and HTTP body. The HTTP resource identifier is the string that identifies a service or resource that the server makes available to the client. The RESTCONF request contains the Universal Resource Identifier or URI which starts with `/rest/restconf/data/` or `/rest/restconf/operations/`.

YANG Model

YANG is the data modeling language used for modeling configuration and state data for manipulation by using remote procedure calls (RPCs). The RESTCONF interface is generated with YANG Data Model. The YANG model is based on Open config model, which is a non vendor specific model that captures the key components found in multiple vendor solutions. RESTCONF is described by the Internet Engineering Task Force (IETF) in RFC 8040.

RESTCONF Authentication

RESTCONF uses the CLI user account and supports both local and remote authentication. Local authentication uses the local CLI user account while remote authentication can use either a RADIUS or TACACS+ server.

You can only use a CLI account with the RWA access level.

With RADIUS or TACACS+ enabled, if the remote server is not available, authentication falls back to local authentication and uses the local CLI user on the switch.

When the RESTCONF client posts for authentication, the HTTP server validates the login user name and password if you have not enabled CLI remote authentication. If the remote server is not reachable, the HTTP server uses the local user for login validation.

For HTTPS access to the RESTCONF server, you must enable TLS and install a certificate. For the certificate, first you need to create a leaf certificate in .pem format and convert the private key to PKCS8 format. Then you need to create a .pem container that contains both the leaf certificate and the PCKS8 private key. The container will be used in the RESTCONF configuration and will have the following format:

```
-----BEGIN CERTIFICATE-----  
-----END CERTIFICATE-----  
-----BEGIN PRIVATE KEY-----  
-----END PRIVATE KEY-----
```

To install the certificate, copy your certificate and private key into the container between the appropriate BEGIN and END lines.

RESTCONF APIs

You can access the RESTCONF API documentation on your switch using the following URL:

`http(s)://<IP>:<tcp-port>/apps/restconfdoc/`

Replace <IP> with the management IP address of your switch and <tcp-port> with the TCP port configured for RESTCONF. For example, `http://192.0.2.16:8080/apps/restconfdoc/`.

The on-switch URL works only if you enable the RESTCONF feature on the switch.

You can also access the RESTCONF API documentation online through the Developer Center (<https://www.extremenetworks.com/support/documentation-api/>).

Server Support

The RESTCONF server in the network operating system (NOS) supports the following actions:

HTTP Action	VOSS Instrumentation
GET	Corresponds to SHOW
POST	Corresponds to SET for creation
PATCH	Corresponds to SET for modification
DELETE	Corresponds to SET for deletion

The following table details modules supported by RESTCONF:

Modules	
OpenConfig	Network Service
OpenConfig	Relay Agent (DHCP support)
OpenConfig	Interfaces Port: POE, port attributes, such as auto-sense, default-vlan-id, flex-uni, qos, untag-port-default-vlan
OpenConfig	Interfaces LAG: attributes, such as flex-uni
OpenConfig	Platform: ports, CPU, fans, power supply, optical devices - GET operations only
OpenConfig	Network Instance: VLAN interface - VRF association, CVLAN I-SID, IS-IS redistribute direct, IPVPN, I-SID, and IP DHCP relay forward path)
OpenConfig	STP: STP global information and port interface bpduguard state, RSTP global and port level information, MSTP global, MST instance level state - GET operations only
OpenConfig	System (aaa)

Modules	
OpenConfig	LLDP
OpenConfig	VLAN

The RESTCONF feature is disabled by default. The RESTCONF server uses the same management IP address as the other applications and TCP port. The default TCP port that RESTCONF server listens to is port 8080. The TCP port delivers the message to the HTTP server for RESTCONF.

RESTCONF configuration using CLI

This section provides procedures to configure Representational State Transfer Configuration Protocol (RESTCONF) using the command line interface (CLI).

Enable the RESTCONF Server

Before You Begin

Run the **show application restconf conflict-ifname** command to see if any conflict in interface names exist. To enable RESTCONF, the interface names (VLAN name, MLT name, and Port interface name) must be unique.

Run the **show application restconf invalid-name mlt** and **show application restconf invalid-name vlan** commands to see if any MLT or VLAN names contain special characters. To enable RESTCONF, VLAN and MLT names cannot contain special characters other than underscore (_), en dash (-), and colon (:).

About This Task

Use the following procedure to enable the RESTCONF server.

Procedure

1. Enter Application Configuration mode:

```
enable
```

```
configure terminal
```

```
application
```

2. Enable the RESTCONF server:

```
restconf enable
```



Note

If the interface names (VLAN, MLT, and Port) are not unique, or if VLAN or MLT names contain prohibited special characters, an error occurs indicating that you cannot enable RESTCONF. You must change the interface names before you enable RESTCONF.

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CTRL/Z.
Switch:1(config)#application
Switch:1(config-app)#restconf enable
```

Configure HTTPS Access to the RESTCONF Server

Before You Begin

Ensure that you have the certificate file in the `/intflash` directory on the switch.

About This Task

By default, the RESTCONF server uses HTTP. If you need to use HTTPS, generate a certificate file and transfer the certificate file to the `/intflash` directory on the switch.

For more information on generating certificate files, see [Manage an SSL Certificate](#) on page 3357.

Procedure

1. Enter Application Configuration mode:
`enable`
`configure terminal`
`application`
2. If RESTCONF is enabled, disable RESTCONF:
`no restconf enable`
3. Install the certificate file for the RESTCONF server:
`restconf install-cert-file WORD<1-128>`
4. Enable HTTPS:
`restconf tls`
5. Enable RESTCONF:
`restconf enable`

Example

```
Switch:1>enable
Switch:1# configure terminal
Enter configuration commands, one per line. End with CTRL/Z.
Switch:1(config)#application
Switch:1(config-app)#no restconf enable
Switch:1(config-app)#restconf install-cert-file /intflash/.cert/restconf-cert.pem
Switch:1(config-app)#restconf tls
Switch:1(config-app)#restconf enable
```

Variable Definitions

The following table defines parameters for the **restconf** command.

Variable	Value
<code>enable</code>	Enables the RESTCONF Server.
<code>install-cert-file WORD<1-128></code>	Installs the certificate file for the RESTCONF server.
<code>tcp-port <1024-49151></code>	Configures the RESTCONF Server TCP port number.
<code>tls</code>	Enables TLS for the RESTCONF server. The default is disabled.
<code>trap-notification</code>	Enables trap notification.

Modifying the RESTCONF Server Settings

About This Task

Use this procedure to modify the RESTCONF server settings.



Note

These steps are considered optional and RESTCONF can operate with the default configuration of these values.

Procedure

1. Enter Application Configuration mode:
`enable`

`configure terminal`

`application`
2. Disable trap notification when the RESTCONF server is not available:
`no restconf trap-notification`
3. Modify the TCP port number for the RESTCONF server:
 - a. Disable RESTCONF: `no restconf enable`
 - b. Change the value of the TCP port: `restconf tcp-port <1024-49151>`
 - c. Enable RESTCONF: `restconf enable`
4. Disable TLS for the RESTCONF server:
 - a. Disable RESTCONF: `no restconf enable`
 - b. Disable TLS: `no restconf tls`
 - c. Enable RESTCONF: `restconf enable`

Variable Definitions

Use the data in the following table to modify the RESTCONF server settings.

Variable	Value
<i>enable</i>	Enables or disables the RESTCONF server. The default is disabled.
<i>tcp-port</i> <1024-49151>	Specifies the TCP port to use for the RESTCONF server. The default is 8080.
<i>trap-notification</i>	Enables or disables trap notification when the RESTCONF server is not available. The default is enabled.

Showing the RESTCONF Configuration Information

About This Task

Use this procedure to show the RESTCONF configuration information and operation status.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the RESTCONF configuration:

```
show application restconf
```

Example

```
Switch:1>show application restconf
=====
                                RESTCONF Info
=====
Admin State           : true
TCP Port              : 8080
Certificate File Status : install
TLS Enable            : false
Trap Notification     : true
Oper State            : up
Web Server Version    : 1.0.1.11
RESTCONF Server Version : 1.0.1.39
=====
```

Show Conflicting Interface Name Information

About This Task

To enable RESTCONF, the interface name (VLAN name, MLT name, and Port interface name) must be unique. Use this procedure to display conflicting interface name information.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Show the RESTCONF conflicting interface name information:

```
show application restconf conflict-ifname
```

Example

```
Switch:1>show application restconf conflict-ifname
-----
Conflicting Interface IfName - Port, VLAN Name and MLT Name
-----
Mlt 1 name is same as Vlan 1001 name - "Interface-1"
Mlt 2 name is same as Vlan 1002 name - "VLAN-1002"
Vlan 1003 name is Mlt 1 Default Name - "MLT-1"
-----
Total Conflict Count: 3
```

What to Do Next

If a conflict exists, change the conflicting interface name to a unique name.

Show Special Characters in VLAN or MLT Names

About This Task

To enable RESTCONF, VLAN and MLT names cannot contain special characters other than underscore (_), en dash (-), and colon (:). Use this procedure to display VLAN or MLT names that contain prohibited special characters.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Show the RESTCONF VLAN or MLT names that contain prohibited special characters:

```
show application restconf invalid-name vlan
```

```
show application restconf invalid-name mlt
```

Example

```
Switch:1>show application restconf invalid-name mlt
-----
Invalid MLT names - Only "-", "_" and ":" special characters are allowed
-----
Mlt 3 name has special characters - "gigi#g"
Mlt 4 name has special characters - "my%mlt"
Mlt 5 name has special characters - "isa.text"
-----
Total Invalid Names Count: 3
```

What to Do Next

If any of the names contain prohibited special characters, change the names to remove the special characters.

RESTCONF Configuration using EDM

This section contains procedures for configuring RESTCONF with Enterprise Device Manager (EDM).

Configure the RESTCONF Server

About This Task

To configure the server, you must enable RESTCONF. RESTCONF is disabled by default.

After RESTCONF is enabled, you must disable RESTCONF to modify some of the RESTCONF parameters.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability**.
2. Select **RESTCONF**.
3. Select the **RESTCONF** tab.
4. Select the **GlobalEnable** check box to enable the RESTCONFserver.
5. Configure optional parameters as required.
6. Select **Apply**.

RESTCONF Field Descriptions

Use the data in the following table to use the **RESTCONF** tab.

Name	Description
GlobalEnable	Enables or disables the RESTCONF server. The default is disabled (cleared).
TcpPort	Specifies the TCP port to use for the RESTCONF server. The default is 8080. The RESTCONF status must be disabled before you can modify this field.
TlsEnable	Enables or disables TLS/SSL if you require HTTPS access to the RESTCONF server. The default is disabled. The RESTCONF status must be disabled before you can modify this field.
CertificateFilename	If HTTPS access is required, specifies the file name and path of the TLS/SSL certificate. The certificate file must be in the <code>/intflash</code> directory on the switch.
CertificateAction	Installs or uninstalls the TLS/SSL certificate file. It also shows the current status of the certificate installation.
NotificationEnable	Enables or disables trap notification when the RESTCONF server is not available. The default is enabled.
OperStatus	Shows the operational status of the RESTCONF server.
WebServerVersion	Shows the RESTCONF web server version that is running on the server.
RestConfServerVersion	Shows the RESTCONF server version that is running on the server.

Use Representational State Transfer Configuration Protocol (RESTCONF) to Configure a Switch

The documentation does not include information about how to use RESTCONF clients. This example documents some common tasks using Python.

For additional details, see [Extreme API with Python](#).

Before You Begin

Configure the RESTCONF server on the switch.

Procedure

1. Import classes, define variables, and prepare the session object:

```
#!/usr/bin/env python
import sys
import json
import requests
from requests import Request, Session
from requests.auth import HTTPBasicAuth
from requests.packages.urllib3.exceptions import InsecureRequestWarning
requests.packages.urllib3.disable_warnings(InsecureRequestWarning)
#####
Host      = '192.0.2.1'
TcpPort   = '8080'
UserName  = 'rwa'
Password  = 'rwa'
LoginUrl  = 'https://%s:%s/auth/token' % (Host, TcpPort)
QueryUrl  = 'https://%s:%s/rest/restconf/data/' % (Host, TcpPort)
vlan      = 123
#####
session   = Session()
session.verify = False
session.timeout = 5

session.headers.update({
    'Accept':          'application/json',
    'Accept-Encoding': 'gzip, deflate, br',
    'Connection':      'keep-alive',
    'Cache-Control':   'no-cache',
    'Pragma':          'no-cache', })
```

2. Learn the authentication token:



Note

The token expires after 24 hours.

```
body = '{"username": "%s", "password" : "%s" }' % (UserName, Password)

response = session.put( LoginUrl, body )

if response.status_code != 200:
    print "ERROR: Login failed"
    sys.exit(1)
else:
    session.headers.update({
        'X-Auth-Token': response.headers['X-Auth-Token']
    })
    print "INFO: login passed"
```

3. Query all VLANs:

```
response = session.get( QueryUrl + 'openconfig-vlan:vlan' )
if response.status_code != 200:
    print 'ERROR: can't fetch VLANs'
```

4. Query specific VLANs:

```
response = session.get( QueryUrl + 'openconfig-vlan:vlan=vlan=%s' % vlan )

if response.status_code != 200:
    print 'INFO: VLAN %s doesn't exists' % vlan
else:
    print 'INFO: VLAN %s exists' % vlan
```

5. Access data:

```
inbound_data = json.loads( response.text )

for vlan in inbound_data['openconfig-vlan:vlan']['vlan']:
    print 'VLAN: %s [%s]' % ( vlan['state']['name'], vlan['vlan-id'] )
```

6. Present data:

```
inbound_data = json.loads( response.text )

for dataVlan in inbound_data[ dataObject ]['vlan']:
    print ''
    print 'VLAN: ' + dataVlan['state']['name'] + '[' + dataVlan['vlan-id'] + ']'
    interfaces = ''
    if 'members' in dataVlan :
        for interface in dataVlan['members']['member ']:
            interfaces = interfaces + interface['interface-ref']['state']['interface']
    + ','

    print interfaces
```

7. Add a VLAN:

```
dataObject = 'openconfig-vlan:vlan'

data = { "openconfig-vlan:vlan": [
    {
        "config":
        { "name": "Test-VLAN",
          "extreme-mod-oc-vlan:stg-id": 1,
          "vlan-id": vlan
        }
    }
]

response = session.post( QueryUrl + dataObject, json=data )

if response.status_code != 201:
    print "ERROR: add VLAN %s failed" % vlan
else:
    print "INFO: VLAN %s added " % vlan
```

8. Update a VLAN:

```
dataObject = 'openconfig-vlan:vlan=vlan=%s/config' % vlan

data = { "openconfig-vlan:config":
    { "name": "New-VLAN-Name" }
}

response = session.patch( QueryUrl + dataObject )

if response.status_code != 204:
```



```
    print "ERROR: update VLAN %s fails" % vlan
else:
    print "INFO: VLAN %s updated" % vlan
```

9. Delete a VLAN:

```
dataObject = 'openconfig-vlan:vlan=vlan=%s' % vlan

response = session.delete( QueryUrl + dataObject )

if response.status_code != 204:
    print "ERROR: delete VLAN %s fails" % vlan
else:
    print "INFO: VLAN %s deleted" % vlan
```

RESTCONF Operational Behavior Examples

This section provides an example of RESTCONF operational behavior.

In the following example, RESTCONF sends a request that involves multiple commands where one command operation fails and other command operations proceed.

Example:

```
POST http://{vossip}:8080/rest/restconf/data/extreme-network-service:network-services
{
  "extreme-network-service:network-services": [
    {
      "network-service": {
        "id": "990099",
        "type": "elanTransparent",
        "name": "isid-990099",
        "interfaces": {
          "interface": [
            {
              "name": "1/8"
            }
          ]
        }
      }
    }
  ]
}

{
  "ietf-restconf:errors": {
    "error": [
      {
        "error-message": "Unable to push data - Manual configuration not allowed on Auto-Sense port",
        "error-tag": "operation-failed",
        "error-type": "protocol"
      }
    ]
  }
}
```

The following sample output displays the configured I-SID and the created ELAN-transparent based service. The port does not display because the command operation failed.

```
Switch:1(config)#i-sid name 990099 isid-990099
```

```
Switch:1(config)#i-sid 990099 elan-transparent
Switch:1(elan-tp:990099)#port 1/8
```

```
Switch:1>show i-sid
```

```
*****
```

=====										
ISID	ISID		PORT	MLT	ORIGIN					ISID
ID	TYPE	VLANID	INTERFACES	INTERFACES						NAME

99999	ELAN_TR	N/A	-	-	C	---	-	---	-	isid-99999
990099	ELAN_TR	N/A	-	-	C	---	-	---	-	isid-990099



RIP

- [RIP fundamentals](#) on page 3150
- [RIPng Fundamentals](#) on page 3154
- [RIP configuration using CLI](#) on page 3156
- [RIPng Configuration using CLI](#) on page 3167
- [RIP configuration using EDM](#) on page 3173

[RIPng Configuration using EDM](#) on page 3181

Table 253: Routing Information Protocol product support

Feature	Product	Release introduced
Routing Information Protocol (RIP)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 254: RIPng product support

Feature	Product	Release introduced
RIPng	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

RIP fundamentals

Use the information in these sections to help you understand the Routing Information Protocol (RIP). For more information about the Border Gateway Protocol (BGP), see [BGP](#) on page 536.

Routing Information Protocol

In routed environments, routers communicate with one another to track available routes. Routers can dynamically learn about available routes using the RIP. The switch software implements standard RIP to exchange IP route information with other routers.

RIP uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. Each router advertises routing information by sending a routing information update every 30 seconds (one interval). If RIP does not receive information about a network for 180 seconds, the metric associated with the network rises to infinity (U); that is, the metric resets to 16, which means the network becomes unreachable. If RIP does not receive information about a network for 120 seconds, it removes the network from the routing table.

RIP is a distance vector protocol. The vector is the network number and next hop, and the distance is the cost associated with the network number. RIP identifies network reachability based on cost, and cost is defined as hop count. One hop is the distance from one router to the next. This cost or hop count is the metric (see the following figure).




Figure 207: Hop count or metric in RIP

RIP version 1 (RIPv1) advertises default class addresses without subnet masking. RIP version 2 (RIPv2) advertises class addresses explicitly, based on the subnet mask.

The switch supports RIPv2, which supports variable length subnet masks (VLSM) and triggered router updates. RIPv2 sends mask information. If RIP does not receive information about a network for 90 seconds, the metric associated with the network rises to infinity (U); that is, the metric resets to 16, which means the network becomes unreachable. If RIP does not receive information about a network for 180 seconds (six update intervals), it removes the network from the routing table. You can change the default timers by configuring the RIP interface timeout timer and the holddown timer.

A directly connected network has a metric of zero. An unreachable network has a metric of 16. Therefore, the highest metric between two networks can be 15 hops or 15 routers.

RIP Updates

Choose one of three options for the switch to receive RIP updates:

- `rip1OrRip2`—accepts RIPv1 or RIPv2 updates
- `rip1`—accepts RIPv1 updates only
- `rip2`—accepts RIPv2 updates only

The following table describes the four RIP send modes that the switch supports. You can configure RIP send modes on all router interfaces.

Table 255: RIP send modes

Send mode	Description	Result
<code>rip1Compatible</code>	Broadcasts RIPv2 updates using RFC1058 route consumption rules. This is the default mode.	<ul style="list-style-type: none"> • Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 192.0.2.255) • RIP update is formed as a RIP-2 update, including network mask • RIP version = 2
<code>ripVersion1</code>	Broadcasts RIP updates that are compliant with RFC 1058	<ul style="list-style-type: none"> • Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff • Destination IP is a broadcast for the network (for example, 198.0.2.255) • RIP update is formed as a RIP-1 update, no network mask included • RIP version = 1
<code>ripVersion2</code>	Broadcasts multicast RIPv2 updates	<ul style="list-style-type: none"> • Destination MAC is a multicast, 01-00-5e-00-00-09 • Destination IP is the RIP-2 multicast address, 224.0.0.9 • RIP update is formed as a RIP-2 update including network mask • RIP version = 2
<code>doNotSend</code>	Does not send RIP updates on the interface	None

RIP Supply and Listen

The following table indicates the relationship between switch action and the RIP supply and listen configuration.

Table 256: RIP supply and listen configuration and switch action

RIP supply configuration		RIP listen configuration		Switch action
Supply	Default supply	Listen	Default listen	
Disabled (false)	Disabled (false)			Sends no RIP updates.
Enabled (true)	Disabled (false)			Sends RIP updates except the default.
Disabled (false)	Disabled (false)			Sends only the default (default route must exist in routing table).
Enabled (true)	Enabled (true)			Sends RIP updates including the default route (if it exists).
		Disabled (false)	Disabled (false)	Does not listen to RIP updates.
		Enabled (true)	Disabled (false)	Listens to all RIP updates except the default.
		Disabled (false)	Enabled (true)	Listens only to the default.
		Enabled (true)	Enabled (true)	Listens to RIP updates including the default route (if it exists).

RIP and Route Redistribution

Redistribution imports routes from one protocol to another. Redistribution sends route updates for a protocol-based route through another protocol. For example, if RIP routes exist in a router and they must travel through a BGP network, configure redistribution of RIP routes through BGP. Redistribution sends RIP routes to a router that uses BGP.

You can redistribute routes

- on an interface basis
- on a global basis between protocols on a single VRF instance (intraVRF)
- between the same or different protocols on different VRF instances (interVRF)

To configure interface-based redistribution, configure a route policy, and then apply it to the interface. Configure the match parameter to the protocol from which to learn routes.

You can redistribute routes on a global basis, rather than for every interface. The switch adds support for global RIP redistribution. Use the **ip rip redistribute**

command to accomplish the (intraVRF) redistribution of routes through RIP, so that RIP redistribution occurs globally on all RIP-enabled interfaces. This redistribution does not require a route policy, but you can use one for more control.

If you configure redistribution globally and on an interface, redistribution through the route policy takes precedence.

You can redistribute routes from a protocol in one VRF to RIP in another VRF. You can use a route policy for redistribution control. If you enable route redistribution between VRF instances, ensure that IP addresses do not overlap.

RIPng Fundamentals

Routing Information Protocol next generation (RIPng) allows routers to exchange information for computing routes through an IPv6-based network. You should implement RIPng only on routers. IPv6 provides neighbor router information required by RIPng protocol to function as intended. A RIPng router is assumed to have interfaces in several networks and the protocol relies primarily on the metric of each network to compute routes using the distance vector algorithm.

RIP identifies network reachability based on cost, and cost is defined as hop count. One hop is the distance from one router to the next. This cost, or hop count, is the metric.

RIPng-enabled routers use UDP port 521 (the RIPng port) to exchange routing information. RIPng responds to a request by sending a message to the port from which the request originates. Specific queries can be sent from ports other than the RIPng port, but they must be directed to the RIPng port on the target machine.

Each router advertises routing information by sending an update every 30 seconds (one interval). If RIPng does not receive information about a network for 180 seconds, the metric associated with the network rises to infinity (U); that is, the metric resets to 16, which means the network becomes unreachable. If RIP does not receive information about a network for 120 seconds, it removes the network from the routing table.



Note

These time interval values are default values which are configurable by the user.

Each router that implements RIPng contains a routing table. This table contains one entry for every destination that is reachable throughout the system operating RIPng. At a minimum, each routing table entry contains the following information:

- The IPv6 prefix of the destination.
- A metric that represents the total cost of getting a datagram from the router to that destination. The metric is the sum of the costs of traversing the networks to arrive at the destination.
- The IPv6 address of the next router in the path to the destination (the next hop). The next-hop IPv6 address is a linklocal address.
- The VLAN or brouter port on which the RIPng routes were learned.
- The age of the RIPng route.

RIPng protocol implementation is specified in IETF document RFC 2080.

RIPng messages and packet format

RIPng-enabled routers use UDP port 521 (the RIPng port) to send and receive datagrams.

The following figure shows the RIPng packet format:




Figure 208: RIPng packet format

A RIPng packet header consists of the following components:

- **command:** Specifies the purpose of the message.
- **version:** The version of RIPng.

Originate Default route

Generally you use a default route when it is not convenient to list every possible network in RIPng updates, and one or more routers in the system are able to handle traffic to networks that RIPng does not explicitly list.

RIPng is enabled with the default route only option. When you enable default route only on an interface, it suppresses all other routes in the update sent for the interface, and advertises only the default route.

Timers

RIPng states four different timer intervals for protocol operation:

- **Update timer:** The RIPng process sends a complete routing table to each neighboring router every 30 seconds. To prevent collisions on broadcast networks, the process adds an offset value to the timer.
- **Timeout time interval:** This is a 180 second time interval associated with every route. If the time interval expires, the metric for this route updates to the value of infinity.

(16) and the route is no longer valid. However, the routing table retains the value for another 120 seconds.

- **Garbage collection time interval:** After the timeout time interval expires and the route becomes invalid, it remains in the routing table until the garbage collection time interval expires. The garbage collection time interval is 120 seconds. Until the garbage collection time interval expires all updates sent by this router include the invalid route. When the garbage collection timer expires, the process removes the route from the routing table.
- **Triggered update time interval:** The triggered update time interval is set to a random value between 1 and 5 seconds after a triggered update is sent. A single update is sent even if multiple triggered updates occur before the timer expires.

Configuration of timers or time intervals is supported only at the CLI/SNMP/EDM level. Configuration of timers or time intervals is not supported at the interface/port level.

F

RIP configuration using CLI

Use Routing Information Protocol (RIP) to perform dynamic routing within an autonomous system. This section describes how you use the command line interface (CLI) to configure and manage RIP.

Configure RIP Globally

Configure RIP parameters on the switch so you can control RIP behavior on the system.

Before You Begin

- You configure RIP on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode and the prefix **ip rip**. The VRF must have an RP Trigger of RIP. Not all parameters are configurable on non0 VRFs.

About This Task

All router interfaces that use RIP use the RIP global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

You can configure RIP on interfaces while RIP is globally disabled. This way, you can configure all interfaces before you enable RIP for the switch.

Procedure

1. Enter RIP Router Configuration mode:

```
enable  
  
configure terminal  
  
router rip
```
2. Define the default-import-metric for the switch:

```
default-metric <0-15>
```

3. (Optional) Configure one or more timer values:
`timers basic timeout <15-259200> [holddown <0-360>] [update <1-360>]`
4. Enable RIP on an IP network:
`network {A.B.C.D}`
5. Exit to Global Configuration mode:
`exit`
6. After the configuration is complete, enable RIP globally:
`router rip enable`
- 7.
8. Check that your configuration is correct:
`show ip rip [vrf WORD<1-16>] [vrfids WORD<0-512>]`

Example

Define the default-import-metric as 1, the timeout interval as 180, the holddown time as 120, and the update time as 30. Enable RIP on an IP network, and ensure your configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#router rip
Switch:1(config-rip)#default-metric 1
Switch:1(config-rip)#timers basic timeout 180 holddown 120 update 30
Switch:1(config-rip)#network 192.0.2.11
Switch:1(config-rip)#exit
Switch:1(config)#router rip enable
Switch:1(config)#show ip rip
=====
                        RIP Global - GlobalRouter
=====
Default Import Metric : 1
      HoldDown Time   : 120
        Queries       : 0
          Rip          : Enabled
      Route Changes    : 0
    Timeout Interval   : 180
        Update Time    : 30
```

Variable definitions

The following table defines parameters for the RIP commands.

Variable	Value
<code>default-metric <0-15></code>	Configures the value of default import metric to import a route into a RIP domain. To announce OSPF internal routes into RIP domain, if the policy does not specify a metric value, the default is used. For OSPF external routes, the external cost is used. The default is 8.
<code>domain <0-39321></code>	Specifies the RIP domain from 0–39321. The default is 0.
<code>holddown <0-360></code>	Configures the RIP hold-down timer value, the length of time (in seconds) that RIP continues to advertise a network after it determines that the network is unreachable. The default is 120.

Variable	Value
<i>network</i> {A.B.C.D}	Enables RIP on an IP network.
<i>timeout</i> <15-259200>	Configures the RIP timeout interval. The default is 180.
<i>update</i> <1-360>	Configures the RIP update timer. The update time is the time interval, in seconds, between RIP updates. The default is 30.

The following table defines parameters for the **show ip rip** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies a VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies a range of VRF IDs.

Configure RIP on an Interface

Configure RIP on Ethernet ports and VLANs so that they can participate in RIP routing.

Before You Begin

- Assign an IP address to the port or VLAN.
- Configure RIP and enable it globally.
- Configure in and out policies.

About This Task

RIP does not operate on a port or VLAN until you enable it both globally and on the port or VLAN.

To configure RIP on a VRF instance for a port or VLAN, you configure RIP on the port or VLAN, and then associate the port or VLAN with the VRF.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Define the cost:

```
ip rip cost <1-15>
```

3. Specify an in policy for filtering inbound RIP packets:

```
ip rip in-policy WORD<0-64>
```

4. Specify an out policy for filtering outbound RIP packets:
`ip rip out-policy WORD<0-64>`
5. Enable RIP:
`ip rip enable`
6. Specify the send mode:
`ip rip send version <notsend|rip1|rip1comp|rip2>`
7. Specify the receive mode:
`ip rip receive version <rip1|rip2|rip1orrip2>`
8. Change other RIP parameters from their default values as required.

Example

The following configuration example shows how to configure the switch (labeled R1) to operate only in RIP version 2 mode.




Figure 209: Configuration example-RIPv2 only

Enable RIPv2 send mode on VLAN 2:

```
Switch:1(config-if)# ip rip send version rip2
```

Enable RIPv2 receive mode on VLAN 2:

```
Switch:1(config-if)# ip rip receive version rip2
```

Repeat these commands on VLAN 3 and the port interfaces.

Variable Definitions

The following table defines parameters for the **ip rip** command.

Variable	Value
<i>advertise-when-down enable</i>	Enables or disables AdvertiseWhenDown. If enabled, RIP advertises the network on this interface as up, even if the port is down. The default is disabled. If you configure a port with no link and enable advertise-when-down, it does not advertise the route until the port is active. RIP advertises the route even when the link is down. To disable advertising based on link status, you must disable this parameter.
<i>auto-aggregation enable</i>	Enables or disables automatic route aggregation on the port. If enabled, the switch automatically aggregates routes to their natural mask when an interface in a different class network advertises them. The default is disable.
<i>cost <1-15></i>	Configures the RIP cost for this port (link).
<i>default-listen enable</i>	Enables DefaultListen. The switch accepts the default route learned through RIP on this interface. The default is disabled.
<i>default-supply enable</i>	Enables DefaultSupply. If enabled, this interface must advertise a default route. The default is false. RIP advertises the default route only if it exists in the routing table.
<i>enable</i>	Enables RIP routing on the port.
<i>holddown <0-360></i>	Configures the RIP holddown timer value, the length of time (in seconds) that RIP continues to advertise a network after it determines that the network is unreachable. The default is 120.
<i>in-policy WORD<0-64></i>	Configures the policy name for inbound filtering on this RIP interface. This policy determines whether to learn a route on this interface and specifies the parameters of the route when RIP adds it to the routing table.
<i>listen enable</i>	Specifies that the routing switch learns RIP routes through this interface. If enabled, the switch listens for a default route without listening for all routes. The default is enable.
<i>out-policy WORD<0-64></i>	Configures the policy name for outbound filtering on this RIP interface. This policy determines whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement. <i>WORD<0-64></i> is a string of length 0–64 characters.

Variable	Value
<code>poison enable</code>	<p>Enables Poison Reverse. If you disable Poison Reverse (no <code>poison enable</code>), Split Horizon is enabled.</p> <p>By default, Split Horizon is enabled. If you enable Split Horizon, the interface does not advertise IP routes learned from an immediate neighbor back to the neighbor. If you enable Poison Reverse, the RIP updates sent to a neighbor from which a route is learned are poisoned with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network. These mechanisms prevent routing loops.</p>
<code>port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>receive version <rip1 rip2 rip1orrip2></code>	Indicates which RIP update version to accept on this interface. The default is rip1orrip2.
<code>send version <notsend rip1 rip1comp rip2></code>	Indicates which RIP update version the router sends from this interface. ripVersion1 implies sending RIP updates that comply with RFC1058. rip1comp implies broadcasting RIP2 updates using RFC1058 route subassumption rules. The default is rip1Compatible.
<code>supply enable</code>	Specifies that the switch advertises RIP routes through the port. The default is enable.
<code>timeout <15-259200></code>	Configures the RIP timeout interval in seconds. The default is 180.
<code>triggered enable</code>	Enables automatic triggered updates for RIP.

Configure Route Redistribution to RIP

Configure a redistribute entry to announce certain routes into the RIP domain, including static routes, direct routes, RIP, Open Shortest Path First (OSPF), IS-IS, or Border Gateway Protocol (BGP). Optionally, use a route policy to control the redistribution of routes.

Before You Begin

- Enable RIP globally.
- Configure a route policy.

Procedure

1. Enter RIP Router Configuration mode:

```
enable

configure terminal

router rip
```
2. Create the redistribution instance:

```
redistribute <bgp|direct|isis|ospf|rip|static> [vrf-src WORD<1-16>]
```
3. Apply a route policy, if required:

```
redistribute <bgp|direct|isis|ospf|rip|static> route-map WORD<0-64>
[vrf-src WORD<1-16>]
```
4. Configure other parameters.
5. Enable the redistribution:

```
redistribute <bgp|direct|isis|ospf|rip|static> enable [vrf-src
WORD<1-16>]
```
6. Ensure that the configuration is correct:

```
show ip rip redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```
7. Exit to Global Configuration mode:

```
exit
```
8. Apply the redistribution:

```
ip rip apply redistribute <bgp|direct|isis|ospf|rip|static> [vrf
WORD<1-16>] [vrf-src WORD<1-16>]
```

Example

Create the redistribution instance, apply a route policy, enable the redistribution, and apply the redistribution.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router rip
Switch:1(config-rip)#redistribute rip
Switch:1(config-rip)#redistribute rip route-map test1
Switch:1(config-rip)#redistribute rip enable
Switch:1(config-rip)#exit
Switch:1(config)#ip rip apply redistribute rip
```

Variable definitions

The following table defines parameters for the **redistribute** command.

Variable	Value
<i>metric</i> <0-65535>	Configures the metric to apply to redistributed routes.
<i>route-map</i> WORD<0-64>	Configures the route policy to apply to redistributed routes.

Variable	Value
<i>[vrf-src WORD<1-16>]</i>	Specifies the optional source VRF instance. You can use this variable with the other command variables.
<i>WORD<0-32></i>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

The following table defines parameters for the **show ip rip redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF instance.
<i>vrfids WORD<1-512></i>	Specifies a range of VRF IDs.

The following table defines parameters for the **ip rip apply redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF instance.
<i>vrf-src WORD<1-16></i>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<i><bgp direct isis ospf rip static></i>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Configure InterVRF Route Redistribution for RIP

Configure a redistribute entry to announce certain routes into the RIP domain, including static routes, direct routes, RIP, OSPF, IS-IS, or BGP. Use a route policy to control the redistribution of routes.

Before You Begin

- Enable RIP globally.
- Configure a route policy.
- Configure the VRFs.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the redistribution instance:

```
ip rip redistribute <bgp|direct|isis|ospf|rip|static>
```

3. Apply a route policy, if required:

```
ip rip redistribute <bgp|direct|isis|ospf|rip|static> route-map  
WORD<0-64> [vrf-src WORD<1-16>]
```

4. Configure other parameters.

5. Enable the redistribution:

```
ip rip redistribute <bgp|direct|isis|ospf|rip|static> enable [vrf-src  
WORD<1-16>]
```

6. Ensure that the configuration is correct:

```
show ip rip redistribute [vrf WORD<1-16>] [vrfids WORD<1-512>]
```

7. Exit to Global Configuration mode:

```
exit
```

8. Apply the redistribution:

```
ip rip apply redistribute <bgp|direct|isis|ospf|rip|static> [vrf  
WORD<1-16>] [vrf-src WORD<1-16>]
```

Example

Create the redistribution instance, apply a route policy, enable the redistribution, and apply the redistribution.

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1#router vrf red  
Switch:1(router-vrf)#ip rip redistribute ospf  
Switch:1(router-vrf)#ip rip redistribute ospf route-map test1  
Switch:1(router-vrf)#ip rip redistribute ospf enable  
Switch:1(router-vrf)#exit  
Switch:1(config)#ip rip apply redistribute ospf
```

Variable definitions

The following table defines parameters for the **ip rip redistribute** *<bgp|isis|ospf|static|direct|rip>* command.

Variable	Value
<i><bgp direct isis ospf rip static></i>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.
<i>vrf-src WORD<1-16></i>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<i>metric <0-65535></i>	Configures the metric to apply to redistributed routes.
<i>route-map WORD<0-64></i>	Configures the route policy to apply to redistributed routes.

The following table defines parameters for the **show ip rip redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF instance.
<i>vrfids WORD<1-512></i>	Specifies a range of VRF IDs.

The following table defines parameters for the **ip rip apply redistribute** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF instance.
<i>vrf-src WORD<1-16></i>	Specifies the source VRF instance. You do not need to configure this parameter for redistribution within the same VRF.
<i><bgp direct isis ospf rip static></i>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, rip, ospf, or static.

Force a RIP Update for a Port or VLAN

Force RIP to update the routing table so that the port or VLAN uses the latest routing information.

About This Task

If you perform this procedure, you also update the tables for all VRFs associated with the port or VLAN.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable the triggered-update flag:

```
ip rip triggered enable
```



Note

You can enable this flag in either the GigabitEthernet or VLAN Interface Configuration mode. However, you can update the RIP routes in the GigabitEthernet Interface Configuration mode only.

3. Update the routing table:

```
action triggerRipUpdate
```

Example

Update the routing table:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface vlan 1  
Switch:1(config-if)#ip rip triggered enable
```

Display the RIP Redistribution Configuration Information

Displays the RIP redistribution configuration information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the RIP redistribution configuration information:

```
show ip rip redistribute [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

View the RIP redistribution configuration information:

```
Switch(config-ospf)#show ip rip redistribute
```

```
=====
RIP Redistribute List - GlobalRouter
=====
```

SRC-VRF	SRC	MET	ENABLE	RPOLICY

GlobalRouter	ISIS	0	FALSE	

Variable definitions

The following table defines parameters for the **show ip rip redistribute** command or the **show ipv6 rip redistribute** command.

Variable	Value
<i>vrf</i> WORD<1-16>	Specifies a VRF by name.
<i>vrfids</i> WORD<0-512>	Specifies a range of VRF IDs.

RIPng Configuration using CLI

Use the procedures in this section to configure RIPng using CLI.

Configure RIPng Globally

Configure RIPng parameters on the router so you can control RIPng behavior on the system.

Before You Begin

You can configure RIPng only on a global router. You cannot configure RIPng on a VRF instance.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable RIPng globally:
`router rip ipv6-enable`

Configure RIPng on an Interface

Configure RIPng on Ethernet ports and VLANs so that they can participate in RIPng routing.

Before You Begin

- Assign an IP address to the port or VLAN.
- Configure RIPng and enable it globally.

About This Task

RIPng does not operate on a port or VLAN until you enable it both globally and on the port or VLAN.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Create a RIPng interface:

```
ipv6 rip
```

3. Enable the RIPng interface:

```
ipv6 rip enable
```

4. Verify the operational status of the RIPng interface:

```
show ipv6 rip interface
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 22
Switch:1(config-if)#ipv6 rip
Switch:1(config-if)#ipv6 rip enable
Switch:1(config-if)#show ipv6 rip interface
```

Total RIPng interfaces: 2

RIPng Interface - GlobalRouter						
IFINDX	COST	POISON STATUS	SEND DEFAULT	ADMIN STATUS	OPER STATUS	
257 (2/2)	2	disable	disable	enable	enable	
2070 (22)	5	disable	disable	enable	disable	

2 out of 2 Total Num of RIPng interfaces displayed

Variable Definitions

The following table defines parameters for the **ipv6 rip** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This variable applies only to VLAN interfaces, not ports.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configure RIPng Custom Values

Configure custom values for RIPng parameters to replace default values.

Before You Begin

- Configure RIPng and enable it globally.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable RIPng poison:

```
ipv6 rip poison enable
```

3. Specify the RIPng cost:
`ipv6 rip cost <1-15 Cost>`
4. Access router RIPng configuration mode:
`router rip`
5. Specify the RIPng holddown timer value:
`ipv6 timers basic holddown <0-360>`
6. Specify the RIPng timeout timer value:
`ipv6 timers basic timeout <15-259200>`
7. Specify the RIPng update timer value:
`ipv6 timers basic update <1-360>`
8. Specify the default route metric value:
`ipv6 default-information metric <1-15>`
9. Enable default information globally:
`ipv6 default-information enable`
10. Ensure the configuration is correct:
`show ipv6 rip`

Example

Configure custom values for RIPng.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#router rip
Switch:1(config-rip)#ipv6 default-information metric 1
Switch:1(config-rip)#ipv6 default-information enable
Switch:1(config-rip)#ipv6 timers basic update 30
Switch:1(config-rip)#ipv6 timers basic timeout 180
Switch:1(config-rip)#ipv6 timers basic holddown 120
Switch:1(config-rip)#router rip ipv6-enable
Switch:1(config)#show ipv6 rip
```

```
=====
                        RIPng Global - GlobalRouter
=====

      Rip : Enabled
      HoldDown Time : 120
      Timeout Interval : 180
      Update Time : 30
      Default Info Metric : 1
      Default Info State : Enabled
      Default Import Metric : 1
```


Variable definitions

Use the data in the following table to use the **ipv6 rip poison**, **ipv6 default-information** and **ipv6 timers basic** commands.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This variable applies only to VLAN interfaces, not ports.
port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
poison enable	Enables Poison Reverse. If you disable Poison Reverse (no poison enable), Split Horizon is enabled. By default, Split Horizon is enabled. If you enable Split Horizon, the interface does not advertise IP routes learned from an immediate neighbor back to the neighbor. If you enable Poison Reverse, the RIP updates sent to a neighbor from which a route is learned are poisoned with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network. These mechanisms prevent routing loops.
<1-15 Cost>	Configures the RIPng cost for this port (link).
holddown <0-360>	Configures the RIPng holddown timer value, the length of time (in seconds) that RIPng continues to advertise a network after it determines that the network is unreachable. The default is 120.
timeout <15-259200>	Configures the RIPng timeout interval. The default is 180.
update <1-360>	Configure the RIPng update timer. The update time is the time interval between RIPng updates.
default-information <1-15>	Configure the default route metric value.

Configure RIPng Route Distribution

Configure a redistribute entry to announce certain routes into the RIPng domain, including static routes, direct routes, Open Shortest Path First (OSPFv3), IS-IS, or Border Gateway Protocol (BGP+).

Before You Begin

- Enable RIPng globally.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Access router RIP configuration mode:

```
router rip
```
3. Enable the redistribution:

```
ipv6 redistribute {direct|isis|static|ospf|bgp} enable
```
4. Ensure the configuration is correct:

```
show ipv6 rip redistribute
```

Example

Enable the redistribution instance.

```
Switch:1#enable
Switch:1#configure terminal
Switch:1(config)#router rip
Switch:1(config-rip)#ipv6 redistribute bgp enable
Switch:1(config-rip)#ipv6 redistribute direct enable
Switch:1(config-rip)#ipv6 redistribute isis enable
Switch:1(config-rip)#ipv6 redistribute ospf enable
Switch:1(config-rip)#ipv6 redistribute static enable
Switch:1(config-rip)#show ipv6 rip redistribute
```

```
=====
                        RIPng Redistribute List
=====
```

direct	: enabled
static	: enabled
ospf	: enabled
bgp	: enabled
isis	: enabled

```
-----
```

Variable definitions

Use the data in the following table to use the **ipv6 redistribute** command.

Variable	Value
<bgp direct isis ospf static>	Specifies the type of routes to redistribute (the protocol source). Valid options are bgp, isis, direct, ospf, or static.

RIP configuration using EDM

Use Routing Information Protocol (RIP) to perform dynamic routing within an autonomous system. This section describes how you use Enterprise Device Manager (EDM) to configure and manage RIP.

Configure RIP Globally

Configure RIP global parameters on the switch so you can control RIP behavior on the system.

Before You Begin

- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

About This Task

All router interfaces that use RIP use the RIP global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

You can configure RIP on interfaces while RIP is globally disabled. This way, you can configure all interfaces before you enable RIP for the switch.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RIP**.
3. Select the **Globals** tab.
4. Select the **enable** option button.
5. Configure other global RIP parameters as required.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Operation	Enables or disables RIP on all interfaces. The default is disabled.
UpdateTime	Specifies the time interval between RIP updates for all interfaces. The default is 30 seconds, and the range is 0–360.
RouteChanges	Specifies the number of route changes RIP made to the IP route database. RouteChanges does not include the refresh of a route age.
Queries	Specifies the number of responses sent to RIP queries received from other systems.

Name	Description
HoldDownTime	Configures the length of time that RIP continues to advertise a network after the network is unreachable. The range is 0–360 seconds. The default is 120 seconds.
TimeOutInterval	Configures the RIP timeout interval. The range is 15–259200 seconds. The default is 180 seconds.
DeflImportMetric	Configures the default import metric used to import a route into a RIP domain. To announce OSPF internal routes into a RIP domain, if the policy does not specify a metric, you must use the default import metric. OSPF external routes use the external cost. The range is 0–15 and the default is 8.

View RIP status

View RIP status.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RIP**.
3. Select the **Status** tab.

Status Field Description

Use the data in the following table to use the **RIP Status** tab.

Name	Description
Address	Specifies the IP address of the router interface.
RcvBadPackets	Specifies the number of RIP response packets received by the RIP process which were subsequently discarded; for example, version 0 packet, or an unknown command type.
RcvBadRoutes	Specifies the number of routes, in valid RIP packets, that are ignored; for example, unknown address family, or invalid metric.
SentUpdates	Specifies the number of triggered RIP updates sent on this interface, that do not include full updates sent containing new information.

Configure RIP Interface Compatibility

Configure RIP parameters on an interface so you can control RIP behavior on the interface. You can specify the RIP version to use on interfaces that you configure to send (supply) or receive (listen to) RIP updates.

Before You Begin

- Configure a routing interface (either a router port or a virtual routing interface).
- Assign an IP address to the interface.

- Enable RIP globally.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

About This Task

On an interface, RIP does not operate until you enable it globally and on the interface.

Although visible, the switch does not support the AuthType and AuthKey parameters.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RIP**.
3. Select the **Interface** tab.
4. Double-click the **Send** value to edit it, and then select the RIP version datagrams the router sends.
5. Double-click the **Receive** value to edit it, and then select the RIP version datagrams for which the router listens.
6. Select **Apply**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Address	Specifies the IP address of the router interface.
Domain	Specifies the value inserted into the routing domain parameter of all RIP packets sent on this interface.
AuthType	Specifies the type of authentication to use on this interface.
AuthKey	Specifies the authentication key whenever AuthType is not noAuthentication.
Send	<p>Specifies the update version the router sends on this interface:</p> <ul style="list-style-type: none"> • DoNotSend—no RIP updates sent on this interface • ripVersion1—RIP updates compliant with RFC1058 • rip1Compatible—broadcast RIPv2 updates using RFC1058 route subassumption rules • ripVersion2—multicast RIPv2 updates <p>The default is rip1compatible.</p>
Receive	<p>Indicates which versions of RIP updates to accept:</p> <ul style="list-style-type: none"> • rip1 • rip2 • rip1OrRip2 <p>The default is rip1OrRip2. Rip2 and rip1OrRip2 imply receipt of multicast packets.</p>

Configure RIP on an Interface

Configure RIP parameters to control and optimize RIP routing for the interface.

Before You Begin

- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RIP**.
3. Select the **Interface Advance** tab.
4. Double-click a RIP parameter to edit it, as required.
5. Select **Apply**.

Interface Advance Field Descriptions

Use the data in the following table to use the RIP **Interface Advance** tab.

Name	Description
Address	Shows the address of the entry in the IP RIP interface table.
Interface	Indicates the index of the RIP interface.
Enable	Shows if the RIP interface is enabled or disabled.
Supply	Enables (true) or disables (false) the ability to send RIP updates on this interface.
Listen	Configures whether the switch learns routes on this interface.
Poison	Configures whether to advertise RIP routes learned from a neighbor back to the neighbor. If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, RIP poisons the RIP updates, sent to the neighbor from which a route is learned, with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network.
DefaultSupply	Enables (true) or disables (false) an advertisement of a default route on this interface. This command takes effect only if a default route exists in the routing table.
DefaultListen	Enables (true) or disables (false) the switch to accept the default route learned through RIP on this interface. The default is disabled. Enable DefaultListen to add a default route to the route table if another route advertises it.
TriggeredUpdate	Enables (true) or disables (false) the switch to send RIP updates from this interface.

Name	Description
AutoAggregate	Enables (true) or disables (false) automatic route aggregation on this interface. If enabled, the switch automatically aggregates routes to their natural mask when an interface advertises them. The default is disabled.
InPolicy	Determines if RIP can learn routes on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if RIP advertises a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The range is 1–15. The default is 1.

Configure RIP on a Port

Configure RIP on a port so that the port can participate in RIP routing.

Before You Begin

- Assign an IP address to the port.
- Configure RIP and enable it globally.

Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

- Enable RIP on the interface.

About This Task

On an interface, RIP does not operate until you enable it globally and on the interface.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **RIP** tab.
5. Configure the RIP parameters as required.
6. Select **Apply**.

RIP Field Descriptions

Use the data in the following table to use the **RIP** tab.

Name	Description
Enable	Enables or disables RIP on the port.
Supply	Specifies that the routing switch advertises RIP routes through the interface. The default is enable.

Name	Description
Listen	Specifies that the routing switch learns RIP routes through this interface. The default is enable.
Poison	If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, the RIP update sent to a neighbor from which a route is learned is poisoned with a metric of 16. In this manner, the route entry is not passed to the neighbor, because 16 is infinity in terms of hops on a network. The default is disable.
DefaultSupply	Enables or disables DefaultSupply. Enable DefaultSupply if a default route exists in the routing table. The default is false. RIP advertises the default route only if it exists in the routing table.
DefaultListen	Enables or disables DefaultListen. Enable DefaultListen if this interface must learn a default route after another router that connects to the interface advertises it. The default is false (disabled). Enable DefaultListen to add a default route to the route table if another router advertises it.
TriggeredUpdateEnable	Enables or disables triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Enables or disables RIP automatic aggregation. RIPv2 automatically aggregates routes to their natural mask. You can enable automatic aggregation only in RIPv2 mode or RIPv1 compatibility mode. The default is false.
AdvertiseWhenDown	Enables or disables AdvertiseWhenDown. If true, RIP advertises the network on this interface as up, even if the port is down. The default is false. If you configure a port with no link and enable AdvertiseWhenDown, the port does not advertise the route until the port is active. RIP advertises the route even when the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
InPolicy	Determines whether the RIP can learn a route on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if this interface advertises a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The default is 1, and the range is 1–15.
HolddownTime	Configures the length of time that RIP continues to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds
TimeOutInterval	Configures the RIP timeout interval in seconds. The range is 15–259200 seconds. The default is 180 seconds.

Configure RIP on a VLAN

Configure RIP on a VLAN so that the VLAN acts as a routed VLAN (a virtual router).

Before You Begin

- Configure the VLAN.
- Assign an IP address to the VLAN.
- Enable RIP globally.
- Enable RIP on the interface.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **RIP** tab.
7. Configure the VLAN RIP parameters as required.
8. Select **Apply**.

RIP Field Descriptions

Use the data in the following table to use the **RIP** tab.

Name	Description
Enable	Enables or disables RIP on the VLAN.
Supply	Specifies that the routing switch advertises RIP routes through the interface. The default is enable.
Listen	Specifies that the routing switch learns RIP routes through this interface. The default is enable.
Poison	If disabled, the interface invokes Split Horizon and does not advertise IP routes learned from an immediate neighbor back to the neighbor. If enabled, the RIP update sent to a neighbor from which a route is learned is poisoned with a metric of 16. In this manner, the route entry is not passed to the neighbor, because 16 is infinity in terms of hops on a network. The default is disable.
DefaultSupply	Enables or disables DefaultSupply. Enable DefaultSupply if a default route exists in the routing table. The default is false. RIP advertises the default route only if it exists in the routing table.

Name	Description
DefaultListen	Enables or disables DefaultListen. Enable DefaultListen if this interface must learn a default route after another router that connects to the interface advertises it. The default is false (disabled). Enable DefaultListen to add a default route to the route table if another router advertises it.
TriggeredUpdateEnable	Enables or disables triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Enables or disables RIP automatic aggregation. RIPv2 automatically aggregates routes to their natural mask. You can enable automatic aggregation only in RIPv2 mode or RIPv1 compatibility mode. The default is false.
AdvertiseWhenDown	Enables or disables AdvertiseWhenDown. If true, RIP advertises the network on this interface as up, even if the interface is down. The default is false. If you configure a VLAN with no link and enable AdvertiseWhenDown, the VLAN does not advertise the route until the VLAN is active. RIP advertises the route even when the link is down. To disable advertising based on link-states, disable AdvertiseWhenDown.
InPolicy	Determines whether the RIP can learn a route on this interface. This variable also specifies the parameters of the route when RIP adds it to the routing table.
OutPolicy	Determines if this interface advertises a route from the routing table. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. The default is 1, and the range is 1–15.
HolddownTime	Configures the length of time that RIP continues to advertise a network after determining it is unreachable. The range is 0–360 seconds. The default is 120 seconds
TimeOutInterval	Configures the RIP timeout interval in seconds. The range is 15–259200 seconds. The default is 180 seconds.

Configure Route Redistribution to RIP

Configure a redistribute entry to announce routes of a certain source protocol type into the RIP domain, for example, static, RIP, or direct. Use a route policy to control the redistribution of routes.

Before You Begin

- Enable RIP globally.
- Configure a route policy.
- Change the VRF instance as required to configure RIP on a specific VRF instance. The VRF must have an RP trigger of RIP. Not all parameters are configurable on non-default VRFs.

About This Task



Important

Changing the RIP redistribute context is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. As a best practice, change default preferences for a RIP redistribute context, you must do so before you enable the protocols.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RIP**.
3. Select the **Redistribute** tab.
4. Select **Insert**.
5. Configure the source of the routes to redistribute.
6. Select **enable**.
7. Select the route policy to apply to redistributed routes.
8. Configure a metric value.
9. Select **Insert**.

Redistribute Field Descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrfId	Specifies the destination VRF instance. You cannot configure this variable.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.
SrcVrfId	Specifies the source VRF instance.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) a RIP redistribute entry for a specified source type.
RoutePolicy	Configures the route policy (by name) that redistributes external routes from a specified source into an RIP domain. Select the ellipsis (...) and choose from the list in the Route Policy dialog box.
Metric	Configures the RIP route redistribution metric for basic redistribution. The value can be in the range 0–65535. A value of 0 indicates to use the original cost of the route.

RIPng Configuration using EDM

Use the procedures in this section to configure RIPng using EDM.

Configure RIPng Globally

Configure RIPng global parameters on the switch so you can control RIPng behavior on the system.

About This Task

All router interfaces that use RIPng use the RIPng global parameters. Both brouter ports and VLAN virtual routing interfaces use the same RIPng global parameters.

You can configure RIPng on interfaces while RIPng is globally disabled. This way, you can configure all interfaces before you enable RIPng for the switch

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 RIPng**.
3. Select the **Globals** tab.
4. Select **enable**.
5. Configure other global RIPng parameters as required.
6. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AdminState	Enables or disables RIPng globally. The default is disabled.
UpdateTime	Specifies the time interval between RIPng updates for all interfaces. The default is 30 seconds, and the range is 1–360.
GlobalHoldDownTime	Configures the length of time that RIPng continues to advertise a network after the network is unreachable. The range is 0–360 seconds. The default is 120 seconds.
GlobalTimeOutInterval	Configures the RIPng timeout interval. The range is 15–259200 seconds. The default is 180 seconds.
DefaultInfoMetric	RIPng default-information metric.
DefaultInfoState	Default-information enable or disable at the global level
DefaultImportMetric	Specifies the RIPng default import metric.

Configure an IPv6 RIPng Interface

Configure RIPng parameters on an interface so you can control RIPng behavior on the interface.

Before You Begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.

- Enable RIPng globally.

About This Task

RIPng does not operate on an interface until you enable it globally and on the interface.

You can also configure an IPv6 RIPng interface for a brouter port by selecting **Device Physical View**, selecting a port, and following the **Edit > Port > IPv6** navigation path. You can configure an IPv6 RIPng interface for a VLAN through the **VLAN > VLANs > Basic > IPv6** navigation path. This procedure uses the main IPv6 RIPng navigation path where you can create both types of interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 RIPng**.
3. Select the **Interfaces** tab.
4. Select **Insert**.
5. In **IfIndex**, type a value to identify the IPv6 interface.
6. In **RipAdminStatus**, select **enable**.
7. Configure other parameters as required.
8. Select **Insert**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	RIPng interface index.
RipAdminStatus	Enable or disable RIPng on an interface.
DefaultInfoState	Enable or disable default information at the interface level.
Cost	Specifies the RIPng metric cost.
Poison	Enable or disable poison reverse on an RIPng interface.
RipOperStatus	Enable or disable the RIPng operational state on an interface.

Configure an IPv6 RIPng VLAN Interface

Configure RIPng parameters on a VLAN interface so you can control RIPng behavior on the interface.

Before You Begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIPng globally.

About This Task

RIPng does not operate on an interface until you enable it globally and on the interface.

You can also configure an IPv6 RIPng interface for a brouter port by selecting **Device Physical View**, selecting a port, and following the **Edit > Port > IPv6** navigation path.

Procedure

1. In the navigation pane, expand **Configuration >VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a row, and then select **IPv6**.
5. Select **Insert**.
6. Configure other parameters, as required.
7. Select **Insert**.
8. Select **Apply**.

IPv6 Interfaces VLAN Field Descriptions

Use the data in the following table to use the **IPv6 Interfaces** tab.

Name	Description
Interface	Specifies the port or VLAN.
Identifier	Shows the IPv6 address interface identifiers. This value is a binary string of up to 8 octets in network byte-order.
IdentifierLength	Shows the length of the identifier, in bits.
Descr	Specifies a description of the interface. The network management system also configures this string.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. This value corresponds to the lower 12 bits of the IEEE 802.1Q VLAN tag.
Type	Shows the interface type.
ReasmMaxSize(MTU)	Specifies the maximum size of the MTU of this IPv6 interface. This value must be the same for all the IP addresses defined on this interface. The default value is 1500.
PhysAddress	Specifies the media-dependent physical address. For Ethernet, this is a MAC address.
AdminStatus	Specifies if IPv6 is active on this interface. The default is false (disabled).
OperStatus	Specifies the current operational status of the interface.

Name	Description
ReachableTime	Specifies the time, in milliseconds, that a neighbor is considered reachable after receiving a reachability confirmation. The default is 30000.
RetransmitTimer	Specifies the time, in milliseconds, between retransmissions of neighbor solicitation messages to a neighbor when resolving the address, or when probing the reachability of a neighbor. The default is 1000.
CurHopLimit	Specifies the current hop limit field sent in router advertisements from this interface. The value must be the current diameter of the Internet. A value of zero indicates that the advertisement does not specify a value for CurHopLimit. The default is 64.
MulticastAdminStatus	The option to select MulticastAdminStatus is disabled. You cannot configure the administrative status for multicast in this context.
MacOffset	Requests a particular MAC for an IPv6 VLAN. You can specify a MAC offset when you configure IPv6 on a VLAN, or the system can assign a MAC address from within the available range. Different hardware platforms support different MAC offset ranges.
ForwardingEnabled	Indicates whether IPv6 forwarding is enabled. The default is enabled.
RSMLTEnable	Shows whether RSMLT is enabled on the interface. The default value is disabled (false).

Configure an IPv6 RIPng Brouter Port Interface

Configure RIPng parameters on an interface so you can control RIPng behavior on the interface.

Before You Begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIPng globally and on the interface.

About This Task

RIPng does not operate on an interface until you enable it globally and on the interface.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > EditPort**.
3. Select **IPv6**.
4. Select the **RIPng** tab.
5. In **RipAdminStatus**, select **enable**.

6. Configure other parameters as required.
7. Select **Apply**.

RIPng Field Descriptions

Use the data in the following table to use the **RIPng** tab.

Name	Description
RipAdminStatus	Enable or disable RIPng on an interface.
DefaultInfoState	Enable or disable default information at the interface level. The default is disable.
Cost	Specifies the RIPng metric cost. The default is 1.
Poison	Enable or disable poison reverse on an RIPng interface. The default is disable.
RipOperStatus	Shows the RIPng operational state on an interface.

Configure Route Redistribution to RIPng

Before You Begin

- Enable RIP globally.
- Configure a route policy.

About This Task

Configure a redistribute entry to announce routes of a certain source protocol type into the RIPng domain, for example, static, RIP, or direct. Use a route policy to control the redistribution of routes.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 RIPng**.
3. Select the **Redistribute** tab.
4. Double-click the value in the **Enable** column that corresponds with the source protocol type you want to enable or disable.
5. Select **enable** or **disable**.
6. Select **Apply**.

Redistribute Field Descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrfiled	Specifies the destination VRF ID used in the redistribution.
Protocol	Specifies the dynamic routing protocol that receives the external routing information.

Name	Description
SrcVrflid	Specifies the source VRF ID used in the redistribution.
RouteSource	Specifies the route source protocol for the redistribution entry.
Enable	Enables (or disables) a RIPng redistribute entry for a specified source type.

View Statistics for RIPng Interfaces

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6 RIPng**.
3. Select the **Stats** tab.

Stats Field Descriptions

Use the data in the following table to use the **Stats** tab.

Name	Description
IfIndex	Shows the unique value to identify an IPv6 interface.
RcvBadPackets	The number of RIPng response packets received by the RIP process that were subsequently discarded for any reason (examples: a version 0 packet or an unknown command type).
RcvBadRoutes	The number of routes, in valid RIPng packets, that were ignored for any reason (examples: unknown address family or invalid metric).
SentUpdates	The number of triggered RIPng updates actually sent on this interface.
RcvUpdates	The number of triggered RIPng updates actually received on this interface. This explicitly does not include full updates received containing new information.

Graph IPv6 RIPng Statistics

Use the following procedure to graph RIPng statistics for monitoring RIPng behavior on the interface.

Before You Begin

- Configure a routing interface (either a brouter port or a virtual routing interface).
- Assign an IP address to the interface.
- Enable RIPng globally and on the interface.

About This Task

RIPng does not operate on an interface until you enable it globally and on the interface.

Procedure

1. In the navigation pane, expand **Configuration >Ipv6** .
2. Select **IPv6 RIPng**.
3. Select the **Stats** tab.
4. Select an interface row.
5. Select **Graph**.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.

Name	Description
RcvBadPackets	The number of RIPng response packets received by the RIP process that were subsequently discarded for any reason (examples: a version 0 packet or an unknown command type).
RcvBadRoutes	The number of routes, in valid RIPng packets, that were ignored for any reason (examples: unknown address family or invalid metric).
SentUpdates	The number of triggered RIPng updates actually sent on this interface.
RcvUpdates	The number of triggered RIPng updates actually received on this interface. This explicitly does not include full updates received containing new information.
AbsoluteValue	Displays the counter value.
Cumulative	Displays the total value since you opened the Stats tab.
Average/Sec	Displays the average value for each second.
Minimum/Sec	Displays the minimum value for each second.
Maximum/Sec	Displays the maximum value for each second.
Last Val/Sec	Displays the last value for each second.



Remote Monitoring

[Remote Monitoring](#) on page 3190

[RMON 2](#) on page 3194

[RMON Configuration Using CLI](#) on page 3198

[RMON Configuration Using EDM](#) on page 3211

[RMON Alarm Variables](#) on page 3235

Table 257: Remote Monitoring product support

Feature	Product	Release introduced
Remote Monitoring 1 (RMON1) for Layer 1 and Layer 2	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Remote Monitoring 2 (RMON2) for network and application layer protocols	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Remote Monitoring

Remote Monitoring (RMON) is a management information base (MIB) or a group of management objects that you use to obtain or configure values using the Simple Network Management Protocol (SNMP). Use CLI, or EDM, to globally enable RMON on the system. After you globally enable RMON, you enable monitoring for individual devices on a port-by-port basis.

RMON1 is the original version of the protocol, which collects information for OSI Layer 1 and Layer 2 in Ethernet networks. RMON1 provides traffic statistics at the MAC layer, and provides statistics on Ethernet segments for packets and bytes received and transmitted.

You can use RMON1 to:

- Configure alarms for user-defined events.
- Collect Ethernet statistics.
- Log events.
- Send traps for events.

Within EDM, you can configure RMON1 alarms that relate to specific events or variables. You can also specify events associated with alarms to trap or log-and-trap. In turn, the system traps or logs tripped alarms.

You can view all RMON1 information using CLI or EDM. Alternatively, you can use any management application that supports SNMP traps to view RMON1 trap information.

This section describes RMON1 alarms, RMON1 history, RMON1 events, and RMON1 statistics.

RMON1 Alarms

You can configure alarms to alert you if the value of a variable goes out of range. You can define RMON1 alarms on any MIB variable that resolves to an integer value. You cannot use string variables (such as system description) as alarm variables.

You can use RMON1 alarm to monitor anything that has a MIB OID associated with it and a valid instance.

All alarms share the following characteristics:

- A defined upper and lower threshold value.
- A corresponding rising and falling event.
- An alarm interval or polling period.

After you activate alarms, you can:

- View the activity in a log and/or a trap.
- Create a script directing the system to sound an audible alert at a console.
- Create a script directing the system to send an e-mail.
- Create a script directing the system to call a pager.

The system polls the alarm variable and the system compares the result against upper and lower limit values you select when you create the alarm. If the system reaches or crosses the alarm variable during the polling period, the alarm fires and generates an event that you can view in the event log or the trap log. You can configure the alarm to either create a log, or have the alarm send a Simple Network Management Protocol (SNMP) trap to a Network Management System (NMS). You can view the activity in a log or a trap log, or you can create a script to cause a console to beep, send an e-mail, or call a pager.

The upper limit of the alarm is the rising value, and the lower limit is the falling value. RMON1 periodically samples data based upon the alarm interval. During the first interval that the data passes above the rising value, the alarm fires as a rising event. During the first interval that the data drops below the falling value, the alarm fires as a falling event.

The following figure shows how alarms fire:




Figure 210: How alarms fire

The alarm fires during the first interval that the sample goes out of range. No additional events generate for that threshold until the system crosses the opposite threshold. Therefore, you must carefully define the rising and falling threshold values for alarms. Incorrect thresholds cause an alarm to fire at every alarm interval, or never at all.

You can define one threshold value to an expected, baseline value, and then define the opposite threshold as the out-of-bounds limit. Because of sample averaging, the value is equal to ± 1 baseline unit. For example, assume you define an alarm with octets leaving a port as the variable. The intent of the alarm is to notify you if excessive traffic occurs on that port. You enable spanning tree, and then 52 octets transmit from the port every 2 seconds, which is equivalent to baseline traffic of 260 octets every 10 seconds. This alarm notifies you if you define the lower limit of exiting octets at 260 and you define the upper limit at 320 (or at any value greater than $260 + 52 = 312$).

The rising alarm fires the first time outbound traffic, other than spanning tree Bridge Protocol Data Units (BPDUs), occurs. The falling alarm fires after outbound traffic, other than spanning tree, ceases. This process provides the time intervals of any nonbaseline outbound traffic.

If you define the alarm with a falling threshold of less than 260 and the alarm polling interval is at 10 seconds, for example, 250, then the rising alarm can fire only once, as shown in the following example. The falling alarm (the opposite threshold) must fire for the rising alarm to fire a second time. The falling alarm cannot fire unless the port becomes inactive or you disable spanning tree, which causes the value for outbound

octets to drop to zero, because the baseline traffic is always greater than the value of the falling threshold. By definition, the failure of the falling alarm to fire prevents the rising alarm from firing a second time.

The following figure shows an example of the alarm threshold:




Figure 211: Alarm example, threshold less than 260

When you create an alarm, you select a variable from the variable list and a port, or another system component to which it connects. Some variables require port IDs, card IDs, or other indexes, for example, spanning tree group IDs. You then select a rising and a falling threshold value. The rising and falling values compare to the actual value of the variable that you choose. If the variable falls outside of the rising or falling value range, an alarm triggers, and the system logs an event or trap.

When you create an alarm, you also select a sample type, which can be either absolute or delta. Define absolute alarms for alarms based on the cumulative value of the alarm variable. An example of an absolute alarm value is card operating status. Because this value is not cumulative, but instead represents states, such as card up (value 1) and card down (value 2), you configure the value as the absolute value. Therefore, you can create an alarm with a rising value of 2 and a falling value of 1 to alert you whether the card is up or down.

Configure most alarm variables related to Ethernet traffic as a delta value. Define delta alarms for alarms based on the difference in the value of the alarm variable between the start of the polling period and the end of the polling period.



Note

If you create an alarm that monitors a variable that does not exist, you will receive an error message and the creation will fail. Also, if the variable you are monitoring is no longer valid at the time of sampling, the switch removes the alarm automatically. For example, if you create an alarm that monitors some information about a VLAN, and that VLAN is later removed, then the switch silently removes the associated alarm at the next sampling interval.

RMON1 History

The RMON1 history group records periodic statistical samples from a network. A sample is a history and the system gathers the sample in time intervals referred to as buckets.

You can use RMON1 history for the MAC layer in the network. You cannot use RMON1 history for application and network layer protocols.

You enable and create histories to establish a time-dependent method to gather RMON1 statistics on a port. The following are the default values for history:

- Buckets are gathered at 30-minute intervals.
- The number of buckets gathered is 50.

You can configure both the time interval and the number of buckets. However, after the system reaches the last bucket, the system dumps bucket 1 and recycles the bucket to hold a new bucket of statistics. Then the system dumps bucket 2, and so forth.

RMON1 Events

RMON1 events and alarms work together to notify you when values in your network go out of a specified range. After a value passes the specified range, the alarm fires. The event specifies how the system records the activity.

You can use RMON1 events to monitor anything that has a MIB OID associated with it and a valid instance.

An event specifies whether a trap, a log, or both a trap and a log generates to view alarm activity.

You must create an event before associating it with an alarm, otherwise an error occurs. Also, you cannot delete an event as long as there are alarms associated with it. If you try to do so, an error message displays.

RMON1 Statistics

You can use EDM to gather and graph statistics in a variety of formats, or you can save the statistics to a file and export the statistics to a third-party presentation or graphing application.

RMON1 Scaling Limits

The following tables shows the scaling limits for RMON1 elements.



Note

When the log table reaches the maximum 500 log limit, the oldest third of the logs per event is removed to make room for new events. For all other elements, a message displays when you reach the maximum limit and no other element can be added.

Alarms	100
Events	100
History (entries in the history control table with 2000 buckets shared between them)	20

Logs	500
Statistics (entries in stats table)	100

RMON 2

Remote Monitoring (RMON) is a management information base (MIB) or a group of management objects that you use to obtain or configure values using the Simple Network Management Protocol (SNMP).

Use CLI or EDM to globally enable RMON on the system.

After you globally enable RMON, you enable monitoring for individual devices on a port-by-port basis.

RMON1 is the original version of the protocol, which collects information for OSI Layer 1 and Layer 2 in Ethernet networks. RMON1 provides traffic statistics at the MAC layer, and provides statistics on Ethernet segments for packets and bytes received and transmitted.

RMON2 monitors and counts network layer and application layer protocol packets on configured network hosts, either VLAN or port interfaces, that you enable for monitoring. RMON2 monitors Segmented Management Instances at the mgmt configuration level for out-of-band (OOB), circuitless IP (CLIP), and VLAN interfaces.



Note

RMON2 counters on Segmented Management Instance interfaces are cleared only when a Segmented Management Instance interface is newly enabled, or when RMON2 is newly enabled on a previously enabled Segmented Management Instance interface.

The following figure shows which form of RMON monitors which layers in the OSI model:




Figure 212: OSI model and RMON

The RMON2 feature is a management information base (MIB) or a group of management objects that you use to obtain or configure values using the Simple Network Management Protocol (SNMP). The switch supports a partial implementation of RMON2. The RMON2 feature adds the following MIBs: protocol directory, protocol distribution, address map, network-layer host and application layer host for the traffic passing through the (Control Processor) CP for these MIB tables.

The system only collects statistics for IP packets that pass through the CP. RMON2 does not monitor packets on other interfaces processed on the switch that do not pass through the CP. However, RMON2 monitors packets for applications listed in the RMON2 MIB, whether or not the application is enabled or supported on the switch.

After you globally enable RMON2, enable monitoring for individual devices. Identify the network hosts for the system to monitor with a manual configuration on the interfaces you want to monitor.

The RMON2 feature monitors a list of predefined protocols. The system begins to collect protocol statistics immediately after you enable RMON.

The RMON2 feature collects statistics on:

- Protocols predefined by the system.
- Address mapping between physical and network address on particular network hosts that you configure for monitoring.
- Network host statistics for particular hosts on a network layer protocol (IP) that you configure for monitoring.
- Application host statistics for a particular host on an application layer protocol that you configure for monitoring.

RMON2 MIBs

This section describes the following MIBs, on which RMON2 can collect statistics: protocol directory, protocol distribution, address map, network-layer host, and application layer host.

Protocol Directory MIB

The protocol directory lists all of the protocols RMON2 can monitor. The protocols include network layer, transport layer, and application layer protocols, under the OSI model. The system only monitors statistics for the predefined protocols. You cannot delete or add additional protocols to this table. The protocol directory MIB is enabled by default for the predefined protocols.

The predefined protocols include:

- Internet Protocol (IP)
- Secure Shell version 2 (SSHv2)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol (HTTP)
- Telnet
- Remote login (rlogin)



Note

RMON2 can count application packets received on a platform, even if the application is not enabled or supported, before dropping them.

- Trivial File Transfer Protocol (TFTP)
- Simple Network Management Protocol (SNMP)

Protocol Distribution MIB

The protocol distribution MIB collects traffic statistics that each protocol generates by local area network (LAN) segment. The switch acts as the probe and the system collects protocol statistics for the entire switch as part of the group for all of the protocols predefined in the protocol directory table. The protocol distribution control table is part of this group. The protocol distribution control table is predefined with an entry for the management IP for the switch to represent the network segment where the system collects the statistics.

No CLI or EDM support exists to add or delete entries in this table.

Address Map MIB

The address map MIB maps the network layer IP to the MAC layer address.

The system populates the address map control table MIB with an entry for each host interface that you enable for monitoring on the switch.

Network Layer Host MIB

The network layer host MIB monitors the Layer 3 traffic statistics for each host. The network layer host MIB monitors traffic packets in and out of hosts based on the network layer address. The network layer host controls the network and application layer host tables.

The system populates an entry for the management IP of the switch to represent the network segment where the system collects the statistics. You have to enable each host interface that you want to monitor on the switch.

The system only collects statistics for this group from packets that go to the CP.

Application Layer Host MIB

The application layer host MIB monitors traffic statistics by application protocol for each host.

The system populates an entry for the management IP of the switch to represent the network segment where the system collects the statistics. You have to enable each host interface that you want to monitor on the switch.

The system only collects statistics for this group from packets that go to the CP.

RMON2 Considerations

The following considerations apply to RMON2:

- You must enable RMON globally before you enable RMON2 monitoring for a Segmented Management Instance interface.
- You must configure an IPv4 address for the Segmented Management Instance management interface before you enable RMON2 monitoring.
- You can enable RMON on a maximum of 30 IP interfaces on a host.
- You cannot directly configure RMON for a routing VLAN that is an underlying management VLAN. In this case, RMON must be configured at the mgmt vlan configuration level for a non-routing VLAN.
- RMON2 is not available if DHCP Client is configured on a Management Instance. DHCP Client is not available if RMON2 is configured on a Management Instance.
- You cannot delete the IPv4 manual address from a Segmented Management Instance management interface that is RMON enabled. If the only IPv4 address is deleted outside of the normal configuration process, RMON is administratively disabled on the Segmented Management Instance management interface.

RMON Configuration Using CLI

This section contains procedures to configure RMON using Command Line Interface (CLI).

For information about RMON statistics, see the following sections:

- [Display RMON Statistics for Specific Ports](#) on page 3207
- [View RMON Statistics](#) on page 3207

Configure RMON

Enable RMON1 and RMON2 globally, and configure RMON1 alarms, events, history, statistics, and whether port utilization is calculated in half or full duplex. By default, RMON1 and RMON2 are disabled globally.

For RMON1, you enable RMON globally, and then you can use RMON1 alarm, history, events, and statistics for the MAC layer in the network. You cannot use RMON1 history or statistics for application and network layer protocols.

For RMON2, you enable RMON globally, and then you enable RMON on the host interfaces you want to monitor.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable RMON1 and RMON2 globally:

```
rmon
```
3. Configure an RMON1 alarm:

```
rmon alarm <1-65535> WORD <1-1536> <1-3600> {absolute|delta} [falling-  
threshold <-2147483647-2147483647> event <1-65535>] [owner WORD<1-  
127>] [rising-threshold <-2147483647-2147483647> event <1-65535>]
```
4. Configure an RMON1 event:

```
rmon event <1-65535> [community WORD<1-127>] [description WORD<0-127>]  
[log] [owner WORD<1-127>] [trap] [trap_dest [{A.B.C.D}]] [trap_src  
[{A.B.C.D}]]
```
5. Configure RMON1 history:

```
rmon history <1-65535> {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} [buckets <1-65535>] [interval <1-3600>] [owner  
WORD<1-127>]
```
6. Configure RMON1 statistics:

```
rmon stats <1-65535> {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} [owner <1-127>]
```
7. Configure whether the system calculates port utilization in half or full duplex:

```
rmon util-method [half|full]
```

Example

Configure RMON globally, an RMON1 alarm, and RMON1 event:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#rmon
Switch:1(config)#rmon event 60534 community public description "Rising Event" log trap
Switch:1(config)#rmon alarm 4 rcCliNumAccessViolations.0 10 absolute rising-threshold 2
event 60000
```

Variable Definitions

Use the data in this table to use the **rmon** command.

Variable	Value
<pre>alarm <1-65535> WORD <1-1536> <1-3600> {absolute delta} [falling-threshold <-2147483647-2147483647> event <1-65535>] [owner WORD<1-127>] [rising-threshold <- 2147483647-2147483647> event <1-65535>]</pre>	<p>Creates an alarm interface.</p> <ul style="list-style-type: none"> • <1-65535>— Specifies the interface index number from 1 to 65535. Each entry defines a diagnostics sample at a particular interval for an object on the device. The default is 1. • WORD <1-1536>— Specifies the variable name or OID. The entry is case sensitive and can have a string length of 1 to 1536. • {absolute delta}— Specifies the sample type. • rising-threshold <-2147483648-2147483647> [<event:1-65535>] — Specifies the rising threshold from -2147483648 to 2147483647, which is a threshold for the sampled statistic. After the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, the system generates a single event. The system also generates a single event if the first sample after this entry that becomes valid is greater than or equal to the rising alarm, or the rising or falling alarm. After the system generates a rising event, the system does not generate another such event until the sampled value falls below this threshold and reaches the alarm falling threshold. You cannot modify this object if the associated alarm status is equal to valid. • <1-65535>— Specifies the rising event index, which the system uses after the system crosses a rising threshold. The event entry identified by a particular value of this index is the same as identified by the same value of the event index object. If no corresponding entry exists in the event table, no association exists. In particular, if this value is zero, the system does not generate an associated event, as zero is not a valid event index. You cannot modify this object if the associated alarm status is equal to valid. • falling-threshold <-2147483648-2147483647> [<event:1-65535>] — Specifies the falling threshold from -2147483648 to 2147483647, which specifies a threshold for the sampled statistic. If the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, the system generates a single event. The system also generates a single event if the first sample after this entry that becomes valid

Variable	Value
	<p>is less than or equal to this threshold and the associated alarm startup alarm is equal to falling alarm or rising or falling alarm. After the system generates a falling event, the system does not generate another such event until the sampled value rises above this threshold, and reaches the alarm rising threshold. You cannot modify this object if the associated alarm status is equal to valid.</p> <p><1-65535> – Specifies the index of the event entry that the system uses after a falling threshold is crossed. The event entry identified by a particular value of this index is the same as identified by the same value of the event index object. If no corresponding entry in the event table exists, no association exists. In particular, if this value is zero, the system does not generate an event, as zero is not a valid event index. You cannot modify this object if the associated alarm status is equal to valid. The default is 60535.</p> <ul style="list-style-type: none"> • <i>owner WORD<1-127></i> — Specifies the name of the owner, with a string length 1 to 127. <p>Note: When configuring from CLI, the default owner is cli; when configuring with SNMP, the default owner is snmp. The default command only sets the owner to default. No other parameters can be changed after you create the alarm.</p>
<pre>event <1-65535> [community WORD<1-127>] [description WORD<0-127>] [log] [owner WORD<1-127>] [trap]</pre>	<p>Create an event.</p> <ul style="list-style-type: none"> • <1-65535>— Specifies the event index number. Each entry defines one event that the system generates after the appropriate conditions occur. The default is 1. • <i>log</i> — Specifies if this event stores a log when the event is triggered by the alarm. • <i>trap</i> — Specifies if this event sends a trap when the event is triggered by the alarm. The trap will be sent to all the snmp-server hosts configured in the snmp table. • <i>description WORD<0-127></i>— Specifies the event description, with a string length of 0 to 127. • <i>owner WORD<1-127></i> — Specifies the name of the owner, with a string length of 1 to 127. • <i>community WORD<1-127></i> — Specifies the SNMP community where you can send SNMP traps, with a string length 1 to 127.

Variable	Value
	<p>You can set the community, but the trap is not filtered out. The trap is sent to all configured snmp-server hosts, regardless of the value of this field.</p>
<pre>history <1-65535> {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all] [all]} [buckets <1-65535>] [interval <1-3600>] [owner WORD<1-127>]</pre>	<p>Configures RMON history.</p> <ul style="list-style-type: none"> • <i><1-65535></i> — Specifies the history index number that uniquely identifies an entry in the history control table. Each entry defines a set of samples at a particular interval for an interface on the default. The default value is 1. • {slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all] [all]} <ul style="list-style-type: none"> — Specifies the single port interface. Identifies the source for which the system collects and places historical data in a media-specific table on behalf of this history control entry. The source is an interface on this device. The statistics in this group reflect all packets on the local network segment that attaches to the identified interface. • <i>buckets <1-65535></i>— Specifies the requested number of discrete time intervals where the system saves data in the part of the media-specific table associated with this history control entry. The default value is 50. • <i>interval <1-3600></i>— Specifies the time interval in seconds over which the system samples the data for each bucket in the part of the media-specific table associated with this history control entry. Because the counters in a bucket can overflow at their maximum value with no indication, you must take into account the possibility of overflow in all the associated counters. Consider the minimum time in which a counter can overflow on a particular media type, and then set the history control interval to a value less than this interval, which is typically most important for the octets counter in a media-specific table. The default value is 1800. • <i>owner WORD<1-127></i>— Specifies the name of the owner.

Variable	Value
<code>stats</code> <code><1-65535> {slot/port[/sub-port]</code> <code>[-slot/port[/sub-port]][,...][slot/all]</code> <code>[all]}owner WORD<1-127></code>	Configures RMON statistics. <ul style="list-style-type: none"> • <code><1-65535></code>— Specifies the control Ether statistics entry index number. • <code>{slot/port [/sub-port] [-slot/port [/sub-port] [, ...]}</code>— Specifies the single port interface. • <code>owner WORD<1-127></code> — Specifies the name of the owner.
<code>util-method [half full]</code>	Configures whether port utilization is calculated in half or full duplex to calculate port usage. <ul style="list-style-type: none"> • <code>half</code>—Configures the string to half duplex. • <code>full</code>—Configures the string to full duplex. <p>After you select half for half duplex, RMON uses InOctets and the speed of the port to calculate port usage (this is the standard RMON RFC 1271 convention). After you select full for full duplex, RMON uses InOctets and OutOctets, and 2X the speed of the port to calculate port usage. If you select full, but the port operates in half-duplex mode, the calculation defaults to the RFC 1271 convention. The default is half.</p>

Enable Remote Monitoring on an Interface

Use the following procedure to enable Remote Monitoring (RMON) on an interface.

Before You Begin

- Enable RMON globally.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable RMON on a particular VLAN:

```
vlan rmon <1-4059>
```

3. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

4. Enable RMON on a particular port:

```
rmon
```

Examples

Enable RMON on VLAN 2:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#vlan rmon 2
```

Enable RMON on port 1/8:

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#interface gigabitethernet 1/8  
Switch:1(config-if)#rmon
```

Variable Definitions

Use the data in this table to use the **vlan rmon** command.

Variable	Value
<1-4059>	Specifies the VLAN ID on which to configure RMON.

Enable RMON2 on a Segmented Management Instance Interface

Before You Begin

You must enable RMON globally before you enable RMON2 monitoring for a Segmented Management Instance interface.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enter the configuration mode for the Management Instance:

```
mgmt <clip | oob | vlan>
```

3. Enable RMON2 on the Segmented Management Instance:

```
rmon
```

View RMON Information

View RMON1 and RMON2 information on the switch. You can view information about RMON1 alarms, events, history, logs, and statistics. You can also view RMON2 information about application host statistics, control tables, network host statistics, and protocol distribution statistics.

Procedure

1. View RMON1 information:

```
show rmon {alarm|event|history|log|stats}
```

2. View RMON2 information:

```
show rmon {address-map|application-host-stats WORD<1-64>|application  
protocols|ctl-table|protocol-dist-stats|network-host-stats}
```

Example

View RMON event, log, and statistics information:

```
Switch:1(config)#show rmon event

=====
Rmon Event
=====

INDEX DESCRIPTION      TYPE      COMMUNITY OWNER      LAST_TIME_SENT
-----
60534 Rising Event     log-and-trap public     192.0.2.155 none
60535 Falling Event    log-and-trap public     192.0.2.155 8 day(s), 19:14:32

Switch:1(config)#show rmon log

=====
Rmon Log
=====

INDEX    TIME                DESCRIPTION
-----
60535. 1 8 day(s), 19:14:45  1.3.6.1.4.1.2272.1.19.14.0 (absValue = 0, Falling
                                Threshold = 2, interval = 10)[alarmIndex.1][trap]
                                "Falling Event"
60535. 2 8 day(s), 19:14:45  1.3.6.1.4.1.2272.1.19.14.0 (absValue = 0, Falling
                                Threshold = 1, interval = 10)[alarmIndex.2][trap]
                                "Falling Event"

Switch:1(config)#show rmon stats

=====
Rmon Ether Stats
=====

INDEX PORT  OWNER
-----
1      1/10    monitor
```

Variable Definitions

The following table defines parameters for the **show rmon** command.

Variable	Value
<i>address-map</i>	Displays the RMON2 address map. This RMON2 parameter expands RMON capacity to display information on network, transport, and application layers.
<i>alarm</i>	Displays the RMON1 alarm table.
<i>application-host-stats WORD<1-64></i>	Displays RMON2 application host statistics from monitored protocols. This RMON2 parameter expands RMON capacity to display network, transport, and application layer information. Note: RMON2 can count application packets received on a platform, even if the application is not enabled or supported, before dropping them.
<i>ctl-table</i>	Displays the RMON2 control tables. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
<i>event</i>	Displays the RMON1 event table.
<i>history</i>	Displays the RMON1 history table. This RMON1 parameter displays and is limited to link layer information, including as MAC information.
<i>log</i>	Displays the RMON1 log table.
<i>network-host-stats</i>	Displays RMON2 network-host statistics. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
<i>protocol-dist-stats</i>	Displays RMON2 protocol distribution statistics. This RMON2 parameter expands RMON capacity to display network, transport, and application layers.
<i>stats</i>	Displays the RMON1 statistics table. This RMON1 parameter displays and is limited to link layer information, including as MAC information.

Display RMON Address Maps

View the maps of network layer address to physical address to interface.

The probe adds entries based on the source MAC and network addresses in packets without MAC-level errors.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View RMON address maps:

```
show rmon address-map
```

Example

```
Switch:1# show rmon address-map
```

```
=====
                        Rmon Address Map Table
=====
PROTOIDX  HOSTADDR      SOURCE  PHYADDR      LASTCHANGE
-----
1         192.0.2.11    2060    b0:ad:aa:42:a5:03  10/09/15 17:30:41
```

View RMON Statistics

About This Task

View RMON statistics to manage network performance.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View RMON statistics:

```
show rmon stats
```

Example

```
Switch:1#show rmon stats
```

```
=====
                        Rmon Ether Stats
=====
INDEX  PORT  OWNER
-----
1      1/10  monitor
```

Display RMON Statistics for Specific Ports

Display individual RMON statistics for specific ports to manage network performance.

**Note**

Slot and port information can differ depending on hardware platform.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View statistics for specific ports:
`show interfaces GigabitEthernet statistics rmon {slot/port[/sub-port]} [-slot/port[/sub-port]][,...]`

Example

View RMON statistics:

```
Switch:1>enable
Switch:1#show interfaces gigabitEthernet statistics rmon 1/13
```

```
=====
```

Port Stats Rmon									
PORT NUM	OCTETS	PKTS	MULTI CAST	BROAD CAST	CRC ALIGN	UNDER SIZE	OVER SIZE	FRAG MENT	COLLI SION
1/13	1943	21	8	13	0	0	0	0	0

Variable Definitions

Use the data in the following table to use the **show interfaces GigabitEthernet statistics rmon** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View RMON Application Host Statistics

View application host statistics to see traffic statistics by application protocol for each host.

Procedure

- To enter User EXEC mode, log on to the switch.
- View RMON application host statistics:

```
show rmon application-host-stats WORD<1-64>
```

Example

```
Switch:1# show rmon application-host-stats ?
WORD<1-64>  Select one of these application protocols
              {TCP|UDP|FTP|TELNET|HTTP|RLOGIN|SSH|TFTP|SNMP|HTTPS}
Switch:1# show rmon application-host-stats FTP
```

Rmon Application Host Stats					
HOSTADDR	INPKT	OUTPKT	INOC	OUTOC	CREATETIME
192.0.2.10	0	0	0	0	10/09/15 17:29:54



Note

Protocol support can vary across platforms.

Variable Definitions

The following table defines parameters for the **show rmon application-host-stats** command.

Variable	Value
WORD<1-64>	Specifies one of the following application protocols: TCP, UDP, FTP, TELNET, HTTP, rlogin, SSH, TFTP, SNMP, HTTPS. Note: RMON2 can count application packets received on a platform, even if the application is not enabled or supported, before dropping them.

View RMON Control Tables

View RMON control tables to see the data source for both network layer and application layer host statistics.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View RMON control tables:

```
show rmon ctl-table
```

Example



Note

RMON2 can count application packets received on a platform, even if the application is not enabled or supported, before dropping them.

```
Switch:1# show rmon ctl-table
```

```
=====
                        Rmon Control Table
=====

=====
                        Protocol Directory Table
=====
```

IDX	PROTOCOL	ADDRMAPCFG	HOSTCFG	MATRIXCFG	OWNER
1	IP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
2	TCP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
3	UDP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
4	FTP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
5	SSH	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
6	TELNET	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
7	HTTP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
8	RLOGIN	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
9	TFTP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
10	SNMP	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1
11	HTTPS	SUPPORTED	SUPPORTED	NOT SUPPORTED	Switch-1

Protocol Distribution Control Table				
IDX	DATASOURCE	DROPFRAMES	CREATETIME	OWNER
1	0.0.0.0	0	09/22/15 19:29:13	Switch-1

Address Map Control Table			
IDX	DATASOURCE	DROPFRAMES	OWNER
1	0.0.0.0	0	Switch-1

Host Control Table				
IDX	DATASOURCE	NHDROPFRAMES	AHDROPFRAMES	OWNER
1	0.0.0.0	0	0	Switch-1

Display RMON Network Host Statistics

View network host statistics to see Layer 3 traffic statistics for each host. The network layer host MIB monitors traffic packets in and out of hosts based on the network layer address.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View RMON network host statistics:

```
show rmon network-host-stats
```

View RMON Protocol Distribution Statistics

View protocol distribution statistics to see traffic statistics that each protocol generates by local area network (LAN) segment.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View RMON protocol distribution statistics:

```
show rmon protocol-dist-stats
```

Example



Note

RMON2 can count application packets received on a platform, even if the application is not enabled or supported, before dropping them.

```
Switch:1# show rmon protocol-dist-stats
```

```
=====
Rmon Protocol Dist Stats
```


PROTOCOL	PKTS	OCTETS
IP	0	0
TCP	0	0
UDP	0	0
FTP	0	0
SSH	0	0
TELNET	0	0
HTTP	0	0
RLOGIN	0	0
TFTP	0	0
SNMP	0	0
HTTPS	0	0

Display RMON Status

View the current RMON status on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View RMON status:

```
show rmon
```

Example

```
Switch:1# show rmon

RMON Info :
Status      : enable
```

View the RMON2 Configuration State of Management Interfaces

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View information about the RMON2 configuration state:

```
show mgmt rmon
```

Example

```
Switch:1>show mgmt rmon

=====
Mgmt Rmon Information
=====
INST  DESCR  RMON-ADMIN-ENABLE  RMON-OPER-ENABLE  RMON-IP-ADDR
-----
1 Mgmt-oob1 disable disable 0.0.0.0
3 Mgmt-clip enable enable 192.0.2.72
4 Mgmt-vlan enable enable 198.51.100.72
```

RMON Configuration Using EDM

This section contains procedures to configure RMON using Enterprise Device Manager (EDM).

For information about RMON statistics, see the following sections:

- [Enable RMON Statistics](#) on page 3224
- [View RMON Statistics](#) on page 3225

Enable RMON Globally

About This Task

You must globally enable RMON before you can use RMON2 functions. If you attempt to enable an RMON2 function before the global flag is disabled, EDM informs you that the flag is disabled and prompts you to enable the flag. You can configure RMON1 while RMON is globally disabled.

If you want to use nondefault RMON parameter values, you can configure them before you enable RMON, or as you configure the RMON functions.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Options**.
3. Select the **Options** tab.
4. Select **Enable**.
5. In **UtilizationMethod**, select a utilization method.
6. Select **Apply**.

Options Field Descriptions

Use the data in the following table to use the **Options** tab.

Name	Description
Enable	Enables RMON. If you select Enable , the RMON agent starts immediately. To disable RMON, clear Enable , select Apply , and then restart the device. The default is disabled.
UtilizationMethod	Controls whether RMON uses a half-duplex or full-duplex formula to calculate port usage. After you select halfDuplex, RMON uses InOctets and the speed of the port to calculate port usage (this is the standard RMON RFC 1271 convention). After you select fullDuplex, RMON uses InOctets and OutOctets and 2X the speed of the port to calculate port usage. If you select fullDuplex, but the port operates in half-duplex mode, the calculation defaults to the RFC 1271 convention. The default is halfDuplex.

Enable RMON on a Port or VLAN

Before You Begin

- Enable RMON globally.

About This Task

Use the following procedure to enable RMON on an interface.

Procedure

1. Enable RMON on a VLAN:
 - a. In the navigation pane, expand **Configuration > VLAN**.
 - b. Select **VLANs**.
 - c. Select the **Advanced** tab.
 - d. In the row for the VLAN, double-click **RmonEnable**, and then select **enable**.
 - e. Select **Apply**.
2. Enable RMON on a port:
 - a. In the Device Physical View, select a port.
 - b. In the navigation pane, expand **Configuration > Edit > Port**.
 - c. Select **General**.
 - d. Select the **Interface** tab.
 - e. For **RmonEnable**, select **enable**.
 - f. Select **Apply**.

Enable RMON History

About This Task

Use RMON to establish a history for a port and configure the bucket interval. For example, to gather RMON statistics over the weekend, you must have enough buckets to cover two days. Configure the history to gather one bucket every hour, and cover a 48-hour period. After you configure the history characteristics, you cannot modify them; you must delete the history and create another one.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Control**.
3. In the **History** tab, select **Insert**.
4. In **Port**, select the ellipsis (...).
5. Select a port.
6. Select **OK**.
7. In **Buckets Requested**, type the number of discrete time intervals to save data.
8. In **Interval**, type the interval in seconds.
9. In **Owner**, type the owner information.
10. Select **Insert**.

History Field Descriptions

Use the data in the following table to use the **History** tab.

Name	Description
Index	Specifies an index that uniquely identifies an entry in the historyControl table. Each entry defines a set of samples at a particular interval for an interface on the device. Index value ranges from 1–65535. The default value is 1.
Port	Identifies the source for which the system collects and places historical data in a media-specific table on behalf of this historyControlEntry. The source is an interface on this device. To identify a particular interface, the object identifies the instance of the ifIndex object, defined in (4,6), for the desired interface. For example, if an entry receives data from interface 1, the object is ifIndex 1. The statistics in this group reflect all packets on the local network segment attached to the identified interface. You cannot modify this object if the associated historyControlStatus object is equal to valid(1).
BucketsRequested	Specifies the requested number of discrete time intervals over which the system save data in the part of the media-specific table associated with this historyControlEntry. After this object is created or modified, the probe configures historyControlBucketsGranted as closely to this object as possible for the particular probe implementation and available resources. Values range from 1–65535. The default value is 50.
BucketsGranted	Specifies the number of discrete sampling intervals over which the system save data in the part of the media-specific table associated with this historyControlEntry. After the associated BucketsRequested object is created or modified, the probe sets this object as closely to the requested value as possible for the particular probe implementation and available resources. The probe must not lower this value except as a result of a modification to the associated BucketsRequested object. Occasionally, the actual number of buckets associated with this entry is less than the value of this object. In this case, at the end of each sampling interval, the system adds a new bucket to the media-specific table. After the number of buckets reaches the value of this object and the system is going to add a new bucket to the media-specific table, the agent deletes the oldest bucket associated with this entry so the system can added the new bucket. After the value of this object changes to a value less than the current value, entries are deleted from the media-specific table associated with this entry. The agent deletes the oldest of these entries so that their number remains less than or equal to the new value of this object. After the value of this object changes to a value greater than the current value, the system allows the number of associated media-specific entries to grow.

Name	Description
Interval	Specifies the interval in seconds over which the system samples data for each bucket in the part of the media-specific table associated with this historyControlEntry. You can set this interval between 1–3600 seconds (1 hour). Because the counters in a bucket can overflow at their maximum value with no indication, you must take into account the possibility of overflow in all of the associated counters. Consider the minimum time in which a counter can overflow on a particular media type, and then set the historyControlInterval object to a value less than this interval, which is typically most important for the octets counter in a media-specific table. For example, on an Ethernet network, the etherHistoryOctets counter can overflow in approximately 1 hour at the maximum utilization. You cannot modify this object if the associated historyControlStatus object is equal to valid. The default value is 1800.
Owner	Specifies the entity that configured this entry and uses the assigned resources.

Disable RMON1 History

About This Task

Disable RMON1 history on a port if you do not want to record a statistical sample from that port.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Control**.
3. In **History**, select the row that contains the port ID to delete.
4. Select **Delete**.

View RMON1 History Statistics

About This Task

View RMON1 history statistics when you want to see a statistical sample from the switch. You can create a graph of the statistics in a bar, pie, chart, or line format.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **RMON History** tab.
5. Select the statistics you want to graph.
6. Select the button for the type of graph you require (bar, pie, chart, or line).

RMON History *Field Descriptions*

Use the data in the following table to use the **RMON History** tab.

Parameter	Description
SampleIndex	Identifies the particular sample this entry represents among all samples associated with the same history control entry. This index starts at one and increases by one as each new sample is taken.
Utilization	Specifies the best estimate of the mean physical layer network utilization on this interface during the sampling interval, in hundredths of a percent.
Octets	Specifies the total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets)
Pkts	Specifies the number of packets (including bad packets) received during this sampling interval.
BroadcastPkts	Specifies the number of good packets received during this sampling interval that were directed to the broadcast address.
MulticastPkts	Specifies the number of good packets received during this sampling interval that the system directs to a multicast address. This number does not include packets addressed to the broadcast address.
DropEvents	Specifies the total number of events in which the probe dropped packets due to lack of resources during this sampling interval. This number is not necessarily the number of packets dropped; it is only the number of times the system detects this condition.
CRCAAlignErrors	The number of packets the system receives during this sampling interval that had a length (excluding framing bits but including FCS octets) from 64–1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
UndersizePkts	Specifies the number of packets the system receives during this sampling interval that were less than 64 octets (excluding framing bits but including FCS octets), and were otherwise well formed.
OversizePkts	Specifies the number of packets the system receives during this sampling interval that were longer than 1518 octets (excluding framing bits but including FCS octets), but were otherwise well formed.

Parameter	Description
Fragments	<p>Specifies the total number of packets received during this sampling interval that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</p> <p>It is entirely normal for Fragments to increment because it counts both runts (which are normal occurrences due to collisions) and noise hits.</p>
Collisions	<p>Specifies the best estimate of the total number of collisions on this Ethernet segment during this sampling interval. The value returned depends on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station must detect a collision in the receive mode if three or more stations transmit simultaneously. A repeater port must detect a collision when two or more stations transmit simultaneously. Thus, a probe placed on a repeater port can record more collisions than a probe connected to a station on the same segment.</p> <p>Probe location plays a small role when 10BASE-T. 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can detect only collisions when it transmits. Thus, probes placed on a station and a repeater can report the same number of collisions.</p> <p>An RMON probe inside a repeater can ideally report collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.</p>

Create an RMON Alarm

Before You Begin

- You must globally enable RMON.

About This Task

After you enable RMON globally, you also create a default rising and falling event. The default for the events is log-and-trap, which means that you receive notification through a trap as well as through a log entry.

Procedure

- In the navigation pane, expand **Configuration > Serviceability > RMON**.
- Select **Alarms**.
- Select the **Alarms** tab.
- Select **Insert**.

5. In **Variable**, select a variable for the alarm.
If you select some variables, the system prompts for a port (or other object) on which you want to configure an alarm.
6. In **SampleType**, select a sample type.
7. In **Interval**, type a sample interval in seconds.
8. In **Index**, type an index number.
9. In **RisingThreshold**, type a rising threshold value.
10. In **RisingEventIndex**, type a rising threshold event index.
11. In **FallingThreshold**, type a falling threshold value.
12. In **FallingEventIndex**, type a falling threshold event index.
13. In **Owner**, type the owner of the alarm.
14. Select **Insert**.

Alarms Field Descriptions

Use the data in the following table to use the **Alarms** tab.

Name	Description
Index	Uniquely identifies an entry in the alarm table. Each entry defines a diagnostic sample at a particular interval for an object on the device. The default is 1.
Interval	Specifies the interval, in seconds, over which the data is sampled and compared with the rising and falling thresholds. deltaValue sampling—Configures the interval short enough that the sampled variable is unlikely to increase or decrease by more than $2^{31}-1$ during a single sampling interval.

Name	Description
Variable	<p>Specifies the object identifier of the particular variable to be sampled. Only variables that resolve to an ASN.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) can be sampled.</p> <p>Alarm variables exist in three formats, depending on the type:</p> <ul style="list-style-type: none"> • A chassis, power supply, or fan-related alarm ends in x where the x index is hard-coded. No further information is required. • A card, spanning tree group (STG), or EtherStat alarm ends with a dot (.). You must enter a card number, STG ID, IP address, or EtherStat information. • A port alarm ends with no dot or index and requires that you use the port shortcut menu. An example of a port alarm is ifInOctets (interface incoming octet count). <p>Because the system articulates SNMP access control entirely in terms of the contents of MIB views, no access control mechanism exists to restrict the value of this object to identify only those objects that exist in a particular MIB view. Because no acceptable means of restricting the read access that is obtained through the alarm mechanism exists, the probe must grant only write access to this object in those views that have read access to all objects on the probe.</p> <p>After you configure a variable, if the supplied variable name is not available in the selected MIB view, the system returns a badValue error. After the variable name of an established alarmEntry is no longer available in the selected MIB view, the probe changes the status of this alarmEntry to invalid.</p> <p>You cannot modify this object if the associated alarmStatus object is equal to valid.</p>
SampleType	<p>Specifies the method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is absoluteValue, the value of the system compares the selected variable directly with the thresholds at the end of the sampling interval. If the value of this object is deltaValue, the system subtracts the value of the selected variable at the last sample from the current value, and the system compares the difference with the thresholds. You cannot modify this object if the associated alarmStatus object is equal to valid. The default is deltaValue.</p>
Value	<p>Specifies the value of the statistic during the last sampling period. For example, if the sample type is deltaValue, this value is the difference between the samples at the beginning and end of the period. If the sample type is absoluteValue, this value is the sampled value at the end of the period. This system compares the value with the rising and falling thresholds. The value during the current sampling period is not made available until the period is completed and remains available until the next period is complete.</p>

Name	Description
StartupAlarm	Specifies the alarm that is sent after this entry is first set to valid. If the first sample after this entry becomes valid is greater than or equal to the risingThreshold and alarmStartupAlarm is equal to the risingAlarm or the risingOrFallingAlarm, then the system generates a single rising alarm. If the first sample after this entry becomes valid is less than or equal to the fallingThreshold and alarmStartupAlarm is equal to the fallingAlarm or the risingOrFallingAlarm, then the system generates a single falling alarm. You cannot modify this object if the associated alarmStatus object is equal to valid.
RisingThreshold	Specifies a threshold for the sampled statistic. After the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, the system generates a single event. The system also generates a single event if the first sample after this entry becomes valid is greater than or equal to this threshold and the associated alarmStartupAlarm is equal to risingAlarm or risingOrFallingAlarm. After a rising event is generated, another such event is not generated until the sampled value falls below this threshold and reaches the alarmFallingThreshold. You cannot modify this object if the associated alarmStatus object is equal to valid.
RisingEventIndex	Specifies the index of the eventEntry that is used after a rising threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of the eventIndex object. If no corresponding entry exists in the eventTable, no association exists. In particular, if this value is zero, the system generates no associated event, as zero is not a valid event index. You cannot modify this object if the associated alarmStatus object is equal to valid. Note: You must create the event prior to associating it to an alarm.
FallingThreshold	Specifies a threshold for the sampled statistic. If the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, the system generates a single event. The system also generates a single event if the first sample after this entry becomes valid is less than or equal to this threshold and the associated alarmStartupAlarm is equal to fallingAlarm or risingOrFallingAlarm. After the system generates a falling event, the system does not generate another similar event until the sampled value rises above this threshold and reaches the alarmRisingThreshold. You cannot modify this object if the associated alarmStatus object is equal to valid.

Name	Description
FallingEventIndex	<p>Specifies the index of the eventEntry that the system uses after a falling threshold is crossed. The eventEntry identified by a particular value of this index is the same as identified by the same value of the eventIndex object. If there is no corresponding entry in the eventTable, no association exists. In particular, if this value is zero, the system generates no associated event, as zero is not a valid event index. You cannot modify this object if the associated alarmStatus object is equal to valid.</p> <p>Note: You must create the event prior to associating it to an alarm.</p>
Owner	Specifies the entity that configured this entry and is therefore using the resources assigned to it.
Status	Specifies the status of this alarm entry.

View RMON1 Alarms

About This Task

View the RMON1 alarm information to see alarm activity.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Alarms**.
3. Select the **Alarm** tab.

Delete an RMON1 Alarm

About This Task

Delete an RMON1 alarm if you no longer display it in the log.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Alarms**.
3. Select the alarm to delete.
4. Select **Delete**.

Create an RMON1 Event

About This Task

Create a custom rising and falling RMON1 event to specify if alarm information is sent to a trap, a log, or both.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Alarms**.
3. Select the **Events** tab.
4. Select **Insert**.
5. In **Description**, type an event name.
6. In **Type**, select an event type.

The default configuration is **log-and-trap**. To save memory, configure the event type to **log**. To reduce traffic from the system, configure the event type to **snmp-log**.

If you select **snmp-trap** or **log**, you must configure trap receivers.

7. In **Community**, type an SNMP community.
8. In **Owner**, type the owner of this event.
9. Select **Insert**.

Events Field Descriptions

Use the data in the following table to use the **Events** tab.

Name	Description
Index	Uniquely identifies an entry in the event table. Each entry defines one event that the system generates after the appropriate conditions occur. The default is 1.
Description	Specifies a comment that describes this event entry.
Type	Specifies the type of notification that the probe makes about this event. In the case of a log, the system makes an entry in the log table for each event. In the case of SNMP traps, the system sends an SNMP trap to one or more management stations.
Community	Specifies the SNMP community where you can send SNMP traps.
LastTimeSent	Specifies the value of sysUpTime at the time this event entry last generated an event. If this entry has not generated events, this value is zero.
Owner	Specifies the entity that configured this entry and is therefore using the assigned resources. If this object contains a string starting with monitor and has associated entries in the log table, all connected management stations retrieve those log entries, as they have significance to all management stations connected to this device.

View RMON1 Events

About This Task

View RMON1 events to see how many events occurred.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Alarms**.
3. Select the **Events** tab.

Events Field Descriptions

Use the data in the following table to use the **Events** tab.

Name	Description
Index	Uniquely identifies an entry in the event table. Each entry defines one event that the system generates after the appropriate conditions occur. The default is 1.
Description	Specifies a comment that describes this event entry.
Type	Specifies the type of notification that the probe makes about this event. In the case of a log, the system makes an entry in the log table for each event. In the case of SNMP traps, the system sends an SNMP trap to one or more management stations.
Community	Specifies the SNMP community where you can send SNMP traps.
LastTimeSent	Specifies the value of sysUpTime at the time this event entry last generated an event. If this entry has not generated events, this value is zero.
Owner	Specifies the entity that configured this entry and is therefore using the assigned resources. If this object contains a string starting with monitor and has associated entries in the log table, all connected management stations retrieve those log entries, as they have significance to all management stations connected to this device.

Delete an Event

About This Task

Delete an event after you no longer require the alarm information.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Alarms**.
3. Select the **Events** tab.
4. Select the event to delete.
5. Select **Delete**.

View the RMON Log

About This Task

View the trap log to see which activity occurred.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Alarms**.
3. Select the **Log** tab.

Log Field Descriptions

Use the data in the following table to use the **Log** tab.

Name	Description
EventIndex	Specifies an index that uniquely identifies an entry in the event table. Each entry defines one event that is generated under appropriate conditions.
Index	Specifies an index that uniquely identifies an entry in the log table generated by the same event entries.
Time	Specifies the creation time for this log entry.
Description	Specifies an implementation dependent description of the event that activated this log entry.

Enable RMON Statistics

About This Task

Enable Ethernet statistics collection for RMON.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Control**.
3. Select the **Ethernet Statistics** tab.
4. Select **Insert**.
5. Next to **Port**, select the ellipsis (...).
6. Select a port.
7. Select **OK**.
8. In **Owner**, type the name of the owner entity.
9. Select **OK**.
10. Select **Insert**.

Ethernet Statistics *Field Descriptions*

Use the data in the following table to use the **Ethernet Statistics** tab.

Name	Description
Index	Uniquely identifies an entry in the Ethernet Statistics table. The default is 1.
Port	Identifies the source of the data that this etherStats entry is configured to analyze.
Owner	Specifies the entity that configured this entry and therefore uses the assigned resources.

View RMON Statistics

Before You Begin

- You must enable RMON statistics collection.

About This Task

Use the following procedure to view RMON statistics for each port.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **RMON** tab.
5. Select the statistics you want to graph.
6. Select a graph type:
 - bar
 - pie
 - chart
 - line

RMON Field Descriptions

The following table describes fields on the **RMON** tab.

Name	Description
Octets	<p>Specifies the number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets).</p> <p>You can use this object as a reasonable estimate of Ethernet utilization. If additional precision is desired, sample the Pkts and Octets objects before and after a common interval. The differences in the sampled values are Pkts and Octets, and the number of seconds in the interval is Interval. These values are used to calculate the Utilization as follows:</p> $\text{Pkts} * (9.6 + 6.4) + (\text{Octets} * .8)$ <p>Utilization = Interval * 10,000</p> <p>The result of this equation is the value Utilization, which is the percent utilization of the Ethernet segment on a scale of 0 to 100 percent.</p>
Pkts	Specifies the number of packets (including bad packets, broadcast packets, and multicast packets) received.
BroadcastPkts	Specifies the number of good packets received that were directed to the broadcast address. This number does not include multicast packets.
MulticastPkts	Specifies the number of good packets received that were directed to a multicast address. This number does not include packets directed to the broadcast address.
CRCAAlignErrors	Specifies the number of packets received that had a length (excluding framing bits, but including FCS octets) between 64 and 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
UndersizePkts	Specifies the number of packets received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
OversizePkts	Specifies the number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.

Name	Description
Fragments	<p>Specifies the number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</p> <p>It is entirely normal for Fragments to increment because it counts both runs (which are normal occurrences due to collisions) and noise hits.</p>
Collisions	<p>Specifies the best estimate of the number of collisions on this Ethernet segment. The value returned depends on the location of the RMON probe. Section 8.2.1.3 (10BASE-5) and section 10.3.1.3 (10BASE-2) of IEEE standard 802.3 states that a station must detect a collision in the receive mode if three or more stations are transmitting simultaneously. A repeater port must detect a collision when two or more stations transmit simultaneously. Thus, a probe placed on a repeater port can record more collisions than a probe connected to a station on the same segment.</p> <p>Probe location plays a much smaller role when considering 10BASE-T. 14.2.1.4 (10BASE-T) of IEEE standard 802.3 defines a collision as the simultaneous presence of signals on the DO and RD circuits (transmitting and receiving at the same time). A 10BASE-T station can only detect collisions when it is transmitting. Thus, probes placed on a station and a repeater reports the same number of collisions.</p> <p>An RMON probe inside a repeater reports collisions between the repeater and one or more other hosts (transmit collisions as defined by IEEE 802.3k) plus receiver collisions observed on any coax segments to which the repeater is connected.</p>

[View the Protocol Directory](#)

About This Task

View the protocol directory to see the list of protocols that RMON2 can monitor. You cannot change the list of protocols.

The protocol directory MIB is enabled by default for the predefined protocols.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Protocol Directory**.
3. Select the **Protocol Directories** tab.

Protocol Directories Field Descriptions

The following table defines parameters for the **Protocol Directories** tab.

Name	Description
Index	Shows a unique identifier for the entry in the table.
Protocol	<p>Shows the protocols RMON2 can monitor:</p> <ul style="list-style-type: none"> • Internet Protocol (IP) • Secure Shell version 2 (SSHv2) • Transmission Control Protocol (TCP) • User Datagram Protocol (UDP) • File Transfer Protocol (FTP) • Hypertext Transfer Protocol (HTTP) • Telnet • Remote login (rlogin) • Trivial File Transfer Protocol (TFTP) • Simple Networking Management Protocol (SNMP) <p>Note: RMON2 can count application packets received on a platform, even if the application is not enabled or supported, before dropping them.</p>
AddressMapConfig	<p>Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following:</p> <ul style="list-style-type: none"> • notSupported • supportedOff • supportedOn <p>If the value is supportedOn, the probe adds entries to the Address Map tab that maps the network layer address to the MAC layer address.</p>
HostConfig	<p>Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following:</p> <ul style="list-style-type: none"> • notSupported • supportedOff • supportedOn <p>If the value is supportedOn, the probe adds entries to the Host Control tab to collect statistics for network layer and application layer hosts.</p>

Name	Description
MatrixConfig	Describes and configures the probe support for the network layer and application layer host tables for this protocol. The value can be one of the following: <ul style="list-style-type: none"> notSupported supportedOff supportedOn
Owner	Shows the entity that configured this entry.

View the Data Source for Protocol Distribution Statistics

About This Task

View the Distribution Control tab to see the network segment data source on which the protocol distribution statistics are measured. The management IP mentioned as a data source represents the IP that the SNMP agent uses to access the switch.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Protocol Distribution**.
3. Select the **Distribution Control** tab.

Distribution Control Field Descriptions

Use the data in the following table to use the **Distribution Control** tab.

Name	Description
Index	Shows a unique identifier for the entry in the table.
DataSource	Specifies the source of data for this protocol distribution.
DroppedFrames	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
CreateTime	Shows the value of the sysUpTime when the entry was last activated.
Owner	Shows the entity that configured this entry.

View Protocol Distribution Statistics

About This Task

View protocol distribution statistics to see traffic statistics that each protocol generates by local area network (LAN) segment.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Protocol Distribution**.
3. Select the **Distribution Stats** tab.

Distribution Stats Field Descriptions

Use the data in the following table to use the **Distribution Stats** tab.

Name	Description
LocalIndex	Identifies the protocol distribution an entry is part of, as well as the particular protocol that it represents.
Pkts	Shows the number of packets without errors received for this protocol type. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
Octets	Shows the number of octets in packets received for this protocol type since it was added to the table. This value does not include octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.

View the Host Interfaces Enabled for Monitoring

About This Task

View the entries in the address map control tab to see which host interfaces are enabled for monitoring on the switch. Each entry in this table enables the discovery of addresses on a new interface.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Address Map**.
3. Select the **Address Map Control** tab.

Address Map Control Field Descriptions

Use the data in the following table to use the **Address Map Control** tab.

Name	Description
Index	Shows a unique identifier for the entry in the table.
DataSource	Shows the source of data for the entry.
DroppedFrames	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
Owner	Shows the entity that configured this entry.

View Address Mappings

About This Task

View the mappings of network layer address to physical address to interface.

The probe adds entries on this tab based on the source MAC and network addresses in packets without MAC-level errors.

The probe populates this table for all protocols on the **Protocol Directories** tab with a value of **AddressMapConfig** equal to **supportedOn**.

Procedure

1. In the navigation pane, expand **Configuration** > **Serviceability** > **RMON**.
2. Select **Address Map**.
3. Select the **Address Map** tab.

Address Map Field Descriptions

Use the data in the following table to use the **Address Map** tab.

Name	Description
LocalIndex	Shows a unique identifier for the entry in the table.
HostAddress	Shows the network address for this entry. The format of the value depends on the protocol portion of the local index.
Source	Shows the interface or port on which the network address was most recently seen.

Name	Description
PhysicalAddress	Shows the physical address on which the network address was most recently seen.
LastChange	Shows the value of the sysUpTime when the entry was created or last changed. If this value changes frequently, it can indicate duplicate address problems.

View the Data Source for Host Statistics

About This Task

View the Host Control tab to see the data source for both network layer and application layer host statistics.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Network Layer Host**.
3. Select the **Host Control** tab.

Host Control Field Descriptions

Use the data in the following table to use the **Host Control** tab.

Name	Description
Index	Shows a unique identifier for the entry in the table.
DataSource	Shows the source of data for the associated host table. The statistics in this group reflect all packets on the local network segment that attaches to the identified interface.
NHDropFrames	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.

Name	Description
AHDropFrames	Shows the total number of frames that the probe receives and drops but does not include in the StatsDropEvents value. This event can occur if the probe is out of resources and sheds the load from this collection. This value does not include packets that were not counted because they had MAC-layer errors.
Owner	Shows the entity that configured this entry.

View Network Host Statistics

About This Task

View network host statistics to see Layer 3 traffic statistics for each host. The network layer host MIB monitors traffic packets in and out of hosts based on the network layer address.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Network Layer Host**.
3. Select the **Network Host Stats** tab.

Network Host Stats Field Descriptions

Use the data in the following table to use the **Network Host Stats** tab.

Name	Description
LocalIndex	Shows a unique identifier for the entry in the table.
HostAddress	Shows the host address for this entry.
InPkts	Shows the number of packets without errors transmitted to this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
OutPkts	Shows the number of packets without errors transmitted by this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.

Name	Description
InOctets	Shows the number of octets transmitted to this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
OutOctets	Shows the number of octets transmitted by this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
CreateTime	Shows the value of the sysUpTime when the entry was last activated.

View Application Host Statistics

View application host statistics to see traffic statistics by application protocol for each host.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > RMON**.
2. Select **Application Layer Host**.
3. Select the **Application Host Stats** tab.

Application Host Stats Field Descriptions

Use the data in the following table to use the **Application Host Stats** tab.

Name	Description
Index	Shows a unique identifier for the entry in the table.
HostAddress	Identifies the network layer address of this entry.
LocalIndex	Identifies the network layer protocol of the address.
InPkts	Shows the number of packets for this protocol type, without errors, transmitted to this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.
OutPkts	Shows the number of packets for this protocol type, without errors, transmitted by this address. This value is the number of link-layer packets so a single, fragmented network-layer packet can increment the counter several times.

Name	Description
InOctets	Shows the number of octets transmitted to this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
OutOctets	Shows the number of octets transmitted by this address, excluding octets in packets that contained errors. This value counts octets in the entire packet that contained the protocol, not just the particular protocol frames.
CreateTime	Shows the value of the sysUpTime when the entry was last activated.

RMON Alarm Variables

RMON alarm variables are divided into three categories. Each category can have subcategories.

The topics in this section provide a brief variable description for the following alarm variable categories:

- [RMON Security Alarm Variables](#) on page 3236
- [RMON Error Alarm Variables](#) on page 3236
- [RMON Traffic Alarm Variables](#) on page 3242



Note

In addition to these elements that are offered in a graphical way by EDM, you can manually set any valid OID in the variable field to be monitored by an alarm. For these cases, the name of the variable cannot be translated automatically in OID, the exact OID must be set as a sequence of numbers.

RMON Security Alarm Variables

The following table lists the alarm variables for the security category and provides a brief variable description.

Table 258: RMON security alarm variables

Variable	Definition
rcCliNumAccessViolations.0	The number of CLI access violations detected by the system.
rcWebNumAccessBlocks.0	The number of accesses the web server blocked.
snmpInBadCommunityNames.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed by the SNMP community named in the message.

RMON Error Alarm Variables

The following tables list the alarm variables for subcategories within the Error category.

Table 259: Interface alarm variables

Variable	Definition
ifInDiscards	The number of inbound packets discarded even though no errors were detected to prevent the packets being deliverable to a higher-layer protocol. One possible reason for discarding a packet is to free buffer space.
ifInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors, preventing them from being deliverable to a higher-layer protocol.

Table 259: Interface alarm variables (continued)

Variable	Definition
ifOutDiscards	The number of outbound packets discarded even though no errors were detected to prevent the packets being transmitted. One possible reason for discarding such a packet is to free buffer space.
ifOutErrors	For packet-oriented interfaces, the number of outbound packets that were not transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that were not transmitted because of errors.

Table 260: Ethernet alarm variables

Variable	Definition
dot3StatsAlignmentErrors	A count of frames received on a particular interface that are not an integral number of octets in length and do not pass the Frame Check Sequence (FCS) check. The count represented by an instance of this object increments when the alignmentError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions exist are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
dot3StatsFCSErrors	A count of frames received on a particular interface that are an integral number of octets in length but do not pass the FCS check. The count represented by an instance of this object increments when the frameCheckError status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions occur are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
dot3StatsSingleCollisionFrames	A count of successfully transmitted frames on a particular interface where transmission is inhibited by exactly one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts object, the ifOutMulticastPkts object, or the ifOutBroadcastPkts object, and is not counted by the corresponding instance of the dot3StatsMultipleCollisionFrames object.
dot3StatsMultipleCollisionFrames	A count of successfully transmitted frames on a particular interface where transmission is inhibited by more than one collision. A frame that is counted by an instance of this object is also counted by the corresponding instance of either the ifOutUcastPkts, ifOutMulticastPkts, or ifOutBroadcastPkts object, and is not counted by the corresponding instance of the dot3StatsSingleCollisionFrames object.

Table 260: Ethernet alarm variables (continued)

Variable	Definition
dot3StatsSQETestErrors	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular interface. The SQE TEST ERROR message is defined in section 7.2.2.2.4 of ANSI/IEEE 802.3-1985 and its generation is described in section 7.2.4.6 of the same document.
dot3StatsDeferredTransmissions	A count of frames for which the first transmission attempt on a particular interface is delayed because the medium is busy. The count represented by an instance of this object does not include frames involved in collisions.
dot3StatsLateCollisions	The number of times that a collision is detected on a particular interface later than 512 bit-times into the transmission of a packet; 512 bit-times corresponds to 51.2 microseconds on a 10 Mb/s system. A (late) collision included in a count represented by an instance of this object is also considered as a (generic) collision for purposes of other collision-related statistics.
dot3StatsExcessiveCollisions	A count of frames where the transmission on a particular interface fails due to excessive collisions.
dot3StatsInternalMacTransmitErrors	<p>A count of frames where the transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object.</p> <p>The precise meaning of the count represented by an instance of this object is implementation specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.</p>
dot3StatsCarrierSenseErrors	The number of times the carrier sense condition was lost or never asserted when the switch attempted to transmit a frame on a particular interface. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.

Table 260: Ethernet alarm variables (continued)

Variable	Definition
dot3StatsFrameTooLongs	A count of frames received on a particular interface that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user). Received frames for which multiple error conditions obtained are, according to the conventions of IEEE 802.3 Layer Management, counted exclusively according to the error status presented to the LLC.
dot3StatsInternalMacReceiveErrors	A count of frames where the transmission on a particular interface fails due to an internal MAC sublayer transmit error. A frame is counted by an instance of this object only if it is not counted by the corresponding instance of either the dot3StatsLateCollisions object, the dot3StatsExcessiveCollisions object, or the dot3StatsCarrierSenseErrors object. The precise meaning of the count represented by an instance of this object is implementation-specific. In particular, an instance of this object can represent a count of transmission errors on a particular interface that are not otherwise counted.

Table 261: IP alarm variables

Variable	Definition
ipInHdrErrors.0	The number of input datagrams discarded due to errors in the datagram IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, and errors discovered in processing IP options.
ipInDiscards.0	The number of discarded input IP datagrams where no problems were encountered to prevent continued processing. An example of why they were discarded can be lack of buffer space. This counter does not include any datagrams discarded while awaiting reassembly.
ipOutDiscards.0	The number of output IP datagrams where no problems were encountered to prevent transmission to the destination, but that were discarded (for example, for lack of buffer space). This counter includes datagrams counted in ipForwDatagrams if packets meet this (discretionary) discard criterion.
ipFragFails.0	The number of IP datagrams discarded because they needed to be fragmented at this entity but were not, for example, because the Don't Fragment flag was set.

Table 261: IP alarm variables (continued)

Variable	Definition
ipReasmFails.0	The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so forth). This is not necessarily a count of discarded IP fragments because some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.
icmpInParmProbs.0	The number of ICMP In parameter problem messages received.
icmpOutParmProbs.0	The number of ICMP Out parameter problem messages received.

Table 262: MLT alarm variables

Variable	Definition
rcStatMltEtherAlignmentErrors	The number of frames received on an MLT that are not an integral number of octets in length, but do not pass the FCS check.
rcStatMltEtherFCSErrors	The number of frames received on an MLT that are an integral number of octets in length, but do not pass the FCS check.
rcStatMltEtherSingleCollFrames	The number of successfully transmitted frames on a particular MLT where transmission is inhibited by exactly one collision.
rcStatMltEtherMultipleCollFrames	The number of successfully transmitted frames on a particular MLT where transmission is inhibited by more than one collision.
rcStatMltEtherSQETestError	A count of times that the SQE TEST ERROR message is generated by the PLS sublayer for a particular MLT.
rcStatMltEtherDeferredTransmissions	A count of frames where the first transmission attempt on a particular MLT is delayed because the medium is busy. The count represented by an instance of this object.
rcStatMltEtherLateCollisions	The number of times that a late collision is detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512-bit-times corresponds to 51.2-microseconds on a 10 Mb/s system.
rcStatMltEtherExcessiveCollisions	The number of times that excessive collisions are detected on a particular MLT later than 512 bit-times into the transmission of a packet; 512 bit-times corresponds to 51.2 microseconds on a 10-Mb/s system.

Table 262: MLT alarm variables (continued)

Variable	Definition
rcStatMltEtherMacTransmitError	A count of frames where the transmission on a particular MLT fails due to an internal MAC sublayer transmit error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the LateCollisions object, the ExcessiveCollisions object, or the CarrierSenseErrors object.
rcStatMltEtherCarrierSenseError	The number of times the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular MLT. The count represented by an instance of this object increments at most once for each transmission attempt, even if the carrier sense condition fluctuates during a transmission attempt.
rcStatMltEtherFrameTooLong	A count of frames received on a particular MLT that exceeds the maximum permitted frame size. The count represented by an instance of this object increments when the frameTooLong status is returned by the MAC service to the LLC (or other MAC user).
rcStatMltEtherMacReceiveError	A count of frames for which reception on a particular MLT fails due to an internal MAC sublayer receive error. A frame is only counted by an instance of this object if it is not counted by the corresponding instance of either the FrameTooLongs object, the AlignmentErrors object, or the FCSErrors object.

Table 263: Other alarm variables

Variable	Definition
rcTblArNoSpace	The number of entries not added to the address translation table due to lack of space.
snmplnAsnParseErrs.0	The total number of ASN.1 or BER errors encountered by the SNMP protocol entity when it decodes received SNMP messages.
rcStgPortInBadBpdus	The number of bad BPDUs received by this port.
dot1dTpPortInDiscards	Count of valid frames received that were discarded (that is, filtered) by the forwarding process.

**Note**

In addition to these elements that are offered in a graphical way by EDM, you can manually set any valid OID in the variable field to be monitored by an alarm. For these cases, the name of the variable cannot be translated automatically in OID, the exact OID must be set as a sequence of numbers.

RMON Traffic Alarm Variables

The following tables list the alarm variables for subcategories within the Traffic category.

Table 264: Interface alarm variables

Variable	Definition
ifInOctets	The total number of octets received on the interface, including framing characters.
ifInMulticastPkts	The number of packets, delivered by this sublayer to a higher sublayer, that are addressed to a multicast address at this sublayer. For a MAC layer protocol, this number includes both Group and Functional addresses.
ifInBroadcastPkts	The number of packets, delivered by this sublayer to a higher (sub) layer, that are addressed to a broadcast address at this sublayer.
ifInUnkownProtos	For packet-oriented interfaces, the number of packets received through the interface that are discarded because of an unknown or unsupported protocol. For character-oriented or fixed-length interfaces that support protocol multiplexing, the number of transmission units received through the interface that are discarded because of an unknown or unsupported protocol. For any interface that does not support protocol multiplexing, this counter is always 0.
ifOutOctets	The total number of octets transmitted from the interface, including framing characters.
ifOutMulticastPkts	The total number of packets that higher-level protocols requested be transmitted, and that are addressed to a multicast address at this sublayer, including those that are discarded or not sent. For a MAC layer protocol, this includes both Group and Functional addresses.
ifOutBroadcastPkts	The total number of packets that higher level protocols requested transmitted, and that were addressed to a broadcast address at this sublayer, including those discarded or not sent.
ifLastChange	The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last reinitialization of the local network management subsystem, this object contains a value of zero.

Table 265: RMON Ether Stats alarm variables

Variable	Definition
etherStatsOctets	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets). Use this object as a reasonable estimate of Ethernet utilization. For greater precision, sample the etherStatsPkts and

Table 265: RMON Ether Stats alarm variables (continued)

Variable	Definition
	etherStatsOctets objects before and after a common interval.
etherStatsPkts	The total number of packets (including bad packets, broadcast packets, and multicast packets) received.
etherStatsBroadcastPkts	The total number of good packets received that are directed to the broadcast address. This number does not include multicast packets.
etherStatsMulticastPkts	The total number of good packets received that are directed to a multicast address. This number does not include packets directed to the broadcast address.
etherStatsCRCAlignErrors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of 64 to 1518 octets, inclusive, but had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
etherStatsUndersizePkts	The total number of packets received that are less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
etherStatsOversizePkts	The total number of packets received that are longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
etherStatsFragments	The total number of packets received that are less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). It is entirely normal for etherStatsFragments to increment because it counts both runs (which are normal occurrences due to collisions) and noise hits.
etherStatsCollisions	The best estimate of the total number of collisions on this Ethernet segment.

Table 266: IP alarm variables

Variable	Definition
ipInReceives.0	All incoming IP packets.
ipInAddrErrors.0	The number of bad IP destination addresses.
ipForwDatagrams.0	IP packets forwarded.
ipInUnknownProtos.0	Number of unsupported IP protocols.
ipInDelivers.0	The number of IP In packets delivered.
ipOutRequests.0	The total number of IP datagrams that local IP user protocols supplied to IP in request for transmission.

Table 266: IP alarm variables (continued)

Variable	Definition
ipOutNoRoutes.0	The number of IP datagrams discarded because no route was found to transmit to the destination.
ipFragOKs.0	The number of IP datagrams successfully fragmented.
ipFragCreates.0	The number of IP datagram fragments generated as a result of fragmentation.
ipReasmReqds.0	The number of requests to reassemble fragments.
ipReasmOKs.0	The number of fragments reassembled successfully.

Table 267: ICMP alarm variables

Variable	Definition
icmpInSrcQuenchs.0	The number of ICMP Source Quench messages received.
icmpInRedirects.0	The number of ICMP redirect messages.
icmpInEchos.0	The number of ICMP Echo request messages received.
icmpInEchosReps.0	The number of ICMP Echo reply messages received.
icmpInTimeStamps.0	The number of ICMP timestamp request messages received.
icmpInTimeStampsReps.0	The number of ICMP timestamp reply messages received.
icmpInAddrMasks.0	The number of ICMP mask request messages reviewed.
icmpInAddrMasksReps.0	The number of ICMP mask reply messages reviewed.
icmpInDestUnreachs.0	The number of ICMP destinations unreachable messages received.
icmpInTimeExcds.0	The number of ICMP Time Exceeded messages received.
icmpOutSrcQuenchs.0	The number of ICMP Source Quench messages sent.
icmpOutRedirects.0	The number of ICMP redirect messages sent.
icmpOutEchos.0	The number of ICMP Echo request messages sent.
icmpOutEchosReps.0	The number of ICMP Echo reply messages sent.
icmpOutTimeStamps.0	The number of ICMP Timestamp request messages sent.
icmpOutTimeStampsReps.0	The number of ICMP Timestamp reply messages sent.
icmpOutAddrMasks.0	The number of ICMP Address mask messages sent.
icmpOutAddrMasksReps.0	The number of ICMP Address mask reply messages sent.

Table 267: ICMP alarm variables (continued)

Variable	Definition
icmpOutDestUnreachs.0	The number of ICMP destination unreachable messages sent.
icmpOutTimeExcds.0	The number of ICMP time exceeded messages sent.

Table 268: SNMP alarm variables

Variable	Definition
snmpInPkts.0	The total number of messages delivered to the SNMP entity from the transport service.
snmpOutPkts.0	The total number of SNMP messages passed from the SNMP protocol entity to the transport service.
snmpInBadVersions.0	The total number of SNMP messages delivered to the SNMP protocol entity that were intended for an unsupported SNMP version.
snmpInBadCommunityUses.0	The total number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.
snmpInTooBigs.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
snmpInNoSuchNames.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
snmpInBadValues. 0	The total number of SNMP PDUs received that were generated by the SNMP protocol entity and for which the value of the error-status field is badValue.
snmpInReadOnlys.0	The total number of valid SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is readOnly. It is a protocol error to generate an SNMP PDU that contains the value readOnly in the error-status field; as such, this object is provided as a means of detecting incorrect implementations of the SNMP.
snmpInGenErrs.0	The total number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is genErr.
snmpInTotalReqVars.0	The total number of MIB objects retrieved successfully by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
snmpInTotalSetVars.0	The total number of MIB objects altered successfully by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs.

Table 268: SNMP alarm variables (continued)

Variable	Definition
snmpInGetRequests.0	The total number of SNMP Get-Request PDUs accepted and processed by the SNMP protocol entity.
snmpInGetNexts.0	The total number of SNMP Get-Next PDUs accepted and processed by the SNMP protocol entity.
snmpInSetRequests.0	The total number of SNMP Set-Request PDUs accepted and processed by the SNMP protocol entity.
snmpInGetResponses.0	The total number of SNMP Get-Response PDUs accepted and processed by the SNMP protocol entity.
snmpInTraps.0	The total number of SNMP Trap PDUs accepted and processed by the SNMP protocol entity.
snmpOutTooBigs.0	The total number of SNMP PDUs generated by the SNMP protocol entity and for which the value of the error-status field is tooBig.
snmpOutNoSuchNames.0	The total number of SNMP PDUs generated by the SNMP protocol entity and for which the value of the error-status field is noSuchName.
snmpOutBadValues.0	The total number of SNMP PDUs sent that were generated by the SNMP protocol entity and for which the value of the error-status field is badValue.
snmpOutGenErrs.0	The total number of SNMP PDUs generated by the SNMP protocol entity and for which the value of the error-status field is genErr.
snmpOutGetRequests.0	The total number of SNMP Get-Request PDUs generated by the SNMP protocol entity.
snmpOutGetNexts.0	The total number of SNMP Get-Next PDUs generated by the SNMP protocol entity.
snmpOutSetRequests.0	The total number of SNMP Set-Request PDUs generated by the SNMP protocol entity.
snmpOutGetResponses.0	The total number of SNMP Get-Response PDUs generated by the SNMP protocol entity.
snmpOutTraps.0	The total number of SNMP Trap PDUs generated by the SNMP protocol entity.

Table 269: Bridge alarm variables

Variable	Definition
rcStgTimeSinceTopologyChange	The time (in hundredths of a second) since the last topology change was detected by the bridge entity.
rcStgTopChanges	The total number of topology changes detected by this bridge since the management entity was last reset or initialized.

Table 269: Bridge alarm variables (continued)

Variable	Definition
rcStgMaxAge	The maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded, in hundredths of a second. This is the actual value that this bridge is currently using.
rcStgPortForwardTransitions	The number of times this port transitioned from the Learning state to the Forwarding state.
rcStgPortInConfigBpdus	The number of Config BPDUs received by this port.
rcStgPortInTcnBpdus	The number of Topology Change Notification BPDUs received by this port.
rcStgPortOutConfigBpdus	The number of Config BPDUs transmitted by this port.
rcStgPortOutTcnBpdus	The number of Topology Change Notification BPDUs transmitted by this port.
dot1dTpPortInFrames	The number of frames received by this port from its segment. A frame received on the interface corresponding to this port is counted by this object only if it is for a protocol being processed by the local bridging function, including bridge management frames.
dot1dTpPortOutFrames	The number of frames transmitted by this port to its segment. A frame transmitted on the interface corresponding to this port is counted by this object if and only if it is for a protocol processed by the local bridging function, including bridge management frames.
dot1dTpLearnedEntryDiscards.0	The total number of Forwarding Database entries learned but discarded due to a lack of space to store them in the Forwarding Database. If this counter increases, it indicates that the forwarding database is regularly becoming full (a condition that has negative performance effects on the subnetwork). If this counter has a significant value but does not increase, it indicates that the problem occurred but is not persistent.

Table 270: Utilization alarm variables

Variable	Definition
rcSysBufferUtil.0	Buffer utilization as a percentage of the total amount of buffer space in the system. A high value indicates congestion.
rcSysNVRamUsed.0	Nonvolatile RAM (NVRAM) in use in kilobytes.

Table 270: Utilization alarm variables (continued)

Variable	Definition
rcSysLastChange.0	Last management-initiated configuration change since sysUpTime.
rcSysLastVlanChange.0	Last management-initiated VLAN configuration change since sysUpTime.

Table 271: MLT alarm variables

Variable	Definition
rcStatMltIfExtnIfInMulticastPkts	The total number of multicast packets delivered to this MLT interface.
rcStatMltIfExtnIfInBroadcastPkts	The total number of broadcast packets delivered to this MLT Interface.
rcStatMltIfExtnIfOutMulticastPkts	The total number of MLT interface multicast packets delivered to this MLT interface.
rcStatMltIfExtnIfOutBroadcastPkts	The total number of MLT interface broadcast packets delivered to this MLT interface.
rcStatMltIfExtnIfHCInOctets	The total number of octets received on this MLT interface including framing characters detected by the high-count (64-bit) register.
rcStatMltIfExtnIfHCInUcastPkts	The number of packets delivered by this MLT interface to a higher MLT that were not addressed to a multicast or broadcast address as detected by the high-count (64-bit) register.
rcStatMltIfExtnIfHCInMulticastPkt	The total number of multicast packets delivered to this MLT interface detected by the high-count (64-bit) register.
rcStatMltIfExtnIfHCInBroadcastPkt	The total number of broadcast packets delivered to this MLT interface detected by the high-count (64-bit) register.
rcStatMltIfExtnIfHCOctets	The total number of octets transmitted from the MLT interface, including framing characters.
rcStatMltIfExtnIfHCOOutUcastPkts	The number of packets transmitted by this MLT interface to a higher MLT that were not addressed to a multicast or broadcast address as detected by the high-count (64-bit) register.
rcStatMltIfExtnIfHCOOutMulticast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a multicast address at this sublayer, including those that were discarded or not sent registered by the high-count (64-bit) register.
rcStatMltIfExtnIfHCOOutBroadcast	The total number of packets that higher-level protocols requested be transmitted, and that were addressed to a broadcast address at this sublayer, including those that were discarded or not sent registered by the high-count (64-bit) register.



Route Filtering and IP Policies

[Route Filtering and IP Policies](#) on page 3249

[Prefix list](#) on page 3253

[Route Policy Definition](#) on page 3254

[IP Policy Configuration Using the CLI](#) on page 3259

[IP Policy Configuration using Enterprise Device Manager](#) on page 3273

Table 272: IP Route Policies product support

Feature	Product	Release introduced
IP route policies	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Route Filtering and IP Policies

When the switch routes IP traffic, you can apply a number of filters to manage, accept, redistribute, and announce policies for unicast routing table information. Filters apply differently to different unicast routing protocols.



Note

IPv6 ingress QoS ACL/Filters and IPv6 Egress Security and QoS ACL/Filters are not supported. For information on the maximum number of IPv6 ingress port/vlan security ACL/filters supported on the switch, see *Fabric Engine Release Notes*.

The following figure shows how filters apply to BGP, RIP, and OSPF protocols.

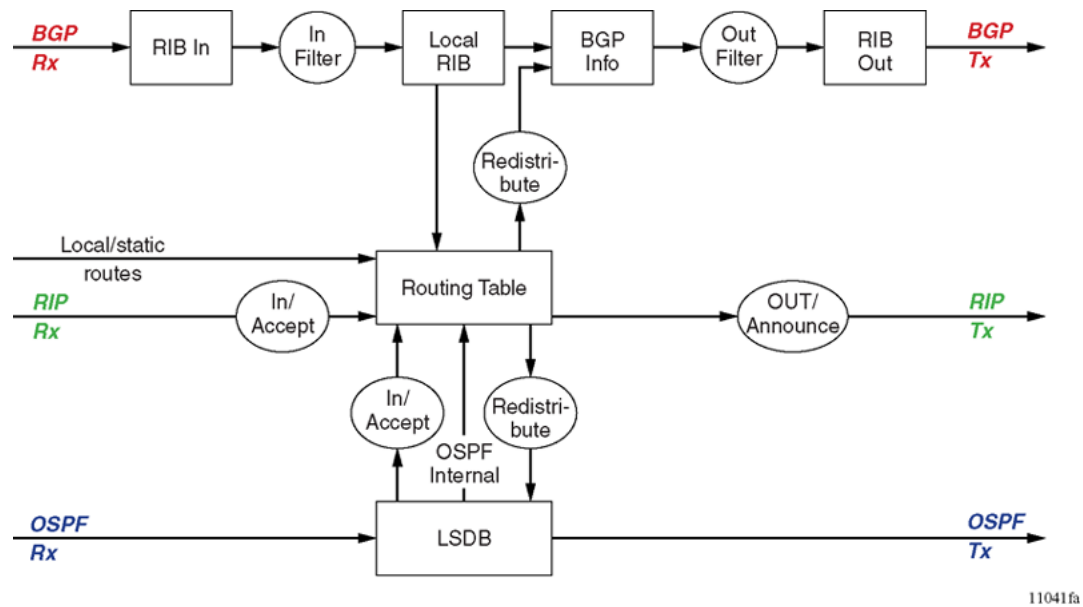


Figure 213: Route filtering for BGP, RIP, and OSPF routing protocols

The following figure shows how filters apply to the IS-IS protocol for Fabric Connect Layer 3 VSNs or IP Shortcuts.

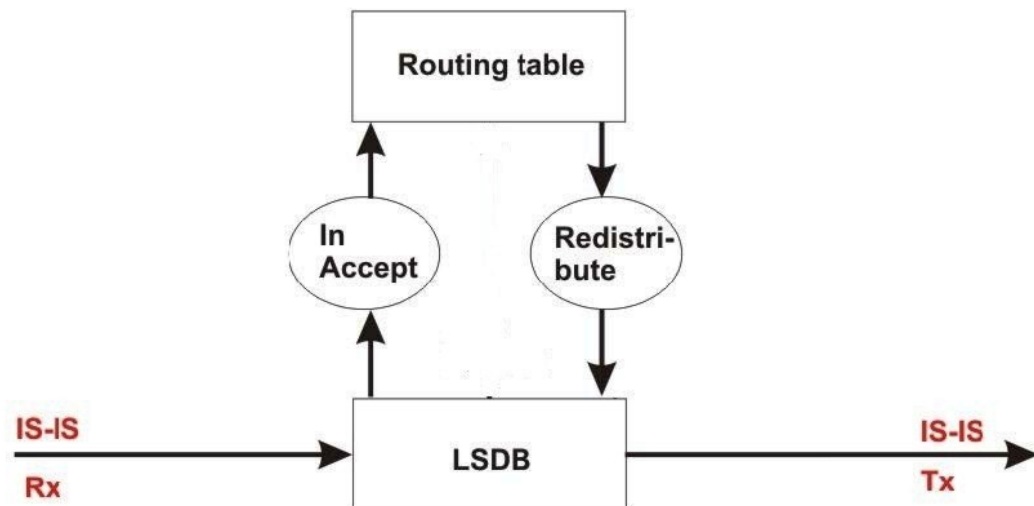


Figure 214: Route filtering for the IS-IS routing protocol

Accept Policies

Accept policies are applied to incoming traffic to determine whether to add the route to the routing table. Accept policies are applied differently to protocols, as follows:

- RIP and BGP—filters apply to all incoming route information.

- OSPF—filters apply only to external route information. Internal routing information is not filtered because otherwise, other routers in the OSPF domain can have inconsistent databases that can affect the router view of the network topology.
- IS-IS —filters apply to all incoming route information.

In a network with multiple routing protocols, you can prefer specific routes from RIP instead of from OSPF. The network prefix is a commonly used match criterion for accept policies.

Redistribution Filters

Redistribution filters notify changes in the route table to the routing protocol (within the device). With redistribution filters, providing you do not breach the protocol rules, you can choose not to advertise everything that is in the protocol database, or you can summarize or suppress route information. By default, no external routes are leaked to protocols that are not configured.

Announce Policies

Announce policies are applied to outgoing advertisements to neighbors or peers in the protocol domain to determine whether to announce specific route information. Out filtering applies to RIP updates and BGP NLRI updates.

In contrast, announce policies are not applied to IS-IS or OSPF information because routing information must always be consistent across the domain. To restrict the flow of external route information in the IS-IS or OSPF protocol database, apply redistribution filters instead of announce policies.

Route Filtering Stages

The following figure shows the three distinct filter stages that are applied to IP traffic.

These stages are:

- Filter stage 1 is the accept policy or in filter that applies to incoming traffic to detect changes in the dynamic (protocol-learned) routing information, which are then submitted to the routing table.
- Filter stage 2 is the redistribution filter that applies to the entries in the routing table to the protocol during the leaking process.
- Filter stage 3 is the announce policy or out filter that applies to outgoing traffic within a protocol domain.

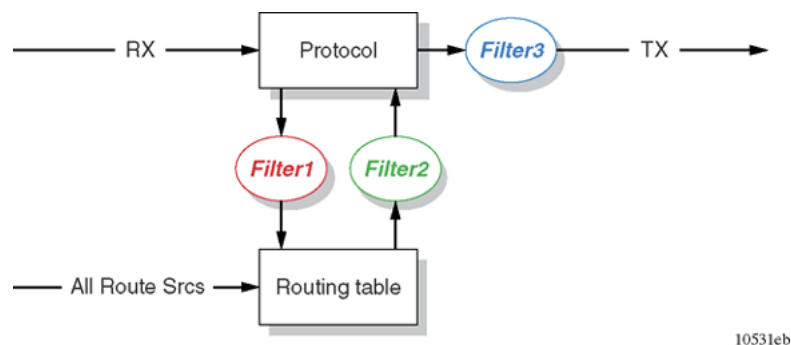


Figure 215: Route filtering stages

The following figure shows the logical process for route filtering on the switch.




Figure 216: Route filtering logic

Prefix list

In the switch software, you can create one or more IP prefix lists and apply these lists to IP route policy.

Route Policy Definition

You can define an IP route policy and its attributes globally, and then apply them individually to interfaces and protocols. You can also form a unified database of route policies that the RIP or OSPF protocol can use for type of filtering purpose. A name or ID identifies a policy.

Under a policy you can have several sequence numbers. If you do not configure a field in a policy, the system displays the field as 0 in CLI show command output. This value indicates that the device ignores the field in the match criteria. Use the clear option to remove existing configurations for the field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference field, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce or redistribute purposes.

You can only apply one policy for each purpose (RIP Announce, for example) on a given RIP interface. In this case, all sequence numbers under the policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

The following tables display the accept, announce, and redistribute policies for RIP, OSPF, IS-IS and BGP. The tables also display which matching criteria apply for a certain routing policy. Superscript numbers in the following table identify the following:

1. denotes advertise router
2. denotes RIP gateway
3. denotes that external type 1 and external type 2 are the only options



Note

IPv4 and IPv6 route-maps cannot be configured on the same match statement.

Table 273: Protocol route policy table for RIP

	Announce				Accept
	OSPF	Direct	RIP	BGP	RIP
Match Protocol	Yes	Yes	Yes	Yes	
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source	Yes ¹		Yes ²		
Match NextHop	Yes	Yes	Yes	Yes	Yes
Match Interface			Yes		
Match Route Type	Yes				
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					

Table 273: Protocol route policy table for RIP (continued)

	Announce				Accept
	OSPF	Direct	RIP	BGP	RIP
Match Community					
Match Community Exact					
MatchTag				Yes	
NssaPbit					
SetRoute Preference					Yes
SetMetric TypeInternal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type					
SetNextHop					
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					Yes
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 274: Protocol route policy table for OSPF

	Redistribute					Accept
	Direct	Static	RIP	BGP	IS-IS	OSPF
Match Protocol				Yes	Yes	
Match Network	Yes	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes ²			
Match NextHop		Yes	Yes	Yes		
Match Interface			Yes			
Match Route Type					Yes ³	
Match Metric	Yes	Yes	Yes	Yes	Yes	Yes
MatchAs Path						

Table 274: Protocol route policy table for OSPF (continued)

	Redistribute					Accept
	Direct	Static	RIP	BGP	IS-IS	OSPF
Match Community						
Match Community Exact						
MatchTag				Yes	Yes	
Set NSSA Bit	Yes	Yes	Yes	Yes	Yes	
SetRoute Preference						
SetMetric TypeInternal						
SetMetric	Yes	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	Yes	
SetNextHop				Yes		
Set Inject NetList	Yes	Yes	Yes	Yes	Yes	Yes
SetMask						
SetAsPath						
SetAsPath Mode						
Set Automatic Tag						
Set CommunityNumber						
Set CommunityMode						
SetOrigin						
SetLocal Pref						
SetOrigin EgpAs						
SetTag						
SetWeight						

Table 275: Protocol route policy table for IS-IS

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
Match Protocol	Yes	Yes	Yes	Yes	Yes
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes		
Match NextHop		Yes	Yes	Yes	Yes
Match Interface			Yes		
Match Route Type					Yes ³
Match Metric	Yes	Yes	Yes	Yes	Yes

Table 275: Protocol route policy table for IS-IS (continued)

	Redistribute				Accept
	Direct	Static	RIP	BGP	OSPF
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
Set NSSA Bit					
SetRoute Preference					
SetMetric Type Internal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	Yes
SetNextHop				Yes	
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 276: Protocol route policy table for BGP

	Redistribute				Accept	Announce
	IPv6 Direct	IPv6 Static	IS-IS	OSPFv3	BGP	BGP
Match as-path					Yes	Yes
Match community	Yes	Yes		Yes	Yes	Yes
Match community-exact					Yes	Yes
Match extcommunity					Yes	Yes
Match interface						

Table 276: Protocol route policy table for BGP (continued)

	Redistribute				Accept	Announce
	IPv6 Direct	IPv6 Static	IS-IS	OSPFv3	BGP	BGP
Match local-preference						
Match metric	Yes	Yes	Yes	Yes	Yes	Yes
Match network	Yes	Yes	Yes	Yes	Yes	Yes
Match next-hop		Yes		Yes	Yes	Yes
Match protocol			Yes			
Match route-source					Yes	
Match route-type			Yes ³	Yes		Yes
Match tag			Yes			
Match vrf						
Match vrfids						
Set as-path					Yes	Yes
Set as-path-mode					Yes	Yes
Set automatic-tag						
Set community					Yes	Yes
Set community-mode					Yes	Yes
Set injectlist	Yes	Yes	Yes	Yes		
Set ip-preference						
Set local-preference					Yes	Yes
Set mask						
Set metric	Yes	Yes	Yes	Yes	Yes	Yes
Set metric-type			Yes			
Set metric-type-internal						
Set next-hop					Yes	Yes
Set nssa-pbit			Yes			
Set origin						Yes
Set origin-egp-as						
Set Tag						
Set Weight					Yes	

Permit and Deny Rules for Inbound and Outbound traffic

Route policies treat permit and deny rules differently for inbound and outbound traffic:

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- To permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

IP Policy Configuration Using the CLI

Configure IP policies to form a unified database of route policies that Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) can use for filtering tasks.

A policy is identified by a name or an ID. Under a given policy you can have several sequence numbers, each of which is equal to one policy in the old convention. Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference field, use only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce and redistribute purposes.

You can apply one policy for one purpose, for example, RIP announce on a RIP interface. All sequence numbers under the given policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

Configuring prefix lists

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

About This Task



Important

When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a prefix list:

```
ip prefix-list WORD<1-64> {A.B.C.D/X} [ge <0-32>] [id <1-2147483647>]
[le <0-32>]
```

3. (Optional) Rename an existing prefix list:

```
ip prefix-list WORD<1-64> name WORD<1-64>
```

4. Display the prefix list:

```
show ip prefix-list [prefix {A.B.C.D}] [vrf WORD<1-16>] [vrfids
WORD<0-512>] [WORD <1-64>]
```

Example

Configure a prefix-list. Display the prefix list.

```
Switch> enable
Switch# configure terminal
Switch(config)# ip prefix-list LIST1 192.0.2.1/255.255.255.0
Switch(config)# show ip prefix-list LIST1
=====
                        Prefix List - GlobalRouter
=====

```

	PREFIX	MASK	LEN	FROM	TO
List 1	LIST1:				
	192.0.1.2.1	24	24	24	

```
1 Total Prefix List entries configured
=====
Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
```

Variable definitions

The following table defines parameters for the **ip prefix-list** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats: <ul style="list-style-type: none"> • a.b.c.d/x • a.b.c.d/x.x.x.x • default
ge <0-32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
id <1-2147483647>	Specifies the Prefix list ID.

Variable	Value
<i>le</i> <0-32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
<i>name</i> WORD<1-64>	Renames the specified prefix list. The name length is 1-64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

Use the data in the following table to use the **show ip prefix-list** command.

Variable	Value
{A.B.C.D}	Specifies the prefix to include in the command output.
<i>vrf</i> WORD<1-16>	Specifies the name of the VRF.
<i>vrfids</i> WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0-512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

The following table defines parameters for the **show ip prefix-list** command output.

Variable	Value
PREFIX	Indicates the member of a specific prefix list.
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
TO	Indicates the prefix mask endpoint in bits.

Configure IP Route Policies

Configure a route policy so that the device can control routes that certain packets can take. For example, you can use a route policy to deny certain Border Gateway Protocol (BGP) routes.

The route policy defines the matching criteria and the actions taken if the policy matches.

About This Task

After you create and enable the policy, you can apply it to an interface. You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In this case, all sequence numbers under the given policy apply to that filter.

Create and enable the policy for IS-IS accept policies for Fabric Connect for Layer 3 Virtual Services Networks (VSNs) and IP Shortcuts, then apply the IS-IS accept policy filters.

**Note**

After you create a route-map in Global Configuration mode or VRF Router Configuration mode, the device enters Route-Map Configuration mode, where you configure the action the policy takes, and define other fields the policy enforces.

**Note**

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

**Note**

You cannot configure IPv4 and IPv6 route-maps on the same match statement.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`

**Note**

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create a route-map policy by giving it a name and sequence number:

```
route-map WORD<1-64> <1-65535> [permit | deny]
```

After you enter the command, the switch automatically transitions to Route-Map Configuration mode for the new policy where you can complete the policy configuration. By default, a new policy permits the route.

3. Define the match criteria for the policy:

```
match {as-path WORD<0-256> | community WORD<0-256> | community-
exact enable | extcommunity WORD<0-1027> | interface WORD<0-259> |
local-preference <0-2147483647> |metric <0-65535> | metric-type-isis
<any|internal|external> | multicast-group WORD<0-259> | multicast-
source WORD<0-259> | network WORD<0-259> | next-hop WORD<0-259> |
protocol WORD<0-60> | route-source WORD<0-259> | route-type <any|
local|internal|external|external-1|external-2>| tag WORD<0-256> | vrf
WORD<1-16> | vrfids WORD<0-512> }
```

4. Define the action the policy takes, if not done at time of policy creation:

- Permit the route:

```
permit
```

- Ignore the route:

```
no permit
```

5. (Optional) Define the set criteria for the policy:

```
set {as-path WORD<0-256> | as-path-mode <tag|prepend> | automatic-
tag enable | community WORD<0-256> | community-mode <additive|none|
unchanged>| data-isid <1-15999999> | injectlist WORD<0-1027> | ip-
preference <0-255> | local-preference <0-2147483647> | mask <A.B.C.D>
| metric <0-65535> | metric-type <type1|type2> | metric-type-internal
<0-1> | metric-type-isis <none|internal|external>| metric-type-live-
metric | next-hop WORD<0-256> | nssa-pbit enable | origin <igp|egp|
incomplete> | origin-egp-as <0-65535>| rx-only | tag WORD<0-256> | tx-
only | weight <0-65535> }
```

6. Enable the policy:

```
enable
```

7. Display current information about the IP route policy:

```
show route-map [WORD<1-64>] [<1-65535>] [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```

Example

Enter Route-Map Configuration mode for the named policy. At the route-map prompt, define the fields the policy enforces. Define the action the policy takes. Display current information about the IP route policy.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map RedisStatic 1
Switch:1(route-map)#match metric 0
Switch:1(route-map)#permit
Switch:1(route-map)#show route-map RedisStatic
=====
Route Policy - GlobalRouter
=====
```

NAME	SEQ	MODE	EN
RedisStatic	1	PRMT	DIS

Variable Definitions

The following table defines the parameters for the **match** command.

Variable	Value
<i>as-path</i> WORD<0-256>	Configures the device to match the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified AS-lists. This field is used only for BGP routes and ignored for all other route types. Specify the list IDs of up to four AS-lists, separated by a comma.
<i>community</i> WORD<0-256>	Configures the device to match the community attribute of the BGP routes against the contents of the specified community lists. This field is used only for BGP routes and ignored for all other route types. Specify the list IDs of up to four defined community lists, separated by a comma.
<i>community-exact</i> enable	When disabled, configures the device so match community-exact results in a match when the community attribute of the BGP routes match an entry of a community-list specified in match-community. When enabled, configures the device so match-community-exact results in a match when the community attribute of the BGP routes matches all of the entries of all the community lists specified in match-community. enable enables match community-exact.
<i>extcommunity</i> WORD <0-1027>	Configures the device to match the extended community. Specify an integer value that represents the community list ID you want to create or modify.
<i>interface</i> WORD <0-259>	If configured, configures the device to match the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other route types. Specify the name of up to four defined prefix lists, separated by a comma.
<i>local-preference</i> <0-2147483647>	Configures the device to match the local preference, applicable to all protocols.
<i>metric</i> <0-65535>	Configures the device to match the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored. The default is 0.

Variable	Value
<code>multicast-group WORD<0-259></code>	Configures the device to match the multicast destination group address for the routed multicast policy. Specify the name of one or more prefix lists, separated by a comma..
<code>multicast-source WORD<0-259></code>	Configures the device to match the multicast source address for the routed multicast policy. Specify the name of one or more prefix lists, separated by a comma..
<code>network WORD <0-259></code>	Configures the device to match the destination network against the contents of the specified prefix lists. Specify the name of up to four defined prefix lists, separated by a comma.
<code>next-hop WORD<0-259></code>	Configures the device to match the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes. Specify the name of up to four defined prefix lists, separated by a comma.
<code>protocol WORD<0-60></code>	Configures the device to match the protocol through which the route is learned. <code>WORD <0-60></code> is xxx, where xxx is local, ospf, ebgp, ibgp,isis, rip, static, or a combination separated by a vertical bard ().
<code>route-source WORD<0-259></code>	Configures the system to match the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types. Specify the name of up to four defined prefix lists, separated by a comma.
<code>route-type {any local internal external external-1 external-2}</code>	Configures a specific route type to match (applies only to OSPF routes).
<code>tag WORD<0-256></code>	Specifies a list of tags used during the match criteria process. Contains one or more tag values. This parameter applies to BGP and OSPF.
<code>[vrf WORD<1-16>] [vrfs WORD<0-512>]</code>	Configures a specific VRF to match (applies only to RIP routes).

Use the data in the following table to use the **set** command.

Variable	Value
<i>as-path</i> WORD<0-256>	Configures the device to add the AS number of the AS-list to the BGP routes that match this policy. Specify the list ID of up to four defined AS-lists separated by a comma.
<i>as-path-mode</i> <tag prepend>	Configures the AS path mode. Prepend is the default configuration. The device prepends the AS number of the AS-list specified in set-as-path to the old as-path attribute of the BGP routes that match this policy. Note: Prepend does not apply to an internal BGP (iBGP) peer with outbound route policy. For more information about iBGP, see BGP on page 536.
<i>automatic-tag</i> enable	Configures the tag automatically. Used for BGP routes only.
<i>community</i> WORD<0-256>	Configures the device to add the community number of the community list to the BGP routes that match this policy. Specify the list ID of up to four defined community lists separated by a comma.
<i>community-mode</i> <additive none unchanged>	Configures the community mode. <i>additive</i> —the device prepends the community number of the community list specified in set-community to the old community path attribute of the BGP routes that match this policy. <i>none</i> —the device removes the community path attribute of the BGP routes that match this policy to the specified value.
<i>data-isid</i>	Configures the data service instance identifier (I-SID) for routed multicast policy.
<i>injectlist</i> WORD<0-1027>	Configures the device to replace the destination network of the route that matches this policy with the contents of the specified prefix list.
<i>ip-preference</i> <0-255>	Configures the preference. This applies to accept policies only.
<i>local-preference</i> <0-65535>	Configures the device to match the local preference, applicable to all protocols.
<i>mask</i> <A.B.C.D>	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.

Variable	Value
<i>metric</i> <0-65535>	Configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF for RIP, the original cost of the route or default-import-metric is used (applies to IS-IS routes also).
<i>metric-type</i> {type1 type2}	Configures the metric type for the routes to announce into the OSPF domain that matches this policy. The default is type 2. This field applies only for OSPF announce policies.
<i>metric-type-internal</i> <0-1>	Configures the MED value for routes advertised to ebgp neighbors to the IGP metric value.
<i>metric-type-isis</i> <none internal external>	Configures the metric type for IS-IS routes. The default is none. This field applies only for IS-IS policies.
<i>metric-type-live-metric</i>	Configures the metric type for BGP routes. The default is disabled. This field applies only for BGP policies.
<i>next-hop</i> WORD <1-256>	Specifies the IP address of the next-hop router. Both IPv4 and IPv6 addresses are supported.
<i>nssa-pbit</i> enable	Configures the not so stubby area (NSSA) translation P bit. Applicable to OSPF announce policies only.
<i>origin</i> {igp egp incomplete}	Configures the device to change the origin path attribute of the BGP routes that match this policy to the specified value.
<i>origin-egp-as</i> <0-65535>	Indicates the remote autonomous system number. Applicable to BGP only.
<i>rx-only</i>	Configures rx only for routed multicast policy.
<i>tag</i> <0-65535>	Configures the tag of the destination routing protocol. If not specified, the device forwards the tag value in the source routing protocol. A value of 0 indicates that this parameter is not configured. This parameter applies to BGP.
<i>tx-only</i>	Configures tx only for routed multicast policy.
<i>weight</i> <0-65535>	Configures the weight value for the routing table. For BGP, this value overrides the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. Used for BGP only. A value of 0 indicates that this parameter is not configured.

Use the data in the following table to use the **name** command.

Variable	Value
<i>WORD<1-64></i>	Renames a policy and changes the name field for all sequence numbers under the given policy.

Configure a Policy to Accept External Routes from a Router

Perform this procedure to configure a policy to accept external routes from a specified advertising router.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router ospf
```

2. Create a policy to accept external routes from a specified advertising route:

```
accept adv-rtr <A.B.C.D>
```

3. Exit to the Privileged EXEC mode.

```
exit
```

4. Apply the OSPF accept policy change:

```
ip ospf apply accept adv-rtr <A.B.C.D>
```

5. Confirm your configuration:

```
show ip ospf accept
```

Example

Log on to the OSPF Router Configuration mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
```

Create a policy to accept external routes from a specified advertising route:

```
Switch:1(config-ospf)#accept adv-rtr 192.0.2.122
```

Enable an OSPF accept entry for a specified advertising route:

```
Switch:1(config-ospf)#accept adv-rtr 192.0.2.122 enable
```

Exit to the Privileged EXEC mode:

```
Switch:1(config-ospf)#exit
Switch:1(config)#exit
```

Apply the OSPF accept policy change and confirm your configuration:

```
Switch:1#ip ospf apply accept adv-rtr 192.0.2.122
Switch:1#show ip ospf accept
=====
Ospf Accept - GlobalRouter
=====
ADV_RTR      MET_TYPE ENABLE POLICY
-----
192.0.2.122  -      FALSE
```

Variable Definitions

Use the data in the following table to use the **accept adv-rtr** command.

Variable	Value
<A.B.C.D>	Specifies the IP address.
<i>enable</i>	Enables an OSPF accept entry for a specified advertising router.
<i>metric-type {any type1 type2 }</i>	Indicates the OSPF external type. This parameter describes which types of OSPF external routes match this entry. <i>any</i> means match all external routes. <i>type1</i> means match external type 1 only. <i>type2</i> means match external type 2 only.
<i>route-map WORD<0-64></i>	Specifies the name of the route policy to use for filtering external routes advertised by the specified advertising router before accepting into the routing table.

Apply OSPF Accept Policy Changes

Apply OSPF accept policy changes to enable the configuration changes in the policy to take effect in an OSPF Accept context (and to prevent the device from attempting to apply the changes one by one after each configuration change).

About This Task



Important

Changing OSPF Accept contexts is a process-oriented operation that can affect system performance and network accessibility while you perform the procedures. If you want to change the default preferences for an OSPF Accept or a prefix-list configuration (as opposed to the default preference), do so before enabling the protocols.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Apply an OSPF accept policy change:
`ip ospf apply accept [vrf WORD<1-16>]`
3. Display information about the configured OSPF entries:
`show ip ospf accept [vrf WORD<1-16>] [vrfs WORD<0-512>]`

Example

Apply the OSPF accept policy and confirm the configuration:

```
Switch:1>enable
Switch:1#ip ospf apply accept
Switch:1>show ip ospf accept
=====
Ospf Accept - GlobalRouter
=====
ADV_RTR      MET_TYPE ENABLE POLICY
-----
192.0.2.11   type1    true   test1
```

Variable Definitions

Use the data in the following table to use the **ip ospf apply accept adv-rtr** command.

Variable	Value
<i>adv-rtr</i>	Commits entered changes. Issue this command after you modify a policy configuration that affects an OSPF accept policy.
<i>vrf WORD<1-16></i>	Specifies the name of the VRF.

Configure Inter-VRF Redistribution Policies

Configure redistribution entries to allow a protocol to announce routes of a certain source type, for example, static, RIP, or direct.

Before You Begin

- Ensure the routing protocols are globally enabled.
- You must configure the route policy, if required.
- Ensure the VRFs exist.
- You must create the route policy and prefix list under the source VRF context.



Note

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the redistribution instance:

```
ip <bgp|ospf|rip> redistribute <bgp|direct|ipv6-direct|ipv6-isis|ipv6-
static|isis|ospf|ospfv3|rip|static|dvr>
```

3. Apply a route policy if required:

```
ip <bgp|ospf|rip> redistribute <bgp|direct|ipv6-direct|ipv6-isis|ipv6-
static|isis|ospf|ospfv3|rip|static|dvr> route-map <WORD 0-64> [vrf-src
<WORD 1-16>]
```

4. Use the following variable definitions table to configure other parameters as required.

5. Enable the redistribution:

```
ip <bgp|ospf|rip> redistribute <bgp|direct|ipv6-direct|ipv6-isis|ipv6-
static|isis|ospf|ospfv3|rip|static|dvr> enable [vrf-src <WORD 1-16>]
```

6. Ensure that the configuration is correct:

```
show ip <bgp|ospf|rip> redistribute [vrf WORD<1-16>] [vrfids
WORD<0-512>]
```

For RIPng, use `show ipv6 rip redistribute`.

7. Apply the redistribution:

```
ip <bgp|ospf|rip> apply redistribute <bgp|direct|ipv6-direct|ipv6-
isis|ipv6-static|isis|ospf|ospfv3|rip|static|dvr> [vrf WORD<1-16>]
[vrf-src WORD<1-16>]
```

Example

```
Switch:1>enable
Switch:1#config terminal
```

Log on to the VRF Router Configuration mode:

```
Switch:1(config)#router vrf test
```

Create the redistribution instance:

```
Switch:1(router-vrf)#ip rip redistribute ospf
```

Enable the redistribution

```
Switch:1(router-vrf)#ip rip redistribute ospf enable
```

Ensure that the configuration is correct:

```
Switch:1(router-vrf)#show ip rip redistribute
```

Exit to Global Configuration mode:

```
Switch:1(router-vrf)#exit
```

Apply the redistribution:

```
Switch:1(config)#ip rip apply redistribute ospf
```

Variable definitions

Use the data in the following table to use the redistribution commands.

Variable	Value
<code><bgp direct isis ospf rip static dvr></code>	Specifies the type of routes to redistribute—the protocol source.
<code>vrf WORD<1-16></code>	Specifies the VRF instance.
<code>vrfids WORD<0-512></code>	Specifies a list of VRF IDs.
<code>vrf-src WORD<1-16></code>	Specifies the source VRF instance. This parameter is not required for redistribution within the same VRF.

Use the data in the following table to use the **ip <bgp|ospf|rip> redistribute <bgp|direct|isis|ospf|rip|static|dvr>** command.

Variable	Value
<code>apply [vrf-src WORD<1-16>]</code>	Applies the redistribution configuration.
<code>enable [vrf-src WORD<1-16>]</code>	Enables the OSPF route redistribution instance.
<code>metric <metric-value> [vrf-src WORD<1-16>]</code>	Configures the metric to apply to redistributed routes.
<code>metric-type <type1 type2> [vrf-src WORD<1-16>]</code>	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
<code>route-map <WORD 0-64> [vrf-src WORD<1-16>]</code>	Configures the route policy to apply to redistributed routes.
<code>subnets <allow suppress> [vrf-src WORD<1-16>]</code>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.

IP Policy Configuration using Enterprise Device Manager

You can form a unified database of route policies that the protocols (RIP, OSPF or Border Gateway Protocol [BGP]) can use for any type of filtering task.

For information about configuring a prefix list, see [Configure a Prefix List](#) on page 3274. For more information about community list, see [Configure a Community Access List](#) on page 653. For more information about AS path list, see [Configure an AS Path List](#) on page 652.

A name or an ID identifies a policy. Under a policy you can have several sequence numbers, each of which is equal to one policy in the old convention. If a field in a policy is not configured, the system displays it as 0 or any when the system displays it in Enterprise Device Manager (EDM). This means that the field is ignored in the match criteria. You can use the clear option to remove existing configurations for any field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a set-preference field set, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce and redistribute purposes.

You can apply only one policy for one purpose (for example, RIP Announce on a given RIP interface). In that example, all sequence numbers under the given policy are applicable for that filter. A sequence number also acts as an implicit preference: a lower sequence number is preferred.

Configure a Prefix List

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or non-contiguous routes. Reference prefix lists by name from within a routing policy.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Policy**.
3. Select the **Prefix List** tab.
4. Select **Insert**.
5. For **Id**, type an ID for the prefix list.
6. For **Prefix**, type an IP address for the route.
7. For **PrefixMaskLength**, type the length of the prefix mask.
8. Configure the remaining parameters as required.
9. Select **Insert**.

Prefix List Field Descriptions

Use the data in the following table to use the **Prefix List** tab.

Name	Description
Id	Configures the list identifier.
Prefix	Configures the IP address of the route.
PrefixMaskLen	Configures the specified length of the prefix mask. You must enter the full 32-bit mask to exact a full match of a specific IP address, for example, if you create a policy to match on the next hop.

Name	Description
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name length can use from 1 to 64 characters.
MaskLenFrom	Configures the lower bound of the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Configures the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Configure a Route Policy

Configure a route policy so that all protocols use them for In, Out, and Redistribute purposes.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand: **Configuration > IP**.
2. Select **Policy**.
3. Select the **Route Policy** tab.
4. Select **Insert**.
5. Enter the appropriate information for your configuration in the Insert Route Policy dialog box.
6. Select **Insert**.

Route Policy Field Descriptions

Use the data in the following table to use the **Route Policy** tab.

Name	Description
Id	Specifies the ID of an entry in the Prefix list table.
SequenceNumber	Specifies a policy within a route policy group.
Name	Specifies the name of the policy. This command changes the name field for all sequence numbers under the given policy.

Name	Description
Enable	Indicates whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored. The default is disabled.
Mode	Specifies the action to take when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route). The default is permit.
MatchProtocol	Selects the appropriate protocol. If configured, matches the protocol through which the route is learned. This parameter applies only for RIP Announce purposes. The default is to enable all match protocols.
MatchNetwork	Specifies if the system matches the destination network against the contents of the specified prefix list.
MatchIpRouteSource	<p>Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.</p> <p>Select the ellipsis and choose from the list in the Match Route Source dialog box. You can select up to four entries. To clear an entry, use the ALT key.</p> <p>You can also change this field in the Route Policy tab of the Policy dialog box.</p>
MatchIpRouteDest	Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
MatchNextHop	<p>Specifies if the system matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.</p> <p>Select the ellipsis and choose from the list in the Match Next Hop dialog box. You can select up to four entries. To clear an entry, use the ALT key.</p>

Name	Description
MatchInterface	Specifies if the system matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route. Select the ellipsis and choose from the list in the Match Interface dialog box. You can select up to four entries. To clear an entry, use the ALT key.
MatchRouteType	Configures a specific route type to match (applies only to OSPF routes). Externaltypel and Externaltypel2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes. The default is any.
MatchMetric	Specifies if the system matches the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored. The default is 0.
MatchMetricTypeelsis	Specifies the match metric type field in the incoming IS-IS routes in accept policy.
IpRoutePolicyMatchIpMcastGroup	Specifies the system matches the IP multicast destination group for the routed multicast policy.
IpRoutePolicyMatchIpMcastSource	Specifies the system matches the IP multicast source address for the routed multicast policy.
MatchAsPath	Configures if the system matches the BGP autonomous system path. This parameter applies to BGP and overrides the BGP neighbor filter list information.
MatchCommunity	Filters incoming and outgoing updates based on a Community List. This parameter applies to BGP. The default is disable.
MatchCommunityExact	Indicates if the match must be exact (that is, all of the communities specified in the path must match). This parameter applies to BGP. The default is disabled.
MatchTag	Specifies a list of tags used during the match criteria process. This parameter applies to BGP and OSPF.
MatchVrf	Identifies the source VRFs that leaks routes to the local VRF (applies only to RIP routes).
MatchLocalPref	Specifies the local preference value to be matched.

Name	Description
NssaPbit	Configures or resets the P bit in specified type 7 link state advertisements (LSA). By default, the Pbit is always configured in case the user configures the Pbit to a disable state for a particular route policy other than all type 7. LSAs associated with that route policy have the Pbit cleared. With this intact, not so stubby area (NSSA) area border router (ABR) does not perform translation of these LSAs to type 5. The default is enable.
SetRoutePreference	Configures a preference value. The default value is 0. If the default is configured, the global preference value is used. When configured to a value greater than zero, specifies the route preference value assigned to the routes that matches the policy. This feature applies to accept policies only.
SetMetricTypeInternal	Identifies the MED value for routes advertised to BGP numbers to the Interior Gateway Protocol (IGP) metric value. The value must be 0 or 1. The default is 0.
SetMetricTypeIspis	Sets the metric type IS-IS.
SetMetric	Configures the system to use the metric value for the route while announcing or redistributing. The default-import-metric is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or the default value is used (applies to IS-IS routes also). The default is 0.
SetMetricType	Configures the metric type for the routes to announce into the OSPF routing protocol that matches this policy. Applicable to OSPF protocol only. The default is type 2. This field is applicable only for OSPF announce policies. The default is type2.
SetNextHop	Configures the IP address of the next-hop router. This parameter applies to BGP. The default is 0.0.0.0.
SetInjectNetList	Configures the destination network of the route that matches this policy with the contents of the specified prefix list. Select the ellipsis and choose from the list in the Set Inject NetList dialog box.
SetMask	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.

Name	Description
SetAsPath	<p>Indicates the AS path value to use whether the SetAsPathMode field is Tag or Prepend. This parameter applies to BGP.</p> <p>Note: Prepend is not applicable to an internal BGP (iBGP) peer with outbound route policy.</p>
SetAsPathMode	<p>Configures if the system converts the tag of a route into an AS path. This parameter applies to BGP. The mode is either Tag or Prepend tag. The value is applicable only while redistributing routes to BGP. The default is prepend.</p> <p>Note: Prepend does not apply to an iBGP peer with outbound route policy.</p>
SetAutomaticTag	<p>Enables the automatic tag feature. This parameter applies to BGP. The default is disable.</p>
SetCommunityNumber	<p>Configures the community number for BGP advertisements. This value can be a number (1 to 42949672000) or no-export or no-advertise.</p>
SetCommunityMode	<p>Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged.</p> <ul style="list-style-type: none"> Unchanged—keeps the community attribute in the route path as it is. None—removes the community in the route path additive. Append—adds the community number specified in SetCommunityNumber to the community list attribute.
SetExtCommunity	<p>Configures a BGP community.</p>
SetExtCommunityMode	<p>Configures the extended-community mode. The value can be append, unchanged, or overwrite. The default value is unchanged.</p> <ul style="list-style-type: none"> append — creates another community string. unchanged — keeps the community attribute as it is. overwrite — changes the current value.

Name	Description
SetOrigin	Configures the origin for BGP to IGP, EGP, incomplete, or unchanged. If not configured, the system uses the route origin from the IP routing table (protocol). The default is unchanged.
SetLocalPref	Configures the local preference for BGP. The system uses this value during the route decision process for the BGP protocol. The default is 0.
SetOriginEgpAs	Indicates the remote autonomous system number for BGP. The default is 0.
SetWeight	Configures the weight value for the routing table for BGP. This field must be used with the match as-path condition. For BGP, this value overrides the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. The default is 0.
SetTag	Configures a tag. The default is 0. This parameter applies to BGP.
Ipv6SetNextHop	Specifies the address of the IPv6 next hop router.
SetDatalsid	Configures the data I-SID for SPB routed multicast Policy that associates with the IP multicast groups (sender and receiver). This field should match the route-map.
SetRxOnly	Configures SPB routed multicast policy with permit action that allows only local receivers, not local senders. The default is disabled.
SetTxOnly	Configures a route-map of SPB routed multicast policy with permit action that allows only local senders, not local receivers. The default is disabled.

Apply a Route Policy

Apply route policies to define route behavior.

About This Task



Important

Changing route policies or prefix lists that affect OSPF accept or redistribute is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. Therefore, if you want to change a prefix list or a routing protocol, you configure all route policies and prefix lists before enabling the protocols.

Procedure

1. In the navigation pane, expand: **Configuration > IP**.
2. Select **Policy**
3. Select the **Applying Policy** tab.
4. Select the type of policy to apply.
5. Select **Apply**.

Applying Policy Field Descriptions

Use the data in the following table to use the **Applying Policy** tab.

Name	Description
RoutePolicyApply	Specifies that configuration changes in the policy take effect in an OSPF route policy context. This prevents the system from attempting to apply the changes one by one after each configuration change. The default is enabled.
RedistributeApply	Specifies that configuration changes in the policy take effect for an OSPF Redistribute context. This prevents the system from attempting to apply the changes one-by-one after each configuration change. The default is enabled.

Configure an OSPF Accept Policy

Perform the following procedure to create or configure an OSPF accept policy.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand: **Configuration > IP**.
2. Select **Policy**.
3. Select the **OSPF Accept** tab.
4. Select **Insert**.
5. Configure the parameters as required.
6. Select **Insert**.

OSPF Accept *Field Descriptions*

Use the data in the following table to use the **OSPF Accept** tab.

Name	Description
AdvertisingRtr	Specifies the routing ID of the advertising router.
Enable	Enables or disables the advertising router. You can also enable or disable advertising in the OSPF Accept tab of the Policy dialog box by selecting in the field and selecting enable or disable from the menu. The default is disable.
MetricType	Specifies the OSPF external type. This parameter describes which types of OSPF ASE routes match this entry. <ul style="list-style-type: none">Any means match either ASE type 1 or 2Type1 means match any external type 1Type2 means match any external type 2 You can also select your entry in the OSPF Accept tab of the Policy dialog box by selecting in the field and selecting any, type1, or type2 from the menu. The default is any.
PolicyName	Specifies the name of the OSPF in filter policy. Select the ellipsis and choose from the list in the Policy Name dialog box. To clear an entry, use the ALT key.

Configure Inbound and Outbound Filtering Policies on a RIP Interface

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Configure inbound filtering on a RIP interface to determine whether to learn a route on a specified interface and to specify the parameters of the route when it is added to the routing table. Configure outbound filtering on a RIP interface to determine whether to advertise a route from the routing table on a specified interface and to specify the parameters of the advertisement.

The port on which the multimedia filter is enabled becomes a DIFFSERV access port.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Policy**.
3. Select the **RIP In/Out Policy** tab.
4. In the desired row, double-click the **InPolicy** or **OutPolicy** column.

5. Select a preconfigured In/Out policy, and then select **OK**.

RIP In/Out Policy Field Descriptions

Use the data in the following table to use the **RIP In/Out Policy** tab.

Name	Description
Address	Specifies the IP address of the RIP interface.
Interface	Specifies the internal index of the RIP interface.
InPolicy	Specifies the policy name used for inbound filtering on this RIP interface. This policy determines whether to learn a route on this interface and specifies the parameters of the route when it is added to the routing table.
OutPolicy	Specifies the policy name used for outbound filtering on this RIP interface. This policy determines whether to advertise a route from the routing table on this interface and specifies the parameters of the advertisement.

Delete Inbound or Outbound Filtering Policies on a RIP Interface

Delete a RIP In or Out policy when you no longer want the switch to learn a route on a specified interface or advertise a route from the routing table on a specified interface.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Policy**.
3. Select the **RIP In/Out Policy** tab.
4. In the desired row, double-click the **InPolicy** or **OutPolicy** column for the policy you want to delete.
5. In the **InPolicy** or **OutPolicy** dialog box, press **CTRL** and then, select the policy you want to delete.
6. Select **OK**.
7. Select **Apply**.



Routed Split MultiLink Trunking

[RSMLT](#) on page 3285

[IPv6 RSMLT](#) on page 3286

[SMLT and RSMLT Operation in Layer 3 Environments](#) on page 3287

[RSMLT Configuration using CLI](#) on page 3291

[RSMLT Configuration using EDM](#) on page 3295

[Troubleshooting RSMLT](#) on page 3302

Table 277: RSMLT for IPv4 product support

Feature	Product	Release introduced
IPv4 RSMLT	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 278: RSMLT for IPv6 product support

Feature	Product	Release introduced
IPv6 RSMLT	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

RSMLT

In many cases, core network convergence time depends on the length of time a routing protocol requires to successfully converge. Depending on the specific routing protocol, this convergence time can cause network interruptions that range from seconds to minutes.

Routed Split MultiLink Trunking (RSMLT) permits rapid failover for core topologies by providing an active-active router concept to core SMLT networks.

RSMLT scenarios include SMLT triangles, squares, and SMLT full-mesh topologies, with routing enabled on the core VLANs.

Routing protocols include the following:

- IP Unicast Static Routes
- RIP1
- RIP2
- OSPF
- BGP

In the event of core router failures, RSMLT manages packet forwarding, thus eliminating dropped packets during the routing protocol convergence.

IPv6 RSMLT

Routed Split Multi-Link Trunking (RSMLT) is an enhancement to SMLT that enables the exchange of Layer 3 information between peer nodes in a switch cluster. RSMLT provides two main advantages over SMLT:

- provides backup for the peer after the peer goes down
- routes traffic on behalf of the peer to prevent Virtual Inter-Switch Trunk (vIST) overload

IPv6 RSMLT enables the subsecond failover for IPv6 forwarding.

The overall model for IPv6 RSMLT is essentially identical to that of IPv4 RSMLT. In short, RSMLT peers exchange their IPv6 configuration and track their states by using vIST messages. An RSMLT node always performs IPv6 forwarding on the IPv6 packets destined to the MAC addresses of the peer. If an RSMLT node detects that the RSMLT peer is down, the node forwards IPv6 traffic destined to the IPv6 addresses of the peer.

With RSMLT enabled, an SMLT switch performs IP forwarding on behalf of the SMLT peer, which prevents IP traffic from being sent over the vIST.

IPv6 RSMLT supports the full set of topologies and features supported by IPv4 RSMLT, including SMLT triangles, squares, and SMLT full-mesh topologies, with routing enabled on the core VLANs.

Because you configure RSMLT on a VLAN, not at the IP layer, the configuration applies to both IPv4 and IPv6. You cannot enable or disable RSMLT on a VLAN for IPv6 but not IPv4; or for IPv4 but not IPv6.

With IPv6, you must configure the RSMLT peers to use the same set of IPv6 prefixes.

Supported routing protocols include the following:

- IPv6 static routes

- OSPFv3

**Note**

IPv6 RSMLT is not virtualized, therefore it is not possible to enable IPv6 and RSMLT together on a VLAN which is associated with a VRF.

The system will display the configuration errors if you attempt to perform the following:

- Create an IPv6 interface on a VLAN which is associated with a VRF and RSMLT is enabled on the VLAN.
- Enable RSMLT on an IPv6 enabled VLAN which is associated with a VRF.

IPv6 differences

The following list identifies ways in which the IPv6 implementation of RSMLT differs from the IPv4 implementation of RSMLT.

- After the switch begins to forward traffic on behalf of the peer, duplicate address detection (DAD) is not executed for the IPv6 address of the peer. The implementation assumes that the peer IPv6 address is already known to be unique.
- An RSMLT switch installs a neighbor entry for the peer IPv6 address immediately after the peer disappearance is detected, possibly while a route for the peer still exists. This action can result in packets destined to the peer IPv6 address being delivered to the CP for a short period of time.
- You cannot configure a vIST with IPv6 peer address
- In a dual-stack VLAN, adding or deleting IPv4 or IPv6 does not affect the RSMLT functionality of one another. If you add IPv4 or IPv6 to an existing IPv6 or IPv4 RSMLT VLAN, the RSMLT state for the protocol you add second will be the same as the previous RSMLT state.

SMLT and RSMLT Operation in Layer 3 Environments

[Router R1 Failure](#) on page 3288 shows a typical redundant network example with user aggregation, core, and server access layers. To minimize the creation of many IP subnets, one VLAN (VLAN 1, IP subnet A) spans all wiring closets.

SMLT provides the loop-free topology and forwards all links for VLAN 1, IP subnet A.

The aggregation layer switches are configured with routing enabled and provide active-active default gateway functionality through RSMLT.

After you enable RSMLT on a VLAN (on both aggregation devices), the cluster devices simply inform each other (over vIST messaging) of their physical IP and MAC on that VLAN. Thereafter, the two cluster devices take mutual ownership of their IP addresses on that VLAN. This action means each cluster device routes IP traffic that is directed to the physical MAC of the IP or the physical MAC of the peer IP on that VLAN, and when one of them is down the other cluster device:

- Replies to ARP requests for both the IP and the peer IP on that VLAN
- Replies to pings to the IP and the peer IP on that VLAN

In this case, routers R1 and R2 forward traffic for IP subnet A. RSMLT provides both router failover and link failover. For example, if the Split MultiLink Trunk link between R2 and R4 is broken, the traffic fails over to R1 as well.

For IP subnet A, VRRP with a backup master can provide the same functionality as RSMLT, as long as no additional router is connected to IP subnet A.

RSMLT provides superior router redundancy in core networks (IP subnet B), where OSPF is used for the routing protocol. Routers R1 and R2 provide router backup for each other, not only for the edge IP subnet A, but also for the core IP subnet B. Similarly routers R3 and R4 provide router redundancy for IP subnet C and also for core IP subnet B.

Router R1 Failure

The following figure shows SMLT and RSMLT in Layer 3 environments.




Figure 217: SMLT and RSMIT in Layer 3 environments

R3 and R4 both use R1 as their next hop to reach IP subnet A. Even though R4 sends the packets to R2, they are routed directly at R2 into subnet A. R3 sends its packets to R1 and they are also sent directly into subnet A. After R1 fails, all packets are directed to R2, with SMLT. R2 still routes for R2 and R1. After OSPF convergence, the routing tables in R3 and R4 change their next hop to R2 to reach IP subnet A. You can configure the hold-up timer (that is, for the amount of time R2 routes for R1 in a failure) for a time period greater than the routing protocol convergence, you can configure it as indefinite (that is, the members of the pair always route for each other).

Use an indefinite hold-up timer value for applications that use RSMLT at the edge instead of VRRP.

Router R1 Recovery

When R1 restarts after a failure, it does not route traffic for R2, nor does it provide backup for R2, until the hold-down timer expires. Similar to VRRP, the hold-down timer value must be greater than the time the routing protocol requires to converge its tables.

During the hold-down interval, R1 routes traffic if the destination MAC of the packet is its own routable VLAN MAC. R1 bridges incoming traffic to R2 if the destination MAC of the packet is the routable VLAN MAC of R2. A temporary default route (one having a route preference equal to 4) that points to R2 is installed on R1. R1 uses this temporary route to forward traffic to R2 that it cannot route itself because of the incomplete routing table; the default route is not saved in the configuration file.

After the hold-down timer expires, the temporary default route that points to R2 is deleted; from this moment on, in addition to routing packets destined to itself, R1 starts routing packets for R2, as well.

RSMLT Network Design and Configuration

Because RSMLT is based on SMLT, all SMLT configuration rules apply. In addition, RSMLT is enabled on the SMLT aggregation switches for each VLAN. The VLAN must be a member of SMLT links and vIST's L2VSN. For more information about how to configure SMLT in a Layer 2 environment, see [MultiLink Trunking and Split MultiLink Trunking](#) on page 2657.

The VLAN also must be routable (IP address configured) and you must configure an Interior Gateway Protocol (IGP) such as OSPF on all four routers, although it is independent of RSMLT. All routing protocols, even static routes, work with RSMLT.

The RSMLT pair switches provide backup for each other. As long as one of the two routers of a vIST pair is active, traffic forwarding is available for both next hops R1/R2 and R3/R4.

RSMLT Edge Support

The switch stores the peer MAC and IP address pair in its local configuration file and restores the configuration if the peer does not restore after a simultaneous restart of both RSMLT-peer switches.

The RSMLT edge support feature adds an enhancement whereby the peer MAC (for the IP on the VLAN) is committed to the config.cfg file after you use the **save config** command. If you power off both devices, and then power up only one of them, that single device can still take ownership of its peer IP on that VLAN even if it has not seen that peer switch since it started. This enhancement is necessary if you configure the peer (the device which is still down) IP as the default gateway in end stations.

If you enable RSMLT edge support, you must also ensure that the hold-up timer for RSMLT on those edge VLANs equals infinity (9999). This timer value ensures that if one cluster device fails, the remaining cluster device maintains ownership of the failed peer IP indefinitely.

The edge VLAN can be tagged over SMLT links, single attached links, or more SMLT links.

**Important**

If you clear the peer information the device can stop forwarding for the peer.

RSMLT implementation does not use a virtual IP address but instead uses physical IP addresses for redundancy. At the same time, you can deploy RSMLT in either routed configurations, or edge configurations, where you previously used VRRP (and back-up master). Previously, if a power outage occurred or a shutdown of both switches within a dual core vIST pair, only one device came back up. Clients using the powered-off device IP/MAC as the default gateway lost connectivity to the network. In such a scenario, even with RSMLT enabled on the device, it cannot act as a backup for the peer as it was unaware of the peer IP or MAC address.

After both the dual core vIST switches come back, the vIST is operational. If an RSMLT peer-enabled message is received from the peer, normal RSMLT operation occurs.

If the peer has either an IP or MAC change, you must save the configuration for the RSMLT edge support to operate correctly. However, if the vIST peer up message is not received (for example, if you do not enable RSMLT properly), and you enable the RSMLT edge support flag, the RSMLT hold-down timer starts and permits routing protocols to converge; during this time user operation can be affected. After the hold-down timer expires, saved peer information is picked up and the device starts to act as backup for the peer by adding the previously saved MAC and ARP records.

The hold-up timer starts and after this timer expires the previously added MAC and ARP records are deleted and the device stops acting as backup for the peer, as the peer is not running proper RSMLT for the VLAN. The RSMLT is a parameter for each VLAN, and therefore all affects are on an individual VLAN basis, not necessarily a global device. Edge support mode uses the local values of the hold-down timer (default value of 60 seconds) and hold-up timer (default value of 180 seconds).

RSMLT Configuration using CLI

Use the procedures in this section to configure RSMLT using CLI.

Configure RSMLT on a VLAN

Configure RSMLT on a VLAN to exchange Layer 3 information between peer nodes in a switch cluster. This configuration applies to both IPv4 and IPv6.

Before You Begin

- An IP routing protocol is enabled on VLAN Layer 3 interfaces.
- VLANs with Layer 3 interfaces participate in Split MultiLink Trunking (SMLT).

About This Task

The VLAN can be either IPv4 or IPv6, or both. RSMLT configuration on a VLAN simultaneously affects both IPv4 and IPv6. By default, RSMLT is disabled on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Configure the holddown timer:

```
ip rsmlt holddown-timer <0-3600>
```

3. Configure the holdup timer:

```
ip rsmlt holdup-timer <0-9999>
```

4. Enable RSMLT on the VLAN:

```
ip rsmlt
```

Example

Configure and enable RSMLT on a VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 100
Switch:1(config-if)#ip rsmlt holddown-timer 100
Switch:1(config-if)#ip rsmlt holdup-timer 200
Switch:1(config-if)#ip rsmlt
```

Variable Definitions

The following table defines parameters for the **ip rsmlt** command.

Variable	Value
<i>holddown-timer</i> <0-3600>	<p>Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The default is 60.</p> <p>Configure this value to be longer than the anticipated routing protocol convergence.</p> <p>If you disable RSMLT on a VLAN, non default values for this variable do not save across restarts.</p>
<i>holdup-timer</i> <0-3600 9999>	<p>Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.</p> <p>If you disable RSMLT on a VLAN, non default values for this variable do not save across restarts.</p>

The following table defines parameters for the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Enable RSMLT Edge Support

Enable RSMLT Edge support to store the RSMLT peer MAC and IP address-pair in the local configuration file and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT peer systems.

If enabled, peer MAC and IP information for all RSMLT-enabled VLANs is saved after you next use the **save config** command.

About This Task

RSMLT Edge support configuration applies to both IPv4 and IPv6. You do not configure IPv4 and IPv6 separately.

The RSMLT Edge support default is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable RSMLT Edge support:

```
ip rsmlt edge-support
```

Example

If you have enabled RSMLT Edge Support, disable the feature as follows:

```
Switch:1(config)#no ip rsmlt edge-support
```

View RSMLT Information

Show RSMLT information to view data about all RSMLT interfaces. The output of the command includes IPv6 information for the local and peer nodes.

About This Task

If you use the **show ip rsmlt** command after you delete an RSMLT, the RSMLT still shows in the command output until you restart the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Show RSMLT information about the interface:
show ip rsmlt [local|peer] [vrf WORD<1-16>] [vrfrids WORD<0-512>]
3. View the status of the switch to act as a peer forwarder:
show ip rsmlt edge-support

Examples

```
Switch:1>show ip rsmlt
```

```
=====
                        Ip Rsmlt Local Info - GlobalRouter
=====
```

VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
101	101.1.1.32	00:24:7f:9e:da:01	Enable	Up	100	200
102	102.1.1.32	00:24:7f:9e:da:02	Enable	Up	60	180

VID	SMLT ID
101	101
102	102

VID	IPv6	MAC	ADMIN	OPER	HDTMR	HUTMR
101	1010:0:0:0:0:0:0:0/64 1010:0:0:0:0:0:0:32/64 fe80:0:0:0:224:7fff:fe9e:da01/128	00:24:7f:9e:da:01	Enable	Up	100	200
102	1020:0:0:0:0:0:0:0/64 1020:0:0:0:0:0:0:32/64 fe80:0:0:0:224:7fff:fe9e:da02/128	00:24:7f:9e:da:02	Enable	Up	60	180

VID	SMLT ID
101	101
102	102

```
=====
                        Ip Rsmlt Peer Info - GlobalRouter
=====
```

VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
101	101.1.1.33	00:24:7f:9e:ea:01	Enable	Up	100	200
102	102.1.1.33	00:24:7f:9e:ea:00	Enable	Up	60	180

VID	HDT REMAIN	HUT REMAIN	SMLT ID
101	60	180	101
102	60	180	102

VID	IPv6	MAC	ADMIN	OPER	HDTMR	HUTMR
-----	------	-----	-------	------	-------	-------

```

101          00:24:7f:9e:ea:01  Enable  Up    100    200
    1010:0:0:0:0:0:0:0/64
    1010:0:0:0:0:0:0:33/64
    fe80:0:0:0:224:7fff:fe9e:ea01/128
102          00:24:7f:9e:ea:00  Enable  Up    60    180
    1020:0:0:0:0:0:0:0/64
    1020:0:0:0:0:0:0:33/64
    fe80:0:0:0:224:7fff:fe9e:ea00/128

```

```

VID  HDT  REMAIN  HUT  REMAIN  SMLT ID
-----
101   60         180
101
102   60         180
102

```

```
Switch:1>show ip rsmlt edge-support
```

```

RSMLT Peer Info:
    rsmlt-peer-forwarding : disable

```

Variable definitions

Use the data in the following table to use the **show ip rsmlt** command.

Variable	Value
<i>local</i>	Shows local RSMLT information.
<i>peer</i>	Shows RSMLT information for the peer.
<i>vrf WORD<1-16></i>	Shows information for a specific VRF name.
<i>vrfids WORD<0-512></i>	Shows information for a specific VRF ID.

RSMLT Configuration using EDM

Use the procedures in this section to configure RSMLT using EDM.

Configure RSMLT on a VLAN

Before You Begin

- Enable an IP routing protocol on VLAN Layer 3 interfaces.
- Ensure VLANs with Layer 3 interfaces participate in Split MultiLink Trunking (SMLT).

About This Task

Configure RSMLT on a VLAN to exchange Layer 3 information between peer nodes in a switch cluster. This configuration applies to both IPv4 and IPv6.

The VLAN can be either IPv4 or IPv6, or both. RSMLT configuration on a VLAN simultaneously affects both IPv4 and IPv6.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.

2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **RSMLT** tab.
7. Select **Enable**.
8. In **HoldDownTimer**, type a hold-down timer value.
9. In **HoldUpTimer**, type a holdup timer value.
10. Select **Apply**.

RSMLT Field Descriptions

Use the data in the following table to use the **RSMLT** tab.

Name	Description
Enable	Enables RSMLT. The default is disabled.
HoldDownTimer	<p>Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address.</p> <p>The range of this value is from 0 to 3600 seconds. The default is 60.</p> <p>If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.</p>
HoldUpTimer	<p>Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.</p> <p>If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.</p>

Enable RSMLT Edge support

About This Task

Enable RSMLT Edge support to store the RSMLT peer MAC and IP address-pair in the local configuration file and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT peer systems.

The default is disabled.

RSMLT Edge support configuration applies to both IPv4 and IPv6.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RSMLT**.
3. Select the **Globals** tab.
4. Select **EdgeSupportEnable**.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
EdgeSupportEnable	Enables RSMLT Edge, which stores the peer MAC and IP address pair in the local configuration file, and restores the configuration if the peer does not restore after a simultaneous restart of both RSMLT-peer switches. The default is disabled.

View and Edit IPv4 RSMLT Local Information

About This Task

Perform the following procedure to view and edit RSMLT local VLAN information.

Procedure

1. In the navigation pane, expand **ConfigurationIP**.
2. Select **RSMLT**.
3. Select the **Local** tab.
4. Configure the parameters as required.
5. Select **Apply**.

Local Field Descriptions

Use the data in the following table to use the **Local** tab.

Name	Description
IfIndex	IP interface identification.
VlanId	Specifies the VLAN ID of the chosen VLAN.
IpAddr	Specifies the IP address on the RSMLT VLAN.
MacAddr	Specifies the MAC address of the selected VLAN.
Enable	Displays the RSMLT operating status as enabled or disabled.
OperStatus	Displays the RSMLT operating status as either up or down. The default is down.
HoldDownTimer	Defines how long the recovering/restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The range of this value is from 0 to 3600 seconds. The default is 60.
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.

Name	Description
SmltId	Specifies the ID range for the SMLT. A valid range is 1 to 512.
VrfId	Identifies the VRF.
VrfName	Indicates the VRF name.

Modify the IPv6 RSMLT Local Information

Edit the existing RSMLT configuration for the local node in the cluster.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **RSMLT**.
3. Select the **Local** tab.
4. Select a cell to change the value.
5. Select **Apply**.

Local Field Descriptions

Use the data in the following table to use the **Local** tab.

Name	Description
IfIndex	Shows the route SMLT operation index.
Ipv6Addr	Configures the IPv6 address of the RSMLT interface.
Ipv6PrefixLength	Configures the IPv6 prefix length.
Enable	Enables or disables RSMLT. The default is disabled.
HoldDownTimer	Defines how long the recovering/rebooting switch remains in a non-Layer 3 forwarding mode for the peer router MAC address. The default is 60.
HoldUpTimer	Defines how long the RSMLT switch maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.
OperStatus	Displays the RSMLT operating status as either up or down.
SmltId	Specifies the ID range for the SMLT.
VlanId	Configures the VLAN ID.
MacAddr	Configures the MAC address of the VLAN.

Name	Description
VrfId	Indicates the virtual router ID to which the local RSMLT instance belongs.
VrfName	Indicates the virtual router name to which the local RSMLT instance belongs.

Modify IPv6 RSMLT Peer Information

Edit the existing configuration for the RSMLT peer node in the cluster.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **RSMLT**.
3. Select the **Peer** tab.
4. Select a cell to change the value.
5. Select **Apply**.

Peer Field Descriptions

Use the data in the following table to use the **Peer** tab.

Name	Description
IfIndex	Shows the route SMLT operation index.
Ipv6Addr	Configures the IPv6 address of the RSMLT interface.
Ipv6PrefixLength	Configures the IPv6 prefix length.
AdminStatus	Shows the administrative status of RSMLT on the peer.
HoldDownTimer	Defines how long the recovering/rebooting switch remains in a non-Layer 3 forwarding mode for the peer router MAC address. The default is 0.
HoldDownTimeRemaining	Indicates the time remaining in the HoldDownTimer.
HoldUpTimer	Defines how long the RSMLT switch maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 0.
HoldUpTimeRemaining	Indicates the time remaining in the HoldUpTimer.
OperStatus	Displays the RSMLT operating status as either up or down.
SmltId	Specifies the ID range for the SMLT.

Name	Description
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
MacAddr	Configures the MAC address of the VLAN.
VrfId	Indicates the virtual router ID to which the peer belongs.
VrfName	Indicates the virtual router name to which the peer belongs.

View IPv4 RSMLT Peer Information

About This Task

Perform this procedure to view and edit RSMLT peer information.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **RSMLT**.
3. Select the **Peer** tab.

Peer Field Descriptions

Use the data in the following table to use the **Peer** tab.

Name	Description
IfIndex	IP interface identification.
VlanId	Specifies the VLAN ID of the chosen VLAN.
IpAddr	Specifies the IP address on the RSMLT VLAN.
MacAddr	Specifies the MAC address of the selected VLAN.
Enable	Displays the RSMLT operating status as enabled or disabled.
OperStatus	Displays the RSMLT operating status as either up or down. The default is down.
HoldDownTimer	Defines how long the recovering/restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address. The range of this value is from 0 to 3600 seconds. The default is 0.

Name	Description
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 0.
HoldDownTimeRemaining	Displays the time remaining of the HoldDownTimer. The default is 0.
HoldUpTimeRemaining	Displays the time remaining of the HoldUpTimer. The default is 0.
SmltId	Specifies the ID range for the Split MultiLink Trunk. A valid range is 1 to 32.
VrfId	Identifies the VRF.
VrfName	Indicates the VRF name.

View IPv6 RSMLT Edge Peers

View the RSMLT peers for which the switch acts as a peer forwarder.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **RSMLT**.
3. Select the **Edge Peers** tab.

Edge Peers Field Descriptions

Use the data in the following table to use the **Edge Peers** tab.

Name	Description
PeerVlanId	Specifies the ID of the VLAN associated with this entry.
PeerIpv6Address	Specifies the IPv6 address of the peer RSMLT interface.
PeerIpv6PrefixLength	Specifies the peer IPv6 address prefix.
PeerMacAddress	Specifies the peer MAC address.

View IPv4 RSMLT Edge Support Information

About This Task

View RSMLT edge support information to verify the RSMLT peer MAC/IP address-pair in its local configuration file and restore the configuration if the peer does not restore it after a simultaneous restart of both RSMLT-peer systems.

Procedure

1. In the navigation pane, expand **Configuration > IP**.

2. Select **RSMLT**.
3. Select the **Edge Peers** tab.

Edge Peers Field Descriptions

Use the data in the following table to use the **Edge Peers** tab fields.

Name	Description
VlanId	Specifies the VLAN ID of the chosen VLAN.
PeerIpAddress	Specifies the peer IP address.
PeerMacAddress	Specifies the peer MAC address.
PeerVrfId	Identifies the Peer VRF.
PeerVrfName	Specifies the Peer VRF name.

Troubleshooting RSMLT

The following sections provide information for troubleshooting IPv4 Split Multi-Link Trunking (RSMLT).

RSMLT Peers Not Up

If, after a series of reconfiguration, RSMLT peers do not transition to the up state, use the following procedure to troubleshoot the issue. You can observe this issue on dual-stack VLANs after multiple delete and re-adds of IPv4 interfaces.

Procedure

1. Display the RSMLT configuration. This command shows whether the peers are up:
`show ip rsmlt peer`
2. Enter VLAN Interface Configuration mode:
`enable`
`configure terminal`
`interface vlan <1-4059>`
3. To recover the peers if they are down, disable and reenabling RSMLT on both IST peers:
`no ip rsmlt`
`ip rsmlt`
4. If the problem persists, boot from a saved configuration.

Example

Display the RSMLT configuration:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 1
Switch:1(config-if)#show ip rsmlt peer
```

```

=====
                        Ip Rsmlt Peer Info - GlobalRouter
=====

VID      IP              MAC              ADMIN    OPER    HDTMR    HUTMR
-----
1        192.0.2.1        00:1f:ca:1e:d3:1e      Enable    Up      60      180
2        198.51.100.1    00:1b:ca:1d:e3:1d      Enable    Up      60      180

VID      HDT REMAIN    HUT REMAIN    SMLT ID
-----
1        60      180      10
2        60      180      10, 16

VID      IPv6              MAC              ADMIN    OPER    HDTMR    HUTMR
-----

VID      HDT REMAIN    HUT REMAIN    SMLT ID
-----
Switch:1(config-if)#no ip rsmlt
Switch:1(config-if)#ip rsmlt

```

Enable Trace Messages for RSMLT Troubleshooting

Use the following procedure to obtain additional RSMLT-related information.

Procedure

If the preceding information does not resolve the issue, you can use the following command to obtain additional RSMLT-related information:

```
trace level 173 4
```



Important

Enabling this trace on a loaded system can slow down the CPU, especially if executed through the console. Use Telnet if possible.



SD-WAN

[SD-WAN Local Breakout](#) on page 3306

[ExtremeCloud SD-WAN Configuration Using CLI](#) on page 3312

[ExtremeCloud SD-WAN Configuration Using EDM](#) on page 3315

Table 279: SD-WAN product support

Feature	Product	Release introduced
Fabric Extend Integration with ExtremeCloud SD-WAN	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.10.1 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	Fabric Engine 8.10.1
	5520 Series	Fabric Engine 8.10.1
	5720 Series	Fabric Engine 8.10.1
	7520 Series	Fabric Engine 8.10.1
	7720 Series	Fabric Engine 8.10.1
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.10.1
	VSP 7400 Series	VOSS 8.10.1

Table 279: SD-WAN product support (continued)

Feature	Product	Release introduced
Ability to specify the VRF that the switch uses for SD-WAN	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0.3
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1
Auto-sense port Multi-area SPB support	4220 Series	Not Applicable
	5320 Series	Not Applicable
	5420 Series	Not Applicable
	5520 Series	Fabric Engine 9.0.3
	5720 Series	Fabric Engine 9.0.3
	7520 Series	Fabric Engine 9.0.3
	7720 Series	Fabric Engine 9.0.3
	7830 Series	Not Supported
	VSP 4900 Series	Not Applicable
	VSP 7400 Series	VOSS 9.0.3
<code>mgmt sdwan interface</code> Note: Requires a minimum of ExtremeCloud SD-WAN 24.4	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.1
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1

If you have remote branch sites with only Internet connectivity and you want to extend the capabilities of fabric to that site, or if you have a branch office with a private WAN connection and you want to have a backup link using the public Internet, you can extend the Fabric network over the Internet towards a remote branch by connecting Fabric switches through ExtremeCloud SD-WAN.

Fabric Extend Integration with ExtremeCloud SD-WAN integrates the following features for SD-WAN deployments:

- Automation of Fabric Extend Tunnels over IP WAN infrastructure
- MTU configuration per IS-IS logical interface for Layer 3 interfaces
- SD-WAN infrastructure adjusts packet size to accommodate restricted MTU paths, such as Internet, by fragmenting the end user IP packets.

This deployment provides the following benefits:

- Encrypted transport
- Automated provisioning of Fabric switches - Fabric Extend tunnels, MTU, and tunnel state exchange between the SD-WAN Appliance and Fabric switches. Only one port on the switch supports this capability. For more information, see [Auto-sense Port States](#) on page 28.

**Note**

On 4220 Series and 5320 Series models that support a single VRF, ensure you specify the VRF that Auto-sense uses for SD-WAN. For more information, see [Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration](#) on page 3312.

- Enhanced branch resiliency - one or more Fabric switches must be connected to the LAN port or ports of the SD-WAN Appliance to achieve automated redundancy.
- Visibility of branch application traffic and optimized performance

SD-WAN Local Breakout

Auto-sense SD-WAN Local Breakout automatically provides connectivity to devices attached to a data VRF, through the SD-WAN VRF, towards the Internet. Use a Local Breakout (LBO) configuration for the following use cases:

- Send Internet-destined traffic directly to the Internet at the branch, through the SD-WAN Appliance and not through the Fabric.

- Provide the SD-WAN Appliance management access to a Fabric service for management and for NTP server reachability.

Table 280: SD-WAN Local Breakout product support

Feature	Product	Release introduced
ExtremeCloud SD-WAN Local Breakout configuration applied through Auto-sense	4220 Series	Not Applicable
	5320 Series	Fabric Engine 9.1 Not supported on models with a single active VRF.
	5420 Series	Fabric Engine 9.1
	5520 Series	Fabric Engine 9.1
	5720 Series	Fabric Engine 9.1
	7520 Series	Fabric Engine 9.1
	7720 Series	Fabric Engine 9.1
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.1
	VSP 7400 Series	VOSS 9.1

If you use SD-WAN Orchestrator to enable an LBO branch, Auto-sense on the switch can apply the LBO configuration to select VRFs. The switch configuration creates a BGP peering relationship with the SD-WAN Appliance in the dynamic VRF (SD-WAN), and then redistributes BGP routes it receives into the selected user-created VRFs.

This feature applies only to VRFs that exist on the same Fabric Extend switch that connects to the SD-WAN Appliance, and only for local or direct subnets on the same VRF. IS-IS does not advertise the LBO-injected routes further than the local switch. If you require LBO in the branch, make the directly-attached switch the default gateway with LBO enabled, and operate the other nodes in the branch as Layer 2 BEBs.

Manual Switch Configuration

You must identify which user-created VRFs accept the Local Breakout (LBO) routes. Auto-sense automatically configures BGP and RIP redistribution policies between the selected VRFs and the dynamic SD-WAN VRF. For more information about how to identify a VRF, see [Select a VRF for Local Breakout Configuration](#) on page 3314.

NEW! *Converting Auto-sense SD-WAN Configuration to a Manual Configuration*

Auto-sense automatically creates a default SD-WAN VRF and applies BGP configuration based on LLDP information from the SD-WAN appliance. If you require persistent or customized routing behavior, you can convert this dynamic configuration to a manual configuration.

When you issue the **auto-sense sd-wan vrf WORD<1-16>** command, the switch renames the default Auto-sense SD-WAN VRF. Use the **ip vrf sd-wan** command

to convert the SD-WAN VRF origin from DYNAMIC to CONFIG. The associated BGP instance remains dynamic through Auto-sense until you issue the **ip bgp** command in that VRF. This command converts the BGP instance to static (CONFIG), removes the Auto-sense BGP configuration, and saves the configuration.

Behavior Considerations

Consider the following key behavior rules:

- BGP origin states. A BGP instance can have one of the following origins:
 - AUTO-SENSE— Auto-sense manages the BGP instance and rejects any manual BGP configuration.
 - CONFIG—The BGP instance is user configured (static) and persistent.
 - DEFAULT (GRT only)— The Global Routing Table (GRT) always contains a BGP instance. You cannot delete it, but you can modify or disable it.
- When the SD-WAN port is down, Auto-sense creates the VRF when the **auto-sense sd-wan vrf <name>** command is issued; however, it does not create or activate a dynamic BGP instance until the port transitions to UP.
- When you remove Auto-sense from an SD-WAN VRF, the VRF and BGP origins determine how the switch handles Local Breakout (LBO). Static VRFs preserve user-configured BGP and LBO, while the switch removes dynamic objects that Auto-sense created whenever you relinquish Auto-sense control or the link state dictates removal.
- The Global Routing Table (GRT) always includes a BGP instance and starts with the DEFAULT origin at boot. When you change any SD-WAN applied BGP configuration, you change the GRT origin to CONFIG and Auto-sense stops applying LLDP-derived configuration. When Auto-sense manages the GRT BGP instance, the switch rejects global BGP settings and reports an Auto-sense error. The only exception is the AS-number command, which changes the origin to CONFIG and removes all Auto-sense applied configuration.
- Running configuration views show only BGP instances with origin CONFIG; the BGP configuration view displays AUTO-SENSE, CONFIG, and DEFAULT. Use these views to verify the origin state.

The following tables explain BGP origin behavior for SD-WAN VRFs and the Global Routing Table (GRT), and how the **auto-sense sd-wan vrf WORD<1-16>** command affects VRF creation. Use this reference to understand when Auto-sense applies

dynamic configuration, when manual configuration takes precedence, and how to achieve configuration persistence.

Table 281: Configuration options—SD-WAN VRFs (excludes the GRT)

Use Case	Command issued	SD-WAN port state	Current VRF origin	Current BGP origin	Result
Convert the Auto-sense default SD-WAN VRF from dynamic to static	ip vrf sd-wan	Any	AUTO-SENSE	AUTO-SENSE	Converts the SD-WAN VRF to CONFIG origin (static). Preserves BGP configuration.
Configure BGP on a VRF when BGP is statically configured	ip bgp	Any	CONFIG	CONFIG	Accepts the configuration and saves it.
Convert an Auto-sense configured BGP instance to static	ip bgp	Any	CONFIG	AUTO-SENSE	Removes all Auto-sense applied BGP configuration and changes BGP origin to CONFIG.
Convert a static BGP instance to Auto-sense (SD-WAN port is UP)	no ip bgp	UP	CONFIG	CONFIG	Removes the static BGP instance. Auto-sense applies a dynamic BGP instance and changes the BGP origin to AUTO-SENSE.
Apply Auto-sense to an existing VRF without BGP	auto-sense sd-wan vrf WORD<1-16>	UP	CONFIG		Creates a dynamic BGP instance and changes the BGP origin to AUTO-SENSE.
Apply Auto-sense to an existing VRF with static BGP	auto-sense sd-wan vrf WORD<1-16>	UP	CONFIG	CONFIG	Preserves the static BGP configuration; saves configuration.
Rename the default Auto-sense SD-WAN VRF to a new VRF	auto-sense sd-wan vrf WORD<1-16>	UP	AUTO-SENSE	AUTO-SENSE	Renames the dynamic sd-wan VRF and converts it to origin CONFIG. Preserves the Auto-sense BGP configuration unless a static configuration replaces it later.
Enable Local Breakout on a VRF	sd-wan local-breakout	Any	CONFIG or AUTO-SENSE	CONFIG or AUTO-SENSE	Enables LBO configuration on the specified VRF and preserves current origins.
Remove Auto-sense from the default SD-WAN VRF	no auto-sense sd-wan vrf	Any	AUTO-SENSE	AUTO-SENSE	Removes Auto-sense applied configuration, including BGP and LBO.

Table 281: Configuration options—SD-WAN VRFs (excludes the GRT) (continued)

Use Case	Command issued	SD-WAN port state	Current VRF origin	Current BGP origin	Result
Remove Auto-sense from a non-default VRF	no auto-sense sd-wan vrf WORD<1-16>	Any	CONFIG	CONFIG	Preserves static VRF and BGP configuration. Removes Auto-sense configuration.
SD-WAN port DOWN with non-default VRF		DOWN	AUTO-SENSE	AUTO-SENSE	Removes the VRF and BGP instance (Auto-sense objects)
SD-WAN port DOWN with non-default VRF		DOWN	CONFIG	AUTO-SENSE	Preserves the VRF; removes the Auto-sense BGP instance.
SD-WAN port DOWN with non-default VRF		DOWN	CONFIG	CONFIG	Preserves the VRF; removes LBO configuration.

Table 282: Configuration options – Global Routing Table (GRT)

Use Case	Command	Current GRT BGP origin	Result
GRT BGP origin at boot		DEFAULT	GRT BGP instance is present; origin is DEFAULT at boot.
Change any SD-WAN applied BGP configuration in GRT	ip bgp	DEFAULT	Changes GRT BGP origin to CONFIG. SD-WAN LLDP derived configuration no longer applies.
Attempt to configure GRT BGP when Auto-sense manages it	ip bgp	AUTO-SENSE	Rejects the configuration and indicates that Auto-sense controls the BGP instance.
Convert Auto-sense GRT BGP to static with 4-byte AS assignment	router bgp <as_number> enable	AUTO-SENSE	Removes all Auto-sense applied GRT BGP configuration and changes the origin to CONFIG.
Display behavior	show running-config show ip bgp conf	CONFIG	show running-config displays BGP instances origin CONFIG. show ip bgp conf displays the BGP origin as CONFIG, AUTO-SENSE, or DEFAULT.

Automatic Configuration using Auto-sense

The switch uses Auto-sense to read BGP Autonomous System Number (ASN) and neighbor information from the LLDP FE TLV, and apply a dynamic configuration. Auto-sense configures BGP and RIP for route redistribution purposes for ports with Auto-sense enabled and an ExtremeCloud SD-WAN appliance.

Optionally, you can apply a dynamic configuration in a VRF and a static BGP configuration in the global routing table (GRT) or another static VRF in parallel. To use this co-existence mode, you must globally enable BGP with 4-byte AS numbering format to apply the dynamic BGP configuration to the SD-WAN VRF. If you apply a static BGP configuration to the switch after the Auto-sense port transitions to the SDWAN state, Auto-sense enables BGP 4-byte AS number support, which allows you to apply the static BGP configuration.



Note

If you enable 4-byte AS numbering format, you can continue to use conventional 2-byte AS numbers.

If you apply a static configuration without a 4-byte AS number before you connect the SD-WAN appliance to an Auto-sense port, the Auto-sense port cannot complete the dynamic BGP configuration on the dynamic SD-WAN VRF and the Auto-sense port transitions to the SD-WAN-PENDING state. To resolve this situation, you must globally disable and then reenables BGP with a 4-byte AS number in the static configuration. The remaining static BGP configuration, including the AS number do not need to change.

Auto-sense performs the following dynamic configuration:

```
router bgp as-4-byte enable
router vrf sd-wan
  ip bgp
  ip bgp router-id <local_ip> #local_ip is received through LLDP
  ip bgp vrf-as <local_asn> #local_asn is received through LLDP
  no ip bgp synchronization
  no ip bgp auto-summary
  no ip bgp aggregation
  ip bgp enable
  ip bgp neighbor <vrouters_1> #vrouters_1,2,3 are IP addresses received through LLDP
  ip bgp neighbor <vrouters_1> remote-as <remote_asn> #remote_asn is received through LLDP
  ip bgp neighbor <vrouters_1> enable
exit

router vrf WORD<1-16>
  ip rip
  ip rip enable
  ip rip redistribute bgp vrf-src sd-wan
  ip rip redistribute bgp enable vrf-src sd-wan
  sd-wan-local-breakout
exit
ip rip apply redistribute bgp vrf WORD<1-16> vrf-src sd-wan

router vrf sd-wan
  ip bgp redistribute direct vrf-src WORD<1-16>
  ip bgp redistribute direct enable vrf-src WORD<1-16>
exit
```

```
ip bgp apply redistribute direct vrf sd-wan vrf-src WORD<1-16>  
ip more-specific-non-local-route
```

High Availability

For dual SD-WAN Appliance and switch pairings in High Availability deployments, you must manually configure BGP peering and a routed segment between the two switches. Auto-sense does not configure this deployment.

Restrictions

- You cannot configure LBO in the GRT.
- If you configure the SD-WAN VRF to use the GRT (**auto-sense sd-wan vrf GlobalRouter**), this configuration can interfere with BGP global configuration. For more information, see [Converting Auto-sense SD-WAN Configuration to a Manual Configuration](#) on page 3307.

ExtremeCloud SD-WAN Configuration Using CLI

The topics in this section provide the CLI configuration specific to interoperability with SD-WAN.

Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration

Before You Begin

On switches that support only a single active VRF, ensure a user-defined VRF with a different name does not exist.

About This Task

Perform this procedure to specify the VRF that Auto-sense uses for the SD-WAN configuration. On models that support a single active VRF, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.



Note

If the running switch configuration includes the Auto-sense-managed SD-WAN VRF and you complete this procedure, the switch renames the dynamic SD-WAN VRF to the specified name, and converts the VRF to a static (CONFIG) origin. The associated BGP instance remains Auto-sense-managed until you issue the **ip bgp** command in that VRF, which converts the BGP instance to a static (CONFIG) origin, removes Auto-sense-applied BGP configuration, and makes the configuration persistent.

Procedure

1. Enter Global Configuration mode:

```
enable  
  
configure terminal
```

- Specify the VRF name that Auto-sense uses for ExtremeCloud SD-WAN configuration:

```
auto-sense sd-wan vrf WORD<1-16>
```
- (Optional) To manually configure BGP in the static VRF, see [Configure BGP](#) on page 574.

Configure the IS-IS Area for a Specific Tunnel

About This Task



Note

You can only use this command on a switch with boundary-node abilities.

Perform this procedure to specify the IS-IS area where Auto-sense creates a specific ExtremeCloud SD-WAN-learned tunnel. This configuration overrides a global area configuration for all learned tunnels.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Configure the IS-IS area:

```
auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D} <home | remote>
```

Variable Definitions

The following table defines parameters for the **auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D} <home | remote>** command.

Variable	Value
{A.B.C.D}	Specifies the destination IP address of the SD-WAN logical interface.
<home remote>	Specifies the IS-IS area as home or remote. By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Configure Auto-sense to Create All Learned Tunnels in the Remote Area

About This Task



Note

You can only use this command on a switch with boundary-node abilities.

Perform this procedure to create all ExtremeCloud SD-WAN learned tunnels in the IS-IS remote area. You can also configure exceptions for specific tunnels. For more information, see [Configure the IS-IS Area for a Specific Tunnel](#) on page 3313.

By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Configure the IS-IS area:
`auto-sense sd-wan multi-area remote`

Select a VRF for Local Breakout Configuration

Enable Auto-sense to manage the VRF after it detects an ExtremeCloud SD-WAN peer.



Note

This procedure does not apply to switch models that support only a single active VRF with IP configuration.

Before You Begin

- Create the VRF on the switch.
- Enable Auto-sense on the port that connects to the SD-WAN Appliance.

About This Task

Auto-sense automatically configures BGP and RIP redistribution policies between the selected VRFs and the dynamic SD-WAN VRF.

This feature applies only to VRFs that exist on the same Fabric Extend switch that connects to the SD-WAN Appliance, and only for local or direct subnets on the same VRF. IS-IS does not advertise the LBO-injected routes further than the local switch.



Note

You cannot manually configure BGP on a VRF with an LBO configuration when the BGP instance origin is AUTO-SENSE.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:
`enable`

`configure terminal`

`router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable LBO on the VRF:
`sd-wan-local-breakout`

ExtremeCloud SD-WAN Configuration Using EDM

The topics in this section provide the EDM configuration specific to interoperability with SD-WAN.

Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration

Before You Begin

On switches that support only a single active VRF, ensure a user-defined VRF with a different name does not exist.

About This Task

Perform this procedure to specify the VRF that Auto-sense uses for the SD-WAN configuration. On models that support a single active VRF, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.



Note

If the running switch configuration includes the Auto-sense-managed SD-WAN VRF and you complete this procedure, the switch renames the dynamic SD-WAN VRF to the specified name, and converts the configuration to VRF to a static (CONFIG) origin. The associated BGP instance remains Auto-sense-managed until you enable BGP in that VRF, which converts the BGP instance to a static (CONFIG) origin, removes Auto-sense-applied BGP configuration, and makes the configuration persistent. For information about configuring BGP, see [Configure BGP](#) on page 618.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **SD-WAN** tab.
4. In **SD-WAN VRF**, type the name of the VRF.
5. Select **Apply**.

SD-WAN Field Descriptions

Use the data in the following table to use the **SD-WAN** tab.

Name	Description
SD-WAN Area Note: This field only displays on a switch with boundary-node abilities.	Specifies the IS-IS area, home or remote, where Auto-sense creates all ExtremeCloud SD-WAN learned tunnels. By default, Auto-sense creates SD-WAN logical interfaces in the home area. You can also configure exceptions for specific tunnels. For more information, see Configure the IS-IS Area for a Specific Tunnel on page 3316.
Sd-Wan Vrf	Specifies the VRF name that Auto-sense uses for the SD-WAN configuration. On models that support a single active VRF, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.

Configure the IS-IS Area for a Specific Tunnel

About This Task

**Note**

You can only use this procedure on a switch with boundary-node abilities.

Perform this procedure to specify the IS-IS area where Auto-sense creates a specific ExtremeCloud SD-WAN-learned tunnel. This configuration overrides a global area configuration for all learned tunnels.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Sd-Wan Multi-Area** tab.
4. Select **Insert**.
5. In **Ip**, type the destination IP address.
6. In **Area**, select **home** or **remote**.
7. Select **Insert**.

Sd-Wan Multi-Area Field Descriptions

Use the data in the following table to use the **Sd-Wan Multi-Area** tab.

Name	Description
Ip	Displays the destination IP address for an ExtremeCloud SD-WAN-learned tunnel.
Area	Specifies the IS-IS area where Auto-sense creates the tunnel. By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Configure Auto-sense to Create All Learned Tunnels in the Remote Area

About This Task



Note

You can only use this procedure on a switch with boundary-node abilities.

Perform this procedure to create all ExtremeCloud SD-WAN learned tunnels in the IS-IS remote area.

By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **SD-WAN** tab.
4. Select **remote**.
5. Select **Apply**.

SD-WAN Field Descriptions

Use the data in the following table to use the **SD-WAN** tab.

Name	Description
SD-WAN Area Note: This field only displays on a switch with boundary-node abilities.	Specifies the IS-IS area, home or remote, where Auto-sense creates all ExtremeCloud SD-WAN learned tunnels. By default, Auto-sense creates SD-WAN logical interfaces in the home area. You can also configure exceptions for specific tunnels. For more information, see Configure the IS-IS Area for a Specific Tunnel on page 3316.
Sd-Wan Vrf	Specifies the VRF name that Auto-sense uses for the SD-WAN configuration. On models that support a single active VRF, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.

Select a VRF for Local Breakout Configuration

Enable Auto-sense to manage the VRF after it detects an ExtremeCloud SD-WAN peer.



Note

This procedure does not apply to switch models that support only a single active VRF with IP configuration.

Before You Begin

- Create the VRF on the switch.
- Enable Auto-sense on the port that connects to the SD-WAN Appliance.

About This Task

Auto-sense automatically configures BGP and RIP redistribution policies between the selected VRFs and the dynamic SD-WAN VRF.

This feature applies only to VRFs that exist on the same Fabric Extend switch that connects to the SD-WAN Appliance, and only for local or direct subnets on the same VRF. IS-IS does not advertise the LBO-injected routes further than the local switch.



Note

You cannot manually configure BGP on a VRF with an LBO configuration when the BGP instance origin is AUTO-SENSE.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **VRF** tab.
4. In **Sd-Wan Local Breakout** for the applicable VRF row, select **true** to enable LBO on the VRF.
5. Select **Apply**.

VRF Field Descriptions

Use the data in the following table to help you use the **VRF** tab.

Name	Description
Id	Specifies the ID number of the VRF instance. VRF ID 0 is reserved for the GlobalRouter.
Name	Names the VRF instance.
ContextName	Identifies the VRF. The SNMPv2 Community String or SNMPv3 contextName denotes the VRF context and is used to logically separate the MIB port management.
TrapEnable	Enables the VRF to send VRF Lite-related traps (VrfUp and VrfDown). The default is enabled.
MaxRoutes Note: Exception: not supported on 5320-16P-4XE, 5320-16P-4XE-DC, 5320-24P-8XE, or 5320-24T-8XE.	Configures the maximum number of routes allowed for the VRF, which varies depending on the hardware. For scaling information, see <i>Fabric Engine Release Notes</i> . The default value varies for the GlobalRouter and non-default VRFs, depending on your hardware platform.
RpTrigger	Specifies the Routing Protocol (RP) triggers for the VRF. The triggers are used to initiate or shutdown routing protocols on a VRF. You can act upon multiple RPs simultaneously. You can also use this option to bring individual RPs up in steps.

Name	Description
Segmented Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Specifies if the VRF is a Segmented VRF. The default is false, which indicates the VRF is not segmented.
MaxRoutesTrapEnable	Enables the generation of the VRF Max Routes Exceeded traps. The default is enabled.
Ipv6MaxRoutes Note: Exception: not supported on 5320-16P-4XE, 5320-16P-4XE-DC, 5320-24P-8XE, or 5320-24T-8XE.	Configures the maximum number of IPv6 routes allowed for the VRF, which varies depending on the hardware. For scaling information, see <i>Fabric Engine Release Notes</i> . The default value varies for the GlobalRouter and non-default VRFs, depending on your hardware platform.
Ipv6MaxRoutesTrapEnable	Enables SNMP trap generation after the maximum number of IPv6 routes are reached. The default is enabled.
Active	Displays if the VRF is active (true). Some switch models support only a single active VRF with IP configuration; the VRF can be the GlobalRouter (GRT) or a non-default VRF. For information about VRF support, see <i>Fabric Engine Release Notes</i> or <i>Fabric Engine and VOSS Feature Matrix</i> .
Origin	Displays the VRF origin: <ul style="list-style-type: none"> • config • dynamic • segmented The Default VRF and the MGMT VRF are considered DYNAMIC. Note: Only manually configured VRFs are visible in running-config.

Name	Description
Sd-Wan Local Breakout Note: Not supported on switch models that support only a single active VRF with IP configuration.	Enables Local Breakout on the VRF. The default is false (disabled).
BgpOrigin	Specifies the origin of the BGP configuration in the VRF. config indicates the BGP instance was created manually. autoSense indicates the BGP instance was generated from Auto-sense. default is for Global Routing Table (GRT) only and indicates no BGP configuration.



Secure Shell

[Secure Shell Fundamentals](#) on page 3323

[Configure Secure Shell using CLI](#) on page 3337

[Secure Shell configuration using Enterprise Device Manager](#) on page 3362

Table 283: Secure Shell product support

Feature	Product	Release introduced
Secure Shell (SSH) server (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	Fabric Engine 9.3
Secure Shell (SSH) client (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 283: Secure Shell product support (continued)

Feature	Product	Release introduced
Secure Sockets Layer (SSL) certificate management	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
SSH server (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
SSH client (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 283: Secure Shell product support (continued)

Feature	Product	Release introduced
SSH client disable	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
SSH rekey	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Secure Shell Fundamentals

Methods of remote access such as Telnet or FTP generate unencrypted traffic. Anyone that can see the network traffic can see all data, including passwords and user names. Secure Shell (SSH) is a client and server protocol that specifies the way to conduct secure communications over a network.



Note

7830 Series does not support SSH connection to the backbone MAC.

Secure Shell can replace Telnet. Secure File Transfer Protocol (SFTP) can replace FTP with an encrypted alternative.



Note

If both SSH and SFTP are concurrently active, you have the ability to disable SFTP while allowing SSH to remain active. For more information, see [Disabling SFTP without disabling SSH](#) on page 3358.

The switch software supports Secure CoPy protocol (SCP), which is a secure file transfer protocol. Use SCP to securely transfer files between a local host and a remote host. SCP is in off state by default, but you can turn it on when you enable SSH using the **boot config flags** command in the global config mode. The switch supports SCP only as an SCP server, which means that clients can send files to the switch or can request files from the switch. Secure CoPy (SCP) can replace FTP with an encrypted alternative.

Secure Shell supports a variety of the different public and private key encryption schemes available. Using the public key of the host server, the client and server negotiate to generate a session key known only to the client and the server. This one-time key encrypts all traffic between the client and the server. The switch supports Secure Shell version 2 (SSHv2).

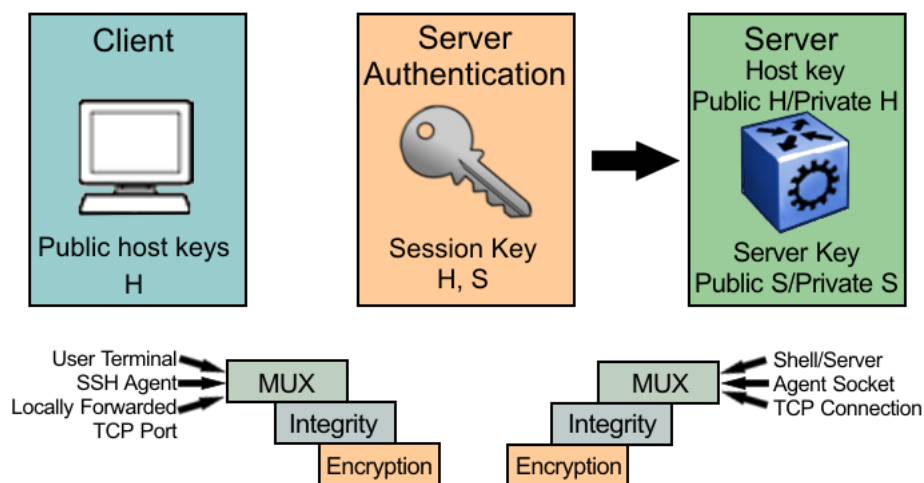


Figure 218: Overview of the SSHv2 protocol

By using a combination of host, server, and session keys, the SSHv2 protocol can provide strong authentication and secure communication over an insecure network, offering protection from the following security risks:

- IP spoofing
- IP source routing
- Domain name server (DNS) spoofing
- adversary-in-the-middle/TCP hijacking attacks
- Eavesdropping and password sniffing

Even if network security is compromised, traffic cannot be played back or decrypted, and the connection cannot be hijacked.

The SSH secure channel of communication does not provide protection against break-in attempts or denial-of-service (DoS) attacks.

With the SSHv2 server in the switch, you can use an SSHv2 client to make a secure connection to the switch and work with commercially available SSHv2 clients. For more information about supported clients, see [Third-Party SSH and SCP Client Software](#) on page 3333. The switch also supports outbound connections to remote SSHv2 servers to provide complete inbound and outbound secure access.

Outbound Connections

The SSHv2 client supports SSHv2 DSA public key authentication and password authentication.

**Note**

For certain switches in enhanced secure mode, all sensitive files are protected. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

**Note**

You must enable SSH globally before you can generate SSH DSA user keys.

The SSHv2 client is a secure replacement for outbound Telnet. Password authentication is the easiest way to use the SSHv2 client feature.

Instead of password authentication, you can use DSA public key authentication between the SSHv2 client and an SSHv2 server. Before you can perform a public key authentication, you must generate the key pair files and distribute the key files to all the SSHv2 server systems. Because passphrase encrypts and further protects the key files, you must provide a passphrase to decrypt the key files as part of the DSA authentication.

**Note**

SSH RSA and DSA public and private keys are copied from `/intflash/shared` to `/intflash/.ssh`.

To attempt public key authentication, the SSHv2 client looks for the associated DSA key pair files in the `/intflash/.ssh` directory. If no DSA key pair files are found, the SSHv2 client automatically prompts you for password authentication. If the SSHv2 client succeeds with the authentication, then a new secured SSHv2 session is established to the remote SSHv2 server. For more information, see [.DSA Authentication Access Level and File Name](#) on page 3335.

**Important**

If you configure the DSA user key with a passphrase but you do not supply the correct passphrase when you try to make the SSHv2 connection, then the system defaults back to the password authentication. If the SSHv2 client succeeds with the authentication, then a new secured SSHv2 session is established to the remote SSHv2 server.

SSH Version 2

SSH version 2 (SSHv2) protocol is a complete rewrite of the SSHv1 protocol. In SSHv2 the functions are divided among three layers:

- SSH Transport Layer (SSH-TRANS)

The SSH Transport Layer manages the server authentication and provides the initial connection between the client and the server. After the connection is established, the Transport Layer provides a secure, full-duplex connection between the client and server.

- SSH Authentication Protocol (SSH-AUTH)

The SSH Authentication Protocol runs on top of the SSH Transport Layer and authenticates the client-side user to the server. SSH-AUTH defines three authentication methods: public key, host-based, and password. SSH-AUTH provides a single authenticated tunnel for the SSH connection protocol.

- SSH Connection Protocol (SSH-CONN)

The SSH Connection Protocol runs on top of the SSH Transport Layer and user authentication protocols. SSH-CONN provides interactive logon sessions, remote execution of commands, forwarded TCP/IP connections, and forwarded X11 connections. These services are multiplexed into the single encrypted tunnel provided by the SSH transport layer.

The following figure shows the three layers of the SSHv2 protocol.




Figure 219: Separate SSH version 2 protocols

The modular approach of SSHv2 improves on the security, performance, and portability of the SSHv1 protocol.



Important

The SSHv1 and SSHv2 protocols are not compatible. The switch does not support SSHv1.

Security Features

The SSHv2 protocol supports the following security features:

- **Authentication.** This feature determines, in a reliable way, the SSHv2 client. During the log on process, the SSHv2 client is queried for a digital proof of identity.

Supported authentications with the switch as a server for SSHv2, are: RSA, DSA, and passwords. Supported authentications with the switch as a client for SSHv2, are: DSA and passwords. The switch does not support RSA when the switch acts as a client.

When the switch acts as an SSH server, by default the switch allows a maximum of only four sessions, although it can accommodate up to eight sessions at a time. However, only one SSH public key encryption per access level is allowed at a time. For instance, if multiple SSH public key encryption clients need to connect to the server with the same access level, such as rwa, then the clients must connect to the server one-by-one as the switch only supports one public key per access level.

- **Encryption.** The SSHv2 server uses encryption algorithms to scramble data and render it unintelligible except to the receiver.
- **Integrity.** This feature guarantees that the data transmits from the sender to the receiver without alterations. If a third party captures and modifies the traffic, the SSHv2 server detects this alteration.



Note

SCP is supported for RWA users only. RW or R level will not work and the switch logs a message on the device.

SSHv2 Considerations using EDM

You must use CLI to initially configure SSHv2. You can use Enterprise Device Manager (EDM) to change the SSHv2 configuration parameters. CLI is the user interface for SSHv2 configuration and use the console port to configure the SSHv2 parameters.

**Important**

Do not enable SSHv2 secure mode using Configuration and Orchestration Manager (COM). If you enable SSHv2 secure mode, then the system disables Simple Network Management Protocol (SNMP). This locks you out of a COM session. Enable SSH secure mode using CLI or EDM.

SSHv2 secure mode is different from enhanced secure mode and hsecure. SSHv2 secure mode disables unsecure management protocols on the device such as FTP, SNMP, Telnet, and TFTP. SSHv2 secure mode is enabled through the **ssh secure** command.

When you enable SSHv2 secure mode, the system disables FTP, SNMPv1, SNMPv2, SNMPv3, Telnet and TFTP. After SSHv2 secure mode is enabled, you can choose to enable individual non-secure protocols. However, after you save the configuration and restart the system, the non-secure protocol is again disabled, even though it is shown as enabled in the configuration file. After you enable SSHv2 secure mode, you cannot enable non-secure protocols by disabling SSHv2 secure mode.

You can disable block-snmp after you enable SSHv2 secure mode, and you can connect again using COM.

User ID Logs

User ID log of an SSH session established by SCP client

The switch logs the user ID of an SSH session initiated by the SCP client. If an SCP client establishes an SSH session, the system displays the message in the following format:

```
CP1 [08/06/15 09:43:42.230:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SSH user
authentication succeeded for user rwa on host 10.68.231.194
CP1 [08/06/15 09:43:42.232:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SSH SCP session
start by user rwa on host 10.68.231.194
CP1 [08/06/15 09:43:44.020:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SCP session
closed by user rwa on host 10.68.231.194
CP1 [08/06/15 09:43:44.021:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SSH session
closed by user rwa on host 10.68.231.194
```

In the preceding example log output, rwa is the user name.

User ID log of an SSH session established by SFTP

The switch logs the user ID of an SSH session initiated by SFTP. If SFTP establishes an SSH session, the system displays the message in the following format:

```
CP1 [08/06/15 09:45:32.903:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SSH user
authentication succeeded for user rwa on host 10.68.231.194
CP1 [08/06/15 09:45:32.905:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SFTP session
start: user rwa on host 10.68.231.194
CP1 [08/06/15 09:45:46.775:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SFTP session
closed by user rwa on host 10.68.231.194
```

```
CP1 [08/06/15 09:45:46.776:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SSH SFTP
session end: user rwa on host 10.68.231.194
CP1 [08/06/15 09:45:46.776:UTC] 0x000d8602 00000000 GlobalRouter SSH INFO SSH session
closed by server for user rwa on host 10.68.231.194
```

In the preceding example log output, rwa is the user name.

User Key Files

Generating keys requires that you have free space on the flash. A typical configuration requires less than 2 kbyte of free space. Before you generate a key, verify that you have sufficient space on the flash, using the **dir** command. If the flash is full when you attempt to generate a key, the system displays an error message and the key is not generated. You must delete some unused files and regenerate the key.

If you remove only the public keys, enabling the SSH does not create new public keys.

SSHv2 password authentication uses the same login and password authentication mechanism as Telnet. The SSHv2 client also supports DSA public key authentication compatible with the switch SSHv2 server and Linux SSHv2 server for SSHv2.

If the switch is the client, use the following table to locate the DSA user key files for DSA authentication for user access level `rwa`.



Note

For certain switches in enhanced secure mode, all sensitive files are protected. The home directory for enhanced secure mode is `/intflash/shared`. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

Table 284: DSA user key files

SSH server	SSH client side	SSH server side
switch with enhanced secure mode disabled	Private and public keys by access level: <ul style="list-style-type: none"> • <code>rwa—/intflash/.ssh/id_dsa_rwa</code> (private key), <code>/intflash/.ssh/id_dsa_rwa.pub</code> (public key) • <code>rw—/intflash/.ssh/id_dsa_rw</code> (private key), <code>/intflash/.ssh/id_dsa_rw.pub</code> (public key) • <code>ro—/intflash/.ssh/id_dsa_ro</code> (private key), <code>/intflash/.ssh/id_dsa_ro.pub</code> (public key) • <code>rw1—/intflash/.ssh/id_dsa_rw1</code> (private key), <code>/intflash/.ssh/id_dsa_rw1.pub</code> (public key) • <code>rw12—/intflash/.ssh/id_dsa_rw12</code> (private key), <code>/intflash/.ssh/id_dsa_rw12.pub</code> (public key) • <code>rw13—/intflash/.ssh/id_dsa_rw13</code> (private key), <code>/intflash/.ssh/id_dsa_rw13.pub</code> (public key) 	Public keys on the server side based on access level: <ul style="list-style-type: none"> • <code>rwa—/intflash/.ssh/dsa_key_rwa</code> (public key) • <code>rw—/intflash/.ssh/dsa_key_rw</code> (public key) • <code>ro—/intflash/.ssh/dsa_key_ro</code> (public key) • <code>rw1—/intflash/.ssh/dsa_key_rw1</code> (public key) • <code>rw12—/intflash/.ssh/dsa_key_rw12</code> (public key) • <code>rw13—/intflash/.ssh/dsa_key_rw13</code> (public key)
switch with enhanced secure mode enabled	Private and public keys by access role level: <ul style="list-style-type: none"> • <code>administrator—/intflash/shared/id_dsa_admin</code> (private key), <code>/intflash/shared/id_dsa_admin.pub</code> (public key) • <code>operator —/intflash/shared/id_dsa_operator</code> (private key), <code>/intflash/shared/id_dsa_operator.pub</code> (public key) • <code>security —/intflash/shared/id_dsa_security</code> (private key), <code>/intflash/shared/id_dsa_security.pub</code> (public key) • <code>auditor —/intflash/shared/id_dsa_auditor</code> (private key), <code>/intflash/shared/id_dsa_auditor.pub</code> (public key) 	Public keys on the server side based on access level: <ul style="list-style-type: none"> • <code>administrator—/intflash/shared/dsa_key_admin</code> (public key) • <code>operator—/intflash/shared/dsa_key_operator</code> (public key) • <code>security—/intflash/shared/dsa_key_security</code> (public key) • <code>privilege—/intflash/shared/dsa_key_priv</code> (public key) • <code>auditor—/intflash/shared/dsa_key_auditor</code> (public key)

Table 284: DSA user key files (continued)

SSH server	SSH client side	SSH server side
	key), /intflash/shared/ id_dsa_auditor.pub (public key) • privilege —/intflash/shared/ id_dsa_priv (private key), /intflash/ shared/id_dsa_priv.pub (public key)	
Linux with Open SSH	~/.ssh/id_dsa (private key) file permission 400 ~/.ssh/id_dsa.pub (public key) file permission 644	~/.ssh/authorized_keys (public key) file

When you attempt to make an SSH connection from the switch, the SSHv2 client looks in its own internal flash for the public key pair files. If the key files exist, the SSHv2 client prompts you for the passphrase to decrypt the key files. If the passphrase is correct, the SSHv2 client initiates the DSA key authentication to the remote SSHv2 server. The SSHv2 client looks for the login user access level public key file on the SSHv2 server to process and validate the public key authentication. If the DSA authentication is successful, then the SSHv2 session is established.

If no matching user key pair files exist on the client side when initiating the SSHv2 session, or if the DSA authentication fails, you are automatically prompted for a password to attempt password authentication.

If the remote SSHv2 server is a Linux system, the server looks for the login user public key file ~/.ssh/authorized_keys by default for DSA authentication. For a Linux SSHv2 client, the user DSA key pair files are located in the user home directory as ~/.ssh/id_dsa and ~/.ssh/id_dsa.pub.

Block SNMP

The boot flag setting for block-snmpp (**boot config flags block-snmpp**) and the runtime configuration of SSH secure (**ssh secure**) each modify the block-snmpp boot flag. If you enable SSH secure mode, the system automatically sets the block-snmpp boot flag to true; the change takes effect immediately. After enabling SSH in secure mode, you can manually change the block-snmpp flag to false to allow both SSH and SNMP access.



Important

The block flag setting for block-snmpp blocks Simple Network Management Protocol (SNMP)v1, SNMPv2, and SNMPv3.

SCP command

Use short file names with the Secure CoPy (SCP) command. The entire SCP command, including all options, user names, and file names must not exceed 80 characters.

The switch supports incoming SCP connections but does not support outgoing connections using an SCP client.

Third-Party SSH and SCP Client Software

Tested Software

The following table describes the third-party SSH and SCP client software that has been tested but is not included with the switch software.

Table 285: Tested software

SSH Client	Secure Shell (SSH)	Secure Copy (SCP)
Tera Term Pro with TTSSH extension MS Windows	<ul style="list-style-type: none"> • Supports SSHv2. • Authentication: <ul style="list-style-type: none"> - RSA is supported when the switch acts as a server. The switch does not support RSA as a client. - DSA - Password • Provides a keygen tool. • It creates both RSA and DSA keys. 	<ul style="list-style-type: none"> • Client distribution does not include SCP client.
Secure Shell Client Windows 2000	<ul style="list-style-type: none"> • Supports SSHv2 client. • Authentication <ul style="list-style-type: none"> ◦ DSA ◦ Password • Provides a keygen tool. • It creates a DSA key in SSHv2 format. • The switch generates a log message stating that a DSA key has been generated. 	<ul style="list-style-type: none"> • Client distribution includes an SCP client that is not compatible with switch.

Table 285: Tested software (continued)

SSH Client	Secure Shell (SSH)	Secure Copy (SCP)
OpenSSH Unix Solaris 2.5 / 2.6	<ul style="list-style-type: none"> • Supports SSHv2 clients. • Authentication: <ul style="list-style-type: none"> - RSA is supported when the switch acts as a server. The switch does not support RSA as a client. - DSA - Password • Provides a keygen tool. • It creates both RSA and DSA keys. 	<ul style="list-style-type: none"> • Client distribution includes an SCP client that is supported on switch.
WinSCP		This SCP client is unsupported on the switch.

Switch As Client

The switch acting as the SSHv2 client generates a DSA public and private server key pair. The public part of the key for DSA is stored in the following location:

```
/intflash/.ssh/dsa_key_rwa
```

**Note**

For certain switches in enhanced secure mode, all sensitive files are protected. The home directory for protected files is `/intflash/shared`. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

The public part of the key must be copied to the SSH server and be named according to the naming requirement of the server.

Consult [DSA Authentication Access Level and File Name](#) on page 3335 for proper naming convention.

If a DSA key pair does not exist, you can generate the DSA key pair using the **ssh dsa-user-key [WORD<1-15>] [size <1024-1024>]** command.

You need to copy the DSA public key to the SSHv2 server that you connect to using the switch as a client. RSA is not supported when using the switch as a client, but you can use RSA when the switch is acting as the server.

**Note**

SSH RSA and DSA public and private keys are copied from `/intflash/shared` to `/intflash/.ssh`.

Switch As Server

After you install one of the SSHv2 clients you must generate a client and server key using the RSA or DSA algorithms.

To authenticate an SSHv2 client using DSA, the administrator must copy the public part of the client DSA key to `/intflash/.ssh` directory on the switch that acts as the SSHv2 server. The file that is copied over to the SSHv2 server must be named according to [DSA Authentication Access Level and File Name](#) on page 3335.

DSA Authentication Access Level and File Name

The following table lists the access levels and file names that you must use to store the SSHv2 client authentication information using DSA onto the switch that acts as the SSHv2 server.

For certain switches in enhanced secure mode, all sensitive files are protected. The home directory for enhanced secure mode is `/intflash/shared`. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

Table 286: DSA authentication access levels and file names

Client key format or WSM	Access level	File name
Client key in non IETF and IETF format with enhanced secure mode disabled Note: The switch supports IETF and non-IETF for DSA.	RWA	/intflash/.ssh/dsa_key_rwa
	RW	/intflash/.ssh/dsa_key_rw
	RO	/intflash/.ssh/dsa_key_ro
	L3	/intflash/.ssh/dsa_key_rwl3
	L2	/intflash/.ssh/dsa_key_rwl2
	L1	/intflash/.ssh/dsa_key_rwl1
Client key in enhanced secure mode	administrator	/intflash/shared/dsa_key_admin
	operator	/intflash/shared/dsa_key_operator
	security	/intflash/shared/dsa_key_security
	privilege	/intflash/shared/dsa_key_priv
	auditor	/intflash/shared/dsa_key_auditor
Client key in enhanced secure mode	administrator	/intflash/.ssh/dsa_key_admin
	operator	/intflash/.ssh/dsa_key_operator
	security	/intflash/.ssh/dsa_key_security
	privilege	/intflash/.ssh/dsa_key_priv
	auditor	/intflash/.ssh/dsa_key_auditor

The switch generates an RSA public and private server key pair. The public part of the key for RSA is stored in `/intflash/.ssh/ssh_key_rsa_pub.key`. If an RSA key pair does not exist, then the switch automatically generates one when you enable the SSH server.

To authenticate a client using RSA, the administrator must copy the public part of the client RSA key to the switch.

For a certain switches in enhanced secure mode, sensitive files are protected. You cannot copy public or private keys directly to `/intflash/.ssh`. You must import the DSA/RSA private and public key from `/intflash/shared`. For more information, see [Import DSA and RSA Private or Public Keys](#) on page 3342.

RSA Authentication Access Level and File Name

For certain switches in enhanced secure mode, all sensitive files are protected. The home directory for enhanced secure mode is `/intflash/shared`. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

The following table lists the access levels and file names you can use for storing the SSH client authentication information using RSA.

Table 287: RSA authentication access levels and file names

Client key format or WSM	Access level	File name
Client key in IETF format with enhanced secure mode disabled.	RWA	/flash/ssh/rsa_key_rwa
	RW	/flash/ssh/rsa_key_rw
	RO	/flash/ssh/rsa_key_ro
	L3	/flash/ssh/rsa_key_rwl3
	L2	/flash/ssh/rsa_key_rwl2
	L1	/flash/ssh/rsa_key_rwl1
Client key with enhanced secure mode enabled	administrator	/intflash/shared/rsa_key_admin
	operator	/intflash/shared/rsa_key_operator
	security	/intflash/shared/rsa_key_security
	privilege	/intflash/shared/rsa_key_priv
	auditor	/intflash/shared/rsa_key_auditor

SSH Rekeying

An SSHv2 server or client can use the SSH rekeying feature to force a key exchange between server and client, while changing the encryption and integrity keys. After you enable SSH rekeying, key exchanges occur after a pre-determined time interval or after the data transmitted in the session reaches the data-limit threshold. You can configure these values using the **ssh rekey** command.

SSH rekeying is optional. You can enable SSH rekey only if SSH is enabled globally. Most SSH clients and servers do not provide a rekey mechanism; do not enable SSH rekey in such cases. Active sessions shut down if the rekey fails.

**Note**

You cannot enable SSH rekey selectively for the SSH client, SSH server, Secure Copy (SCP), or Secure File Transfer Protocol (SFTP); SSH rekey is enabled for all of these functions simultaneously.

Configure Secure Shell using CLI

Use Secure Shell version 2 (SSHv2) to enable secure communications support over a network for authentication, encryption, and network Integrity.

Before You Begin

- Disable the sshd daemon. All SSHv2 commands, except enable, require that you disable the sshd daemon.
- Set the user access level to read/write/all community strings.
- Disable all nonsecure access services. As a best practice, disable the following services: Trivial File Transfer Protocol (TFTP), File Transfer Protocol (FTP), and Telnet. For more information about disabling access services, see [Enable Remote Access Services](#) on page 293.
- Use the console port to configure the SSHv2 parameters.

Enabling the SSHv2 server

Enable the SSHv2 server to provide secure communications for accessing the switch. The switch does not support SSHv1.

Before You Begin

To enable SSH, ensure to enable RSA or DSA authentication, or both using command **ssh rsa-auth** or **ssh dsa-auth**.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable the SSH server:
`boot config flags sshd`
3. Save the configuration file:
`save config`

Example

Enable the SSHv2 server:

```
Switch:1>enable
Switch:1#configure terminal
```

```
Switch:1(config)#boot config flags sshd
Switch:1(config)#save config
```

Changing the SSH server authentication mode

Use this procedure to change the SSH server authentication mode from the default of password-authentication to keyboard-interactive.

About This Task

If you enable keyboard-interactive authentication mode, the server uses that mode over other authentication methods, except for public-key authentication, if the SSH client supports it.

If you enable keyboard-interactive authentication mode, the server generates the password prompts to display to the client rather than the client generating the prompts automatically like with password-authentication.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable keyboard-interactive authentication:
`ssh keyboard-interactive-auth`

Configure SSH Authentication Type

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the authentication type to use:
`ssh authentication-type {[aead-aes-128-gcm-ssh] [aead-aes-256-gcm-ssh]
[hmac-sha1] [hmac-sha2-256]}`

Variable Definitions

The following table defines parameters for the **ssh authentication-type** command.

Variable	Value
<code>[aead-aes-128-gcm-ssh] [aead-aes-256-gcm-ssh] [hmac-sha1] [hmac-sha2-256]</code>	<p>Specifies the authentication type. Select from one of the following:</p> <ul style="list-style-type: none">• aead-aes-128-gcm-ssh• aead-aes-256-gcm-ssh• hmac-sha1• hmac-sha2-256 <p>By default, all types are enabled.</p>

Enable DSA Authentication

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable DSA authentication:
`ssh dsa-auth`

Generate the DSA Host Key

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Generate a new DSA host key:
`ssh dsa-host-key [<1024-1024>]`

Variable Definitions

The following table defines parameters for the **ssh dsa-host-key** command.

Variable	Value
<code><1024-1024></code>	The DSA host key size is 1024.

Generate a DSA User Key

Before You Begin

You must enable SSH globally before you can generate DSA user keys.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create the DSA user key:

```
ssh dsa-user-key WORD<1-15> [size <1024-1024>]
```

Variable Definitions

The following table defines parameters for the **ssh dsa-user-key** command.

Variable	Value
<i>WORD</i> <1–15>	<p>Specifies the user access level.</p> <p>You must enable SSH globally before you can generate SSH DSA user keys.</p> <p>If enhanced secure mode is disabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • rwa — Specifies read-write-all. • rw — Specifies read-write. • ro — Specifies read-only. • rwl1 — Specifies read-write for Layer 1. • rwl2 — Specifies read-write for Layer 2. • rwl3 — Specifies read-write for Layer 3. <p>If you enable enhanced secure mode, the switch uses role-based authentication. You associate each user name with a specific role and the appropriate authorization rights to commands based on that role.</p> <p>If enhanced secure mode is enabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • admin—Specifies a user role with access to all of the configurations, show commands, and the ability to view the log file and security commands. The administrator role is the highest level of user roles. • operator—Specifies a user role with access to all of the configurations for packet forwarding on Layer 2 and Layer 3, and has access to show commands to view the configuration, but cannot view the audit logs and cannot access security and password commands. • auditor—Specifies a user role that can view log files and view all configurations, except password configuration. • security—Specifies a user role with access only to security settings and the ability to view the configurations. • priv—Specifies a user role with access to all of the commands that the administrator has access to, and is referred to as an emergency-admin. However, the user with the privilege role must be authenticated within the switch locally. RADIUS and TACACS+ authentication is not accessible. A user role at the privilege level must login to the switch through the console port only.
<i>size</i> <1024–1024>	Specifies the size of the DSA user key.

Import DSA and RSA Private or Public Keys

About This Task

Use this task to import SSH RSA and DSA public and private keys.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enter the following command to install the DSA or RSA public or private key:

```
ssh install-user-key WORD<1-15> {rwa | rw | ro | rwl1 | rwl2 | rwl3  
| admin | operator | auditor | security | priv} WORD<1-15> {public |  
private} WORD<1-15> {dsa | rsa}
```

Examples

```
Switch:1(config)#ssh install-user-key admin public dsa  
Switch(config)#1 2021-07-22T14:09:43.278+03:00 5520-24X-VOSS CP1 - 0x0041460b - 00000000  
GlobalRouter CLOUD_AGENT INFO  
Successfully installed SSH public key to path </intflash/.ssh/dsa_key_admin>.  
Switch:1(config)#ssh install-user-key admin private dsa  
Info: Successfully installed the private key to path /intflash/.ssh/id_dsa_admin  
Switch:1(config)#1 2021-07-22T14:08:00.354+03:00 5520-24X-VOSS CP1 - 0x0041460b -  
00000000  
GlobalRouter CLOUD_AGENT INFO Successfully installed SSH private key to path </  
intflash/.ssh/id_dsa_admin>.
```

Variable Definitions

The following table defines parameters for the **ssh install-user-key** command.

Variable	Value
<code>WORD<1-15>{ rwa rw ro rwl1 rwl2 rwl3, enhanced-secured mode : admin operator auditor security priv }</code>	<p>Specifies the user access level.</p> <p>You must enable SSH globally before you can generate SSH DSA user keys.</p> <p>If enhanced secure mode is disabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • rwa — Specifies read-write-all. • rw — Specifies read-write. • ro — Specifies read-only. • rwl1 — Specifies read-write for Layer 1. • rwl2 — Specifies read-write for Layer 2. • rwl3 — Specifies read-write for Layer 3. <p>If you enable enhanced secure mode, the switch uses role-based authentication. You associate each user name with a specific role and the appropriate authorization rights to commands based on that role.</p> <p>If enhanced secure mode is enabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • admin—Specifies a user role with access to all of the configurations, show commands, and the ability to view the log file and security commands. The administrator role is the highest level of user roles. • operator—Specifies a user role with access to all of the configurations for packet forwarding on Layer 2 and Layer 3, and has access to show commands to view the configuration, but cannot view the audit logs and cannot access security and password commands. • auditor—Specifies a user role that can view log files and view all configurations, except password configuration. • security—Specifies a user role with access only to security settings and the ability to view the configurations. • priv—Specifies a user role with access to all of the commands that the administrator has access to, and is referred to as an emergency-admin. However, the user with the privilege role must be authenticated within the switch locally. RADIUS and TACACS+ authentication is not accessible. A user

Variable	Value
	role at the privilege level must login to the switch through the console port only.
<i>WORD</i> <1-15>{ <i>public</i> <i>private</i> }	Specifies the public key or the private key type to copy from /intflash/shared to /intflash/.ssh.
<i>WORD</i> <1-15>{ <i>dsa</i> <i>rsa</i> }	Specifies the DSA or RSA signature algorithm for the public key or the private key to copy.

Configure Encryption Type

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the type of encryption to use:

```
ssh encryption-type {[3des-cbc][aead-aes-128-gcm-ssh ][aead-aes-256-  
gcm-ssh] [aes128-cbc][aes128-ctr][aes192-cbc][aes192-ctr][aes256-cbc]  
[aes256-ctr][blowfish-cbc] [rijndael128-cbc][rijndael192-cbc]}
```

Variable Definitions

The following table defines parameters for the **ssh encryption-type** command.

Variable	Value
<code>[3des-cbc] [aead-aes-128-gcm-ssh] [aead-aes-256-gcm-ssh] [aes128-cbc] [aes128-ctr] [aes192-cbc] [aes192-ctr] [aes256-cbc] [aes256-ctr] [blowfish-cbc] [rijndael128-cbc] [rijndael192-cbc]</code>	<p>Configures the encryption-type. Select an encryption-type from one of the following:</p> <ul style="list-style-type: none"> • 3des-cbc • aead-aes-128-gcm-ssh • aead-aes-256-gcm-ssh • aes128-cbc • aes128-ctr • aes192-cbc • aes192-ctr • aes256-cbc • aes256-ctr • blowfish-cbc • rijndael128-cbc • rijndael192-cbc <p>By default, all encryption types are enabled.</p> <p>If you configure the switch in Enhanced Secure Mode, the following encryption types are disabled by default:</p> <ul style="list-style-type: none"> • 3des-cbc • aes192-cbc • aes192-ctr • blowfish-cbc • rijndael128-cbc • rijndael192-cbc

Configure the SSH Key Exchange Method

About This Task

If you configure the switch in Enhanced Secure Mode, the following key exchange methods are disabled by default:

- diffie-hellman-group-exchange-sha256
- diffie-hellman-group14-sha1

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the key-exchange to use:

```
ssh key-exchange-method {[diffie-hellman-group-exchange-sha256]
[diffie-hellman-group14-sha1] [diffie-hellman-group14-sha256] [diffie-
hellman-group16-sha512] [diffie-hellman-group18-sha512]}
```

Configure the Maximum Number of SSH Sessions

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the maximum number of SSH sessions:
`ssh max-sessions <0-8>`

Variable Definitions

The following table defines parameters for the **ssh max-sessions** command.

Variable	Value
<0-8>	Specifies the maximum number of SSH sessions allowed. A value from 0 to 8. Default is 4.

Clear SSH Sessions

About This Task

Perform this procedure to clear SSH sessions on the switch, including Secure Copy (SCP) and Secure FTP (SFTP) sessions.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Clear a session using one of the following commands:
 - Clear an SSH session: `clear ssh <0-7>`
 - Clear an SCP session: `clear ssh scp <0-7>`
 - Clear an SFTP session: `clear ssh sftp <0-7>`

Enable SSH Password Authentication

Enable password authentication. The default is enabled.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable password authentication:
`ssh pass-auth`

Configure the SSH Connection Port

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the SSH connection port:
`ssh port <22,1024..49151>`

Variable Definitions

The following table defines parameters for the **ssh port** command.

Variable	Value
<22,1024-49151>	<p>Specifies the TCP port number. The default is 22.</p> <p>Important: You cannot configure TCP port 6000 as the SSH connection port.</p>

Enable RSA Authentication

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the RSA authentication:
`ssh rsa-auth`

Generate an RSA Host Key

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Generate a new RSA host key:
`ssh rsa-host-key [<1024-2048>]`

Variable Definitions

The following table defines parameters for the **ssh rsa-host-key** command.

Variable	Value
<1024-2048>	Specify an optional key size from 1024 to 2048. The default is 2048.

Generate an RSA User Key

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Generate a new RSA user key.
`ssh rsa-user-key WORD<1-15> [size <1024-2048>]`

Variable Definitions

The following table defines parameters for the **ssh rsa-user-key** command.

Variable	Value
<i>WORD</i> <1-15>	<p>Specifies the user access level.</p> <p>If enhanced secure mode is disabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • rwa — Specifies read-write-all • rw — Specifies read-write • ro — Specifies read-only • rwl1 — Specifies read-write for Layer 1 • rwl2 — Specifies read-write for Layer 2 • rwl3 — Specifies read-write for Layer 3 <p>If you enable enhanced secure mode, the switch uses role-based authentication. You associate each user name with a specific role and the appropriate authorization rights to commands based on that role.</p> <p>If enhanced secure mode is enabled, the value user access levels for the switch are:</p> <ul style="list-style-type: none"> • admin—Specifies a user role with access to all of the configurations, show commands, and the ability to view the log file and security commands. The administrator role is the highest level of user roles. • operator—Specifies a user role with access to all of the configurations for packet forwarding on Layer 2 and Layer 3, and has access to show commands to view the configuration, but cannot view the audit logs and cannot access security and password commands. • auditor—Specifies a user role that can view log files and view all configurations, except password configuration. • security—Specifies a user role with access only to security settings and the ability to view the configurations • priv—Specifies a user role with access to all of the commands that the administrator has access to, and is referred to as an emergency-admin. However, the user with the privilege role must be authenticated within the switch locally. RADIUS and TACACS+ authentication is not accessible. A user role at the privilege level must log in to the switch through the console port only.
<i>size</i> <1024-1024>	Specifies the size of the RSA user key.

Enable SSH Secure Mode

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable SSH secure mode:
`ssh secure`

Configure the SSH version

Configure Secure Shell version 2 (SSHv2). The switch does not support SSHv1.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the SSH version:
`ssh version <v2only>`

Configure the SSH Authentication Timeout

Configure the timeout for SSH connection authentication.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the authentication timeout:
`ssh timeout <1-120>`

Variable Definitions

The following table defines parameters for the **ssh timeout** command.

Variable	Value
<1-120>	Specifies the SSH connection authentication timeout in seconds. Default is 60 seconds.

Resetting the SSH

About This Task

Perform this task to terminate all SSH sessions and restart SSH server.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Reset the SSH.

```
ssh reset
```

Configure X.509 V3 Authentication

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable X.509 V3 authentication:

```
ssh x509v3-auth enable
```

3. Configure X.509 V3 revocation:

```
ssh x509v3-auth revocation-check-method {none | ocsp}
```

4. Configure the X.509 V3 user name:

```
ssh x509v3-auth username {overwrite | strip-domain | use-domain  
WORD<1-254>}
```

5. Configure the X.509 V3 CA trustpoint name:

```
ssh x509v3-auth ca-name WORD<1-45>
```

6. Configure the X.509 V3 digital certificate subject name to use as the identity certificate:

```
ssh x509v3-auth cert-subject-name WORD<1-45>
```

Example

Display the certificate authority details:

```
Switch:1(config)#show certificate ca
```

CA table entry

Name	:	823-pki[auto-installed]
CommonName	:	CaA2-1
KeyName	:	pki
SubjectName	:	823
CaUrl	:	
UsePost	:	0
SubjectCertValidityDays	:	0
Action	:	(null)
LastActionStatus	:	(null)
LastActionFailureReason	:	
CA-Auth Sha256Fingerprint	:	

```
UsedFor          :    SSH-X509

CA table entry
Name             :    a1
CommonName       :    CaA1
KeyName          :    rsa_2048
SubjectName      :
CaUrl            :    http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
UsePost          :    1
SubjectCertValidityDays :    365
Action           :    (null)
LastActionStatus :    (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
bd9bb74b3f4d75e86113222a8d291b6349c7a42c457e487b9be0a48b4f09cc7c
UsedFor          :

CA table entry
Name             :    a2
CommonName       :    CaA2
KeyName          :    pki2
SubjectName      :    822
CaUrl            :    http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
UsePost          :    1
SubjectCertValidityDays :    365
Action           :    (null)
LastActionStatus :    (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
0ccb8d0c38d36cf427187f0e1dd380536c078fd6fae39ec9872187327912056b
UsedFor          :    Default
```

Variable Definitions

The following table defines parameters for the **ssh x509v3-auth** command.

Variable	Value
<code><none oscp></code>	Specifies the X.509 V3 authentication revocation check method. The default is OSCP. <ul style="list-style-type: none"> none - Specifies no revocation check method. oscp - Specifies Online Certificate Status Protocol (OCSP) as revocation check method.
<code>overwrite strip-domain use-domain WORD<1-254></code>	Specifies the X.509 V3 user name configuration. The default is disabled. <ul style="list-style-type: none"> overwrite - Specifies the switch to send the principal name and domain name from the certificate to the RADIUS server for authorization. strip-domain - Specifies the switch to send the principal name from the certificate without the domain name to the RADIUS server for authorization. use-domain WORD<1-254> - Specifies the switch to send the principal name from the certificate, with the domain name you entered to the RADIUS server for authorization.
<code>ca-name WORD<1-45></code>	Specifies the X.509 V3 CA trustpoint name.
<code>cert-subject-nameWORD<1-45></code>	Specifies the digital certificate subject name to be used as the identity certificate.

Verify and Display SSH Configuration Information

Verify that SSH services are enabled on the switch and display SSH configuration information to ensure that the SSH parameters are properly configured.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Verify that SSH services are enabled and view the SSH configuration:

```
show ssh <global|session>
```

Example

Display global system SSH information:

```
Switch:1(config)#show ssh global
Total Active Sessions      : 1
version                   : v2only
port                      : 22
```

```

max-sessions          : 4
timeout              : 60
action rsa-host key   : rsa-hostkeysize 2048
action dsa-host key   : dsa-hostkeysize 1024
rsa-auth              : true
dsa-auth              : true
pass-auth             : true
keyboard-interactive-auth : false
x509-auth             : true
x509-auth Trustpoint CA Name :
x509-auth Identity Subject Name : not configured
x509-auth overwrite   : false
x509-auth strip-domain : false
x509-auth use-domain   : -
x509-auth revocation-check-method : OCSP
sftp enable           : true

client enable         : true

enable                : true
authentication-type    : aead-aes-128-gcm-ssh aead-aes-256-gcm-ssh hmac-sha1 hmac-sha2-256
encryption-type        : 3des-cbc aead-aes-128-gcm-ssh aead-aes-256-gcm-ssh aes128-cbc aes128-ctr
                        aes192-cbc aes192-ctr aes256-cbc aes256-ctr blowfish-cbc rijndael128-cbc
                        rijndael192-cbc
key-exchange-method    : diffie-hellman-group-exchange-sha256 diffie-hellman-group14-sha1
                        diffie-hellman-group14-sha256 diffie-hellman-group16-sha512
                        diffie-hellman-group18-sha512

```

Variable Definitions

The following table defines parameters for the **show ssh** command.

Variable	Value
<i>global</i>	Display global system SSH information.
<i>session</i>	Display the current session SSH information.

Connect to a Remote Host using the SSH Client

Make an SSH connection to a remote host.

Before You Begin

Enable the SSH server on the remote host.

About This Task

You can also use the IS-IS host name or IS-IS system ID with the SSH client to establish a connection to another switch in the SPBM cloud.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Connect to a remote host:

```
ssh WORD<1-256> -l WORD<1-32> [-p <1-49151>]
```

Examples

Connect to the remote host using an IP address:

```
Switch:1>ssh 192.0.2.1 -l rwa
```


Connect to the SSH server on a remote switch using the IS-IS host name or system ID, as displayed in the **show isis spbm unicast-fib** command output. This example shows a partial ssh connection with only the relevant information.

```
Switch:1#show isis spbm unicast-fib
```

```
=====
SPBM UNICAST FIB ENTRY INFO
=====
DESTINATION      BVLAN   SYSID      HOST-NAME      OUTGOING      COST      AREA      AREA-NAME
ADDRESS
-----
00:04:96:b0:78:84 4051     0004.96b0.7884 VSP-4900-48P   1/1           2000      HOME
f2:f7:11:ff:ff:ff 4051     0004.96b0.7884 VSP-4900-48P   1/1           2000      HOME
00:04:96:b0:78:84 4052     0004.96b0.7884 VSP-4900-48P   1/1           2000      HOME
f2:f7:11:ff:ff:ff 4052     0004.96b0.7884 VSP-4900-48P   1/1           2000      HOME
=====
```

```
#Host name:
```

```
Switch:1>ssh VSP-4900-48P -l rwa
```

```
#System ID:
```

```
Switch:1>ssh 0004.96b0.7884 -l rwa
```

```
Trying ssh to ISIS system-id 0004.96b0.7884 as user rwa
```

```
#example text not included
```

```
rwa@fe80::0004:96b0:7884 %linuxsshtobmac's password:
```

```
VSP-4900-48P:1>
```

Variable Definitions

The following table defines parameters for the **ssh** command.

Variable	Value
<code>-l WORD<1-32></code>	Specifies the user login name for the remote SSH server.
<code>-p <1-49151></code>	Specifies the port number to connect to the remote SSH server. The default is 22.
<code>WORD<1-256></code>	Specifies the remote SSH server using one of the following: <ul style="list-style-type: none"> • host name • IP address • IS-IS host name • IS-IS system ID <p>Note: You cannot establish an SSH connection from an IS-IS remote area.</p>

Generate User Key Files

Configure the SSH parameters to generate DSA user key files.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable SSH server.

3. Create the DSA user key file:

```
ssh dsa-user-key [WORD<1-15>][size <1024-1024>]
```

4. Enter the encryption password to protect the key file.

5. Copy the user public key file to the remote SSH servers.

**Note**

For certain switches in enhanced secure mode, the public key is copied from `/intflash/.ssh` to `/intflash/shared` after key pair generation to be available in enhanced secure mode.

6. If you are generating the compatible keys on a Linux system, use the following steps:
 - a. Create the DSA user key file:

```
ssh-keygen -t dsa
```
 - b. Copy the user public key to the remote SSH servers.

**Note**

The DSA pair key files can be generated on the Linux system and used by the SSH client on the switch.

Example

Create the DSA user key file with the user access level set to read-write-all and size of the DSA user key set to 1024 bits:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ssh dsa-user-key rwa size 1024
```

Variable Definitions

The following table defines parameters for the **ssh dsa-user-key** command.

Variable	Value
<code>WORD <1-15></code>	<p>Specifies the user access level. If enhanced secure mode is disabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • <code>rwa</code>—Specifies read-write-all. • <code>rw</code>—Specifies read-write. • <code>ro</code>—Specifies read-only • <code>rw13</code>—Specifies read-write for Layer 3. • <code>rw12</code>—Specifies read-write for Layer 2. • <code>rw11</code>—Specifies read-write for Layer 1. <p>If you enable enhanced secure mode, the switch uses role-based authentication. You associate each user name with a specific role and the appropriate authorization rights to commands based on that role.</p> <p>If enhanced secure mode is enabled, the valid user access levels for the switch are:</p> <ul style="list-style-type: none"> • <code>admin</code>—Specifies a user role with access to all of the configurations, show commands, and the ability to view the log file and security commands. The administrator role is the highest level of user roles. • <code>operator</code>—Specifies a user role with access to all of the configurations for packet forwarding on Layer 2 and Layer 3, and has access to show commands to view the configuration, but cannot view the audit logs and cannot access security and password commands. • <code>auditor</code>—Specifies a user role that can view log files and view all configurations, except password configuration. • <code>security</code>—Specifies a user role with access only to security settings and the ability to view the configurations. • <code>priv</code>—Specifies a user role with access to all of the commands that the administrator has access to, and is referred to as an emergency-admin. However, the user with the privilege role must be authenticated within the switch locally. RADIUS and TACACS+ authentication is not accessible. A user role at the privilege level must login to the switch through the console port only.
<code>size <1024-1024></code>	Specifies the size of the DSA user key. The default is 1024 bits.

Manage an SSL Certificate



Note

For certain switches in enhanced secure mode, all sensitive files are protected. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

The TLS server selects a certificate authority (CA)-signed certificate if the certificate is already installed in the Digital Certificate module.

If the server certificates are not available, the TLS server generates a new self-signed certificate at startup and uses that by default. You can choose to use an online or an offline CA-signed certificate, which takes precedence over the self-signed certificate.

For more information about SSL certificate manipulation, see [Certificate Order Priority](#) on page 3384.

About This Task

If a certificate is already present, you must confirm that it can be deleted before a new one is created.

After you create a certificate, the system logs one of the following INFO alarms:

- New default Server Certificate and Key are generated and installed
- Current Server Certificate and Key are installed

The default certificate key length for a certificate generated on the switch is 2,048 bits.



Note

The **ssl certificate [validity-period-in-days <30-3650>]** command in this procedure does not require a system reboot.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Delete a certificate:

```
no ssl certificate
```



Note

The certificate loaded in memory remains valid until you use the **ssl reset** command or reboot the system.

3. Create and install a new self-signed certificate:

```
ssl certificate [validity-period-in-days <30-3650>]
```

Variable Definitions

The following table defines parameters for the **ssl certificate** command.

Variable	Value
<i>validity-period-in-days</i> <30-3650>	Specifies an expiration time for the certificate. The default is 365 days.

Disabling SFTP without disabling SSH

Disable SFTP while allowing SSH to remain active.

Before You Begin

Enhanced secure mode must be enabled. For information about enabling enhanced secure mode, see [Enable Enhanced Secure Mode](#) on page 324.

Procedure

- 1. Enter Global Configuration mode:
`enable`
`configure terminal`
- 2. Enable the SSHv2 server:
`no ssh sftp enable`
- 3. Save the configuration file:
`save config`

Enable SSH Rekey

You cannot enable SSH rekey selectively for the SSH client, SSH server, Secure Copy (SCP), or Secure File Transfer Protocol (SFTP); SSH rekey is enabled for all of these functions simultaneously.

Before You Begin

Enable SSH globally.

Procedure

- 1. Enter Global Configuration mode:
`enable`
`configure terminal`
- 2. Enter the following command:
`ssh rekey enable`

Variable Definitions

The following table defines parameters for the **ssh rekey** command.

Variable	Value
<i>enable</i>	Enables SSH rekey globally. The default is disabled.

Configure the SSH Rekey Data Limit Threshold

Use the following procedure to limit data transmission during the session.

Procedure

- 1. Enter Global Configuration mode:
`enable`
`configure terminal`
- 2. Enter the following command:
`ssh rekey data-limit <1-6>`

Example

Configure the data limit to 2 GB:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ssh rekey data-limit 2
```

Variable Definitions

The following table defines parameters for the **ssh rekey data-limit** command.

Variable	Value
<1-6>	Specifies the SSH rekey data limit in GB. The default is 1 GB.

Configure the SSH Rekey Time Interval

Use the following procedure to configure a time interval, after which the key exchange takes place.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enter the following command:
`ssh rekey time-interval <1-6>`

Example

Configure the SSH rekey time-interval to 3 hours:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ssh rekey time-interval 3
```

Variable Definitions

The following table defines parameters for the **ssh rekey time-interval** command.

Variable	Value
<1-6>	Specifies the time-interval for SSH rekeying in hours. The default is 1 hour.

Display SSH Rekey Information

Use the following procedure to display the SSH rekey information.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Enter the following command:
`show ssh rekey`

Example

```
Switch:1>show ssh rekey
Rekey Status      : TRUE
Rekey data limit   : 1 GB
Rekey time interval : 1 hours
```

Enable or Disable the SSH Client

Before You Begin

To enable the SSH client functionality, you must enable SSH globally.

About This Task

You can disable the SSH client functionality on the switch. By default, the SSH client functionality is enabled.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable or disable the SSH client functionality:
 - To disable the client, enter the following command: `no ssh client enable`
 - To enable the client, enter the following command: `ssh client enable`

Downgrading or Upgrading from Releases that Support Different Key Sizes

Use this procedure if you need to downgrade or upgrade from a release that supports different key sizes.

Different releases can support different DSA host key, RSA host key, and DSA user key sizes. If you need to upgrade or downgrade to an earlier release that does not support the same key size, you must delete all of the keys from the `.ssh` directory and generate new keys for SSH. If you do not do this, key sizes that are no longer supported will no longer function.

You only need to perform this procedure if you have previously generated DSA host, RSA host, or DSA user keys with a release that supports different key sizes.

Procedure

1. Use the following command to disable SSH:
`no ssh`

2. From the config terminal go to the .ssh directory using the command:

```
cd /intflash/.ssh
```

3. After you upgrade or downgrade, delete the following keys from the .ssh directory.

```
ssh_dss.key
ssh_rsa.key
moc_sshc_dsa_file
moc_sshc_rsa_file
id_dsa_rwa
id_dsa_rwa.pub
id_rsa_rwa
id_rsa_rwa.pub
moc_sshc_dsa_file_fed
moc_sshc_rsa_file_fed
known_hosts
ssh_ecdsa.key
dsa_key <access level like rwa/rw/ro/admin/security/privilege/operator/auditor>,
example: dsa_key_rwa
rsa_key <access level like rwa/rw/ro/admin/security/privilege/operator/auditor>,
example: rsa_key_rwa
```

4. Generate a new DSA host key:

```
ssh dsa-host-key [<1024-1024>]
```

5. Generate a new SSH DSA user key:

```
ssh dsa-user-key WORD<1-15> [size <1024-1024>]
```

6. Generate a new RSA host key:

```
ssh rsa-host-key [<1024-2048>]
```

Secure Shell configuration using Enterprise Device Manager

Use Secure Shell version 2 (SSHv2) to enable secure communications support over a network for authentication, encryption, and network Integrity.

On IPv6 networks, the switch supports SSHv2 server only. The switch does not support outbound SSHv2 client over IPv6. On IPv4, the switch supports both SSHv2 server and SSHv2 client.

For more information, see [Change Secure Shell Parameters](#) on page 3362.

Change Secure Shell Parameters

You can use Enterprise Device Manager to change the SSHv2 configuration parameters. However, as a best practice, use the CLI to perform the initial configuration of SSHv2. The switch does not support SSHv1.

Before You Begin

- The user access level is read/write/all community strings.
- You must disable the SSH service before you configure the SSH service parameters. If the SSHv2 service is enabled, the system displays all fields dimmed until the SSH service is disabled.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **SSH**.
3. Select the **SSH** tab.
4. In the **Enable** field, select the type of SSH service you want to enable.
5. In the **Version** field, select a version.
6. In the **Port** field, type a port number.
7. In the **MaxSession** field, type the maximum number of sessions allowed.
8. In the **Timeout** field, type the timeout value in seconds.
9. From the **KeyAction** field, choose a key action.
10. In the **RsaKeySize** field, type the RSA key size in bytes.
11. In the **DSAKeySize** field, type the DSA key size in bytes.
12. Select the **RsaAuth** check box for RSA authentication.
13. Select the **DsaAuth** check box for DSA authentication.
14. Select the **PassAuth** check box for password authentication.
15. In the **AuthType** section, select the authentication types you want.
16. In the **Encryption Type** section, select the encryption types you want.
17. In the **KeyExchangeMethod** section, select the key exchange method you want.
18. Select **Apply**.

SSH Field Descriptions

Use the data in the following table to use the **SSH** tab.

Name	Description
Enable	<p>Enables, disables, or securely enables SSHv2. The options are:</p> <ul style="list-style-type: none"> • false • true • secure <p>Select false to disable SSHv2 services. Select true to enable SSHv2 services. Select secure to enable SSH and disable access services (SNMP, FTP, TFTP, and Telnet). The default is false.</p> <p>Important: Do not enable SSHv2 secure mode using EDM. Enabling secure mode disables SNMP, which locks you out of the EDM session. Enable SSHv2 secure mode using CLI.</p>
Version	<p>Configures the SSH version. The options are:</p> <ul style="list-style-type: none"> • v2only <p>The default is v2only.</p>

Name	Description
Port	Configures the SSHv2 connection port number. <22 or 1024–49151> is the port range of SSHv2. Important: You cannot configure the TCP port 6000 as SSHv2 connection port.
MaxSession	Configures the maximum number of SSHv2 sessions allowed. The value can be from 0 to 8. The default is 4.
Timeout	Configures the SSHv2 authentication connection timeout in seconds. The default is 60 seconds.
KeyAction	Configures the SSHv2 key action. The options are: <ul style="list-style-type: none"> • none • generateDsa • generateRsa • deleteDsa • deleteRsa
RsaKeySize	Configures SSHv2 RSA key size. The value can be from 1024 to 2048. The default is 2048.
DsaKeySize	Configures the SSHv2 DSA key size. The value can be from 512 to 1024. The default is 1024.
RsaAuth	Enables or disables SSHv2 RSA authentication. The default is enabled.
DsaAuth	Enables or disables SSHv2 DSA authentication. The default is enabled.
PassAuth	Enables or disables SSHv2 RSA password authentication. The default is enabled.
RekeyEnable	Enables SSH rekey globally. The default is disabled. Note: You cannot enable SSH rekey selectively for the SSH client, SSH server, Secure Copy (SCP), or Secure File Transfer Protocol (SFTP); SSH rekey is enabled for all of these functions simultaneously.
RekeyTimeInterval	Configures a time interval, after which the key exchange takes place. The default is 1 hour.
RekeyDataLimit	Configures the limit for data transmission during the session. The default is 1 GB.
SftpEnable	Enables or disables SFTP. You can use this check box to disable SFTP without affecting the SSH status. The default is enabled.
KeyboardInteractiveAuth	Changes the SSH server authentication mode from the default of password authentication to keyboard interactive.
ClientEnable	Enables SSH client functionality on the switch. By default, the SSH client functionality is enabled. To enable the SSH client functionality, SSH must be enabled globally.
X509AuthEnable	Enables SSH x509 authentication. The default is enabled.

Name	Description
X509AuthRevocationCheckMethod	Specifies the X.509 V3 authentication revocation check method. The default is OCSP. <ul style="list-style-type: none"> • none - Specifies no revocation check method. • oscp - Specifies Online Certificate Status Protocol (OCSP) as revocation check method.
X509AuthUserNameOverwrite	Enables the switch to send the principal name and domain name from the certificate to the RADIUS server for authorization. The default is disabled.
X509AuthUserNameStripDomain	Enables the switch to send the principal name from the certificate without the domain name to the RADIUS server for authorization. The default is disabled.
X509AuthUserNameUseDomain	Enables the switch to send the principal name from the certificate, with the domain name you entered to the RADIUS server for authorization.
X509AuthCertificateSubjectName	Specifies the digital certificate subject name used as identity certificate.
X509AuthCertificateCAName	Specifies the digital certificate CA trustpoint name to use.
AuthType	Specifies the authentication type. Select from one of the following: <ul style="list-style-type: none"> • hmacSha1 • hmacSha2256 • aeadAes128GcmSsh • aeadAes256GcmSsh By default, all authentication types are selected.

Name	Description
EncryptionType	<p>Configures the encryption-type. Select an encryption-type from one of the following:</p> <ul style="list-style-type: none"> • aes128Cbc • aes256Cbc • threeDesCbc • aeadAes128GcmSsh • aeadAes256GcmSsh • aes128Ctr • rijndael128Cbc • aes256Ctr • aes192Ctr • aes192Cbc • rijndael192Cbc • blowfishCbc <p>By default, all encryption types are enabled.</p> <p>If you configure the switch in Enhanced Secure Mode, the following encryption types are disabled by default:</p> <ul style="list-style-type: none"> • 3des-cbc • aes192-cbc • aes192-ctr • blowfish-cbc • rijndael128-cbc • rijndael192-cbc
KeyExchangeMethod	<p>Configures the key-exchange type. Select from one of the following:</p> <ul style="list-style-type: none"> • diffieHellmanGroup14Sha1 • diffieHellmanGroupExchangeSha256 • diffieHellmanGroup14Sha256 • diffieHellmanGgroup16Sha512 • diffieHellmanGroup18Sha512 <p>If you configure the switch in Enhanced Secure Mode, the following key exchange methods are disabled by default:</p> <ul style="list-style-type: none"> • diffieHellmanGroup14Sha1 • diffieHellmanGroupExchangeSha256



Security

[Security Fundamentals](#) on page 3367

[Security Configuration Using CLI](#) on page 3386

[Security Configuration using EDM](#) on page 3427

This section provides conceptual and procedural information about security services on the switch.

Security Fundamentals

This section provides conceptual content to help you configure and customize the security services on the switch.

For information about configuring access policies, see [Access Policies for Services](#) on page 283.

Security Overview

Security is a critical attribute of networking devices. Security features are split into two main areas:

- Control path—protects the access to the device from a management perspective.
- Data path—protects the network from malicious users by controlling access authorization to the network resources (such as servers and stations). This protection is primarily accomplished by using filters or access lists.

You can protect the control path using the following mechanism:

- logon and passwords
- access policies to specify the network and address that can use a service or daemon
- secure protocols, such as Secure Shell (SSH), Secure Copy (SCP), and the Simple Network Management Protocol version 3 (SNMPv3)
- the Message Digest 5 Algorithm (MD5) to protect routing updates, Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP)

You can protect the data path using the following mechanism

- Media Access Control (MAC) address filtering
- Layer 3 filtering, such as Internet Protocol (IP) and User Datagram Protocol (UDP)/Transmission Control Protocol (TCP) filtering

- routing policies to prevent users from accessing restricted areas of the network
- mechanisms to prevent denial-of-service (DOS) attacks

Security Modes

The switch support three security modes:

- Enhanced secure
- Hsecure
- SSH secure

Enable SSH secure mode to permit only SSH and disable all other protocols, which include Telnet, FTP, SNMP, TFTP, HTTP, and HTTPS. Enabling this mode disables Telnet, FTP, SNMP, TFTP, HTTP, and HTTPS by setting the boot flags for these protocols to off. You can over-ride the configuration and enable required protocols individually for run-time use. The administrator must enable required protocols individually for run-time use again following a reboot even if you save the configuration. This is because the SSH secure mode enable takes precedence at the time of reboot and the other protocols will be disabled even though the configuration file has them set to enabled.



Note

Disabling SSH secure mode will not automatically enable the OAM protocols that were disabled. The boot flags for the required protocols will have to be individually set to enabled.

The following table lists the differences between enhanced secure mode and hsecure mode.

Table 288: Enhanced secure mode versus hsecure mode

Feature	Enhanced secure	Hsecure
Authentication	Role-based: <ul style="list-style-type: none"> • admin • privilege • operator • security • auditor 	Access-level based: <ul style="list-style-type: none"> • rwa • rw • ro • l3 • l2 • l1
Password length	Minimum of 8 characters with the exception of the Admin, which requires a minimum of 15 characters	10 characters, minimum
Password rules	1 or 2 upper case, lower case, numeric and special characters	Minimum of 2 upper case, 2 lower case, 2 numeric and 2 special characters
Password expiration	Per-user minimum change interval is enforced, which is programmed by the Administrator	Global expiration, configured by the Admin

Table 288: Enhanced secure mode versus hsecure mode (continued)

Feature	Enhanced secure	Hsecure
Password-unique	Previous passwords and common passwords between users are prevented	The same
Password renewal	Automatic password renewal is enforced	The same
Audit logs	Audit logs are encrypted, and authorized users are able to view, modify, and delete.	Standard operation
SNMPv3	Password rules apply to SNMPv3 Auth&Priv. SNMPv3 is required (V1/V2 disabled)	SNMPv1 and SNMPv2 can be enabled.
EDM	Site Admin to enable or disable	Disabled
Telnet and FTP	Site Admin to enable or disable	The same
DOS attack Prevention		Prevents DOS attacks by filtering IP addresses and IP address ranges.

For information on Enhanced secure mode and SSH, see [Enhanced Secure Mode](#) on page 266.

hsecure Mode

The switch supports a flag called high secure (hsecure). hsecure introduces the following behaviors for passwords:

- 10-character enforcement
- aging time
- limitation of failed logon attempts
- protection mechanism to filter certain IP addresses

After you enable the hsecure flag, the software enforces the 10-character rule for all passwords. This password must contain a minimum of two uppercase characters, two lowercase characters, two numbers, and two special characters.

After you enable hsecure, the system requires you to save the configuration file and reboot the system for hsecure to take effect. If the existing password does not meet the minimum requirements for hsecure, the system prompts you to change the password during the first login.

The default user name is rwa and the default password is rwa. In hsecure, the system prompts you to change these during first login because they do not meet the minimum requirements for hsecure.

When you enable hsecure, the system disables Simple Network Management Protocol (SNMP) v1, SNMPv2 and SNMPv3. If you want to use SNMP, you must re-enable SNMP, using the command **no boot config flag block-snmp**.

Aging Enforcement

After you enable the hsecure flag, you can configure a duration after which you must change your password. You configure the duration by using the aging parameter.

For SNMP and File Transfer Protocol (FTP), after a password expires, access is denied. Before you access the system, you must change a community string to a new string consisting of more than eight characters.



Important

Consider the following after you enable the hsecure flag:

- You cannot enable the web server for Enterprise Device Manager (EDM) access.
- You cannot enable the Secure Shell (SSH) password authentication.

Filtering Mechanism

Incorrect IP source addresses as network or broadcast addresses are filtered at the virtual router interface. Source addresses 192.168.168.0 and 192.168.168.255 are discarded.

This change is valid for all IP subnets, not only for /24.

You can filter addresses only if you enable the hsecure mode.

CLI Passwords

Table 289: Password Hash

Feature	Product	Release introduced
SHA512 Password Hashing Note: When upgrading, SHA1 password hashes and custom users are retained, until a factory default reset or until the password hash level is changed. During a factory default reset, SHA2 512-bit becomes the default password hash, all custom users are deleted, and SHA1 passwords are removed.	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.8
	5420 Series	Fabric Engine 8.8
	5520 Series	Fabric Engine 8.8
	5720 Series	Fabric Engine 8.8
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.8
	VSP 7400 Series	VOSS 8.8



Important

The default passwords are documented and well known. Change the default passwords and community strings immediately after you first log on.

The switch ships with default passwords assigned for access to Command Line Interface (CLI) through a console or management session. If you have read/write/all access authority, and you are using SNMPv3, you can change passwords that are in an encrypted format. If you use EDM, you can also specify the number of available Telnet sessions.

After a factory default or if your switch has no primary or backup configuration files, a password change is required to access the CLI. The system provides three attempts to change the password, if unsuccessful you are taken back to the login prompt but are not locked out. You cannot reuse a password and your password cannot be empty. A password change is required irrespective of security mode, console, SSH, or Telnet access.

You can select SHA1 for 160-bit or SHA2 for 512-bit password hash security. You can switch the password hashing with the **password hash** command. After a hashing change, all custom users and passwords are deleted, and on first login each default user must change their default password.



Note

If you upgrade to a release that supports password hash configuration, custom users are retained until a factory default reset or until a password hash level change. During a factory default reset, all customer users are deleted, all SHA1 passwords are removed, and SHA2 becomes the new default password hash.

If you enable enhanced secure mode with the **boot config flags enhancedsecure-mode** command, you enable different access levels, along with stronger password complexity, length, and minimum change intervals. For more information on system access fundamentals and configuration, see [Managing the Switch](#) on page 246.

User-Based Policy Support

You can set up a user-based policy (UBP) system by using Enterprise Policy Manager (EPM), a RADIUS server.

EPM is an application designed to manage the traffic prioritization and network access security for business applications. It provides centralized control of advanced packet classification and the ability to priority mark, police, meter, or block traffic.

EPM 5.0 supports UBPs, which allow security administrators to establish and enforce roles and conditions for each user for all access ports in the network. The UBP feature in EPM works in conjunction with Extensible Access Protocol (EAP) technology to enhance the security of the network. Users log on to the networks and are authenticated as the network connection is established.

The UBP feature works as an extension to the Roles feature in EPM. In a UBP environment, role objects are linked directly to specific users (as RADIUS attributes), as opposed to being linked simply to device interfaces. The role object then links the users to specific policies that control the user's access to the network.

When the RADIUS server successfully authenticates a user, the device sends an EAP session start event to the EPM policy server. The policy server then sends user-based policy configuration information for the new user roles to the interface, based on the role attribute that was assigned to that user on the RADIUS server.

Denial-of-Service Attack Prevention

Table 290: Denial-of-service attack prevention product support

Feature	Product	Release introduced
Directed Broadcast	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 290: Denial-of-service attack prevention product support (continued)

Feature	Product	Release introduced
High Secure mode (hsecure boot configuration flag)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Hsecure

The switch supports a configurable flag, called high secure (hsecure). High secure mode introduces a protection mechanism to filter certain IP addresses, and two restrictions on passwords: 10-character enforcement and aging time.

If the device starts in hsecure mode with default factory settings, and no previously configured password, the system will prompt you to change the password. The new password must follow the rules mandated by high secure mode. After you enable hsecure and restart the system, if you have an invalid-length password you must change the password.

If you enable hsecure for the first time and the password file does not exist, then the device creates a normal default user name (rwa) and password (rwa). In this case, the password does not meet the minimum requirements for hsecure and as a result the system prompts you to change the password.

The following information describes hsecure mode operations:

- When you enable the hsecure flag, after a certain duration you are asked to change your password. If not configured, the aging parameter defaults to 90 days.
- For SNMP and FTP, access is denied when a password expires. You must change the community strings to a new string made up of more than eight characters before accessing the system.
- You cannot enable the web server at any time.
- You cannot enable the SSH password-authentication feature at any time.

Hsecure is disabled by default. When you enable hsecure, the desired behavior applies to all ports.

For more information, see [Prevent Certain Types of DOS Attacks](#) on page 3393.

Prioritization of Control Traffic

The switch uses a sophisticated prioritization scheme to schedule control packets on physical ports. This scheme involves two levels with both hardware and software queues to guarantee proper handling of control packets regardless of the switch load. In turn, this scheme guarantees the stability of the network. Prioritization also guarantees that applications that use many broadcasts are handled with lower priority.

You cannot view, configure, or modify control-traffic queues.

Directed Broadcast Suppression

You can enable or disable forwarding for directed broadcast traffic on an IP-interface basis. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcasts on an interface, you cause all frames sent to the subnet broadcast address for a local router interface to be dropped. Directed broadcast suppression protects hosts from possible DoS attacks.

To prevent the flooding of other networks with DoS attacks, such as the Smurf attack, enable directed broadcast suppression. For more information, see [Configure Directed Broadcast](#) on page 3391.

ARP Request Threshold

The Address Resolution Protocol (ARP) request threshold defines the maximum number of outstanding unresolved ARP requests. The default value for this function is 500 ARP requests. To avoid excessive amounts of subnet scanning that a virus can cause, as a best practice, change the ARP request threshold to a value between 100 and 50. This configuration protects the CPU from causing excessive ARP requests, protects the network, and lessens the spread of the virus to other PCs. The following list provides further ARP threshold values:

- Default: 500
- Severe conditions: 50
- Continuous scanning conditions: 100
- Moderate: 200
- Relaxed: 500

For more information about how to configure the ARP threshold, see [Address Resolution Protocol](#) on page 385.

Multicast Learning Limitation

The Multicast Learning Limitation feature protects the CPU from multicast data packet bursts generated by malicious applications. If more than a certain number of multicast streams enter the CPU through a port during a sampling interval, the port is shut down until the user or administrator takes the appropriate action.

For more information, see [IP Multicast](#) on page 1696.

Authentication for Privileged EXEC Command Mode

Table 291: Authentication for Privileged EXEC Mode product support

Feature	Product	Release introduced
Authentication for Privileged EXEC CLI Command Mode	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.4
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.2
	VSP 7400 Series	VOSS 8.2

For enhanced security, you can enable user authentication to enter Privileged EXEC command mode. Use the **sys priv-exec-password** command to enable password authentication.

After you enable password authentication for Privileged EXEC command mode, the system prompts you to enter a password to access Privileged EXEC command mode from User EXEC command mode. You must enter the same password that you used to log on to the switch.

Considerations

- The same user name and password used to Telnet or SSH to the switch is used to access Privileged EXEC command mode when authentication is enabled.
- RADIUS and TACACS+ protocols are supported. If RADIUS and TACACS+ servers are not reachable, access to Privileged EXEC command mode is denied. You must open a new session and type the same user name and password used to Telnet or SSH to the switch.
- No timeout exists for the password used to log on to the Privileged EXEC command mode.
- No limit to the number of retries if you enter an incorrect password.

Configuration Considerations

Use the information in this section to understand the limitations of some security functions, such as BSAC RADIUS servers and Layer 2 protocols before you attempt to configure security.

Attribute Format for a Third-party RADIUS Server

If you use a third-party RADIUS server and need to modify the dictionary files, you must add a vendor-specific attribute (attribute #26) and use 1584 as vendor code for all the

devices and then send back access-priority vendor-assigned attribute number 192 with a decimal value of 1 to 6, depending upon whether you want read only to read-write-all.

Authentication for Privileged EXEC Command Mode

Authentication for Privileged EXEC command mode supports RADIUS and TACACS+ protocols. If RADIUS and TACACS+ servers are not reachable, access to Privileged EXEC command mode is denied. You must open a new session and type the same user name and password used to Telnet or SSH to the switch.

RADIUS on Management Ports

The management port supports the RADIUS protocol. When RADIUS packets are sent out of the management port, the SRC-IP address is properly entered in the RADIUS header.

For more information about the supported RADIUS servers, see the documentation of the RADIUS server.

RADIUS Server SNMP Accounting

An SNMP query sent by an unreachable RADIUS server configured as used-by snmp and with accounting enabled, can cause a timeout. A timeout can occur if the device that receives the SNMP query attempts to send accounting packets to the unreachable server. You can mitigate the timeout issue by configuring lower retry and timeout values on the RADIUS server. Alternatively, you can configure a higher timeout value for SNMP.

Single Profile Enhancement for BSAC RADIUS Servers

Before enabling Remote Access Dial-In User Services (RADIUS) accounting on the device, you must configure at least one RADIUS server.

The switch software supports Microsoft Radius Servers (NPS Windows 2008, Windows 2003 IAS Server), BaySecure Access Control (BSAC), Merit Network servers and Linux based servers. To use these servers, you must first obtain the software for the server. You must also make changes to one or more configuration files for these servers.

Single Profile is a feature that is specific to BSAC RADIUS servers. In a BSAC RADIUS server, when you create a client profile, you can specify all the returnable attributes. When you use the same profile for different products, you specify all the returnable attributes in the single profile.

SNMP Cloned User Considerations

If the user from which you are cloning has authentication, you can choose for the new user to either have the same authentication protocol as the user from which it was cloned, or no authentication. If you choose authentication for the new user, you must provide a password for that user. If you want a new user to have authentication, you must indicate that at the time you create the new user. You can assign a privacy protocol only to a user that has authentication.

If the user from which you are cloning has no authentication, then the new user has no authentication.

Source IP Configuration

The Source IP is the Internet Protocol address of the device sending the IP data packet. The system limits the Source IP to a maximum of three interfaces:

- the management Out-of-Band (OOB) management interface
- the VLAN management interface
- the Circuitless IP (CLIP) management interface

The system uses separate routing tables for each Segmented Management Instance interface, plus a default main table. Because multiple routing tables are in use, each management interface can have overlapping or identical static routes without interfering with each other. The main table has a super-set of all routes where the weight of the static route can tie-break routes to the same destination going through different segmented management interfaces. By default, the following weights are used and the default route priority is management CLIP, then management VLAN, then management OOB:

- mgmt CLIP - 100
- mgmt VLAN - 200
- mgmt OOB - 300

You can route packets through a different management interface than the default configuration, but you must add a specific static route or change the default weight of the management interface.

For example, if the default route uses mgmt CLIP and you want to use the mgmt OOB interface as the source IP to reach a RADIUS server, you must perform one of the following options:

- Configure a specific source IP static route for the mgmt OOB interface:

```
mgmt oob
ip route 192.0.2.0/24 next-hop 198.51.100.1
```

OR

- Configure the default mgmt OOB route weight lower than the default mgmt CLIP route weight:

```
mgmt oob
no ip route 0.0.0.0/24 next-hop 198.51.100.1
ip route 0.0.0.0/24 next-hop 198.51.100.1 weight 50
```



Note

If you change the default route weight, the management interface with the lowest weight value becomes the default route for all segmented management interface traffic.

Interoperability Configuration

The switch is compatible with RADIUS servers.

Unicast Reverse Path Forwarding (uRPF)

Table 292: Unicast Reverse Path Forwarding product support

Feature	Product	Release introduced
Unicast Reverse Path Forwarding (URPF) checking (IPv4)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Supported on 5320-48P-8XE, 5320-48T-8XE, 5320-24T-24S-4XE-XT, and 5320-16P-2MXT-2X
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Unicast Reverse Path Forwarding (URPF) checking (IPv6)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Supported on 5320-48P-8XE, 5320-48T-8XE, and 5320-24T-24S-4XE-XT
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The Unicast Reverse Path Forwarding (uRPF) feature prevents packet forwarding for incoming unicast IP packets that have incorrect or forged (spoofed) IP addresses. The uRPF feature checks that the traffic received on an interface comes from a valid IP address, thereby preventing address spoofing. On a reverse path check, if the source IP address of the received packet at the interface is not reachable using the FIB, the system drops the packet as the packet may have originated from a misconfigured or a malicious source.

You can configure uRPF for each IP interface or VLAN. When uRPF is enabled on an interface, the switch checks all routing packets that come through that interface. It ensures that the system displays the source address and source interface in the routing table, and that it matches the interface, on which the packet was received.

You can use one of two modes for uRPF:

- **Strict mode:** In strict mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. If the incoming interface is not the best reverse path, the packet check fails and uRPF drops the packet. If the routing engine finds the source IP entry, uRPF further checks if the source IP interface matches the incoming interface of the packet. If they match, the system forwards the packet as usual, otherwise, the system discards the packet.

**Note**

The number of packets dropped due to uRPF check on the ingress interface gets incremented along with other general dropped statistics under the IN-DISCARD column in the output of the command `show interfaces gigabitEthernet error <collision|verbose> {slot/port[-slot/port]}[,...]`.

- **Loose mode:** In loose mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. The packet is dropped only if the source address is not reachable via any interface on that router.

uRPF can be enabled independently for IPv4 and IPv6. However, on a given interface, if uRPF is enabled for both IPv4 and IPv6, the uRPF mode can be either strict-mode or loose-mode for both IPv4 and IPv6. That means we cannot have IPv4 uRPF mode configured differently than that of IPv6.

**Note**

When you enable uRPF mode the MTU values for both IPv4 and IPv6 packets on the same VLAN are matched. Different Layer 3 MTU sizes on the same VLAN are not allowed in uRPF mode.

**Note**

uRPF check cannot detect spoofed source IP address if the source IP address belongs to a known subnet.

Digital Certificate/PKI

Table 293: Digital Certificate/PKI product support

Feature	Product	Release introduced
Digital Certificate/PKI	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Subject alternative name	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0.6
Certificate fingerprint validation	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	VOSS 8.3

Table 293: Digital Certificate/PKI product support (continued)

Feature	Product	Release introduced
Multiple CA Trustpoints and multiple Certificate Identities	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.3.100
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

This section provides information on the digital certificate framework and offline certificate management.

A digital certificate is an electronic document that identifies the subject, proves the ownership of a public key, and is digitally signed by a certificate authority (CA) that certifies the validity of the information in the certificate. A digital certificate is valid for a specific time period.

The switch uses Public Key Infrastructure (PKI) support to obtain and use digital certificates for secure communication in the network.

To be certified, a switch performs the following tasks:

- Generate a certificate signing request.
- Verify that a present certificate has not been revoked.
- Validate the certificate.
- Renew the certificate before it expires.
- Remove the certificate, if required.

Subject

An administrator configures the subject parameters, such as common name, organization name, organization unit, locality, state, country, and subject name for requesting the identity certificate.

Subject Name

You can configure up to 10 distinguished subject names.

Subject Alternative Name

A subject alternative name associates host name values, such as an e-mail address, an IP address, or a Fully qualified domain name (FQDN) with a security certificate. You can protect these additional host names with a single certificate.

Challenge Password

A password is required for Simple Certificate Enrollment Protocol (SCEP) operations, such as the enrollment and renewal of identity certificates. This password is given offline by the CA during end entity registration. The administrator provides this password during enroll and renew operations.

UsePost

There are different types of CAs such as EJBCA, Win2012, and others. The `usePost` parameter enables you to choose the style of HTTP request. The value for the `usePost` parameter can be **True** or **False**.

For example, if Win2012 SCEP does not support the POST mode of HTTP request, configure the `usePost` as `False` for Win2012 and configure `usePost` as `True` for EJBCA.

Root CA Certificate

The Root CA certificate obtained offline from a CA must be installed for SCEP operations. This Root CA certificate is transferred to the device during the certificate installation. SCEP operations cannot be performed if the offline Root CA certificate is not installed and if error messages are logged.

Key Generation

The supported key type is RSA with RSA key of size 2048. There can be only one active key-pair associated with the trustpoint CA and digital certificate. A new key-pair cannot be generated if there is a key-pair already associated with the active digital certificate. The system logs the error message if such new key generation is attempted. In such a case, the certificate must be revoked before a new key-pair is generated.

Trustpoint CA

Use trustpoints to manage and track CAs and certificates. A trustpoint is a representation of a CA or of an identity pair. A trustpoint contains the identity of the CA, CA-specific configuration parameters, and an association with one generated key. The switch can enroll with a trustpoint to obtain an identity certificate. Trustpoint is configured after the RSA key pair is generated and the CA identity and other configuration parameters are available. You can configure up to eight CA trustpoints by providing the CA name.

You can configure a SHA-256 fingerprint to authenticate a received CA certificate that matches the configured common name. The switch first checks for an installed, offline root certificate and validates against it. If no root certificate is present, the switch checks the SHA-256 fingerprint in the received CA certificate. The SHA-256 fingerprint does not authenticate the root certificate.

Certificate Enrollment

Certificate enrollment involves generating a certificate signing request (CSR). Before certificate enrollment, the trustpoint CA must be configured and the user configuration parameters should be available. The key usage extension parameter is required as an input; it indicates the purpose of the key contained in the certificate, that the key can be used for encipherment, digital signature, certificate signing and so on.

The certificate enrollment is not allowed if there is an active certificate already available. If new certificate enrollment is required, the existing active certificate must be revoked first. The system logs the enrollment success or failure responses.

Certificate Renewal

The administrator must renew the certificate before it expires. A trap is configured for a pre-defined period before the expiry date of the certificate, and the system logs the certificate renewal due warning message. A certificate renewal request is not performed if an active certificate is not available. The system replaces the existing certificate with the newly obtained certificate on successful renewal. The system logs the renewal success or failure responses.

Certificate Revocation or Removal

The certificate can be revoked or withdrawn from the specific device for a specific reason at any time. A certificate revocation request is not performed if an active certificate is not available. The system releases the existing certificate on successful revocation. The system logs the revocation success or failure responses.

During boot up, the system checks whether an active installed certificate is available. If a valid certificate is not available, the system logs the warning message.

Offline Certificate Management

Offline certificate management supports switches that cannot communicate with the Certificate Authority to obtain the identity certificate or certificates online by certificate enrollment operation.

The certificate signing request (CSR) is used to obtain the offline identity certificate. Configure the subject and RSA key-pair to obtain the offline identity certificate. You can generate and store up to 10 RSA keys identified by the key name label. To obtain multiple offline certificates, you must specify a distinguished subject-name and key-name.

You must install the Root CA certificate and all the intermediate CA certificates of the certificate chain in the device before installing the offline identity or device certificate. All the intermediate and Root CA certificates are stored in the certificate store and are used for CA certificate chain validation. The CA certificate chain validation is performed starting from the issuing CA certificate to the Root CA certificate during the installation of offline identity certificate. The offline identity certificate is installed only if the CA certificate chain validation, subject, and key match.

Storage

No digital certificate configuration is visible if you use the **show running-config** command. Instead, use the commands appropriate for displaying digital certificate information. For more information, see [View the Certificate Details](#) on page 3410.

Certificate Order Priority

Table 294: Certificate order priority product support

Feature	Product	Release introduced
Certificate order priority	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Use the following information to understand the certificate order priority when the Transport Layer Security (TLS) server and switch connect.

**Note**

For certain switches in enhanced secure mode, all sensitive files are protected. You cannot access any sensitive files using Telnet, SSH, FTP, SFTP, TFTP, and SCP connections. For more information, see [Sensitive File Protection](#) on page 270.

The TLS server selects a certificate authority (CA)-signed certificate if the certificate is already installed in the Digital Certificate module.

If the server certificates are not available, the TLS server generates a new self-signed certificate at startup and uses that by default. You can choose to use an online or an offline CA-signed certificate, which takes precedence over the self-signed certificate.

SAN Entries in Default TLS Certificates

The default TLS certificates add hostname.domainname and management IP address as subject alternative name (SAN) entries:

- If you configure both hostname and domain-name, the self-signed certificate uses `hostname.domain-name` for both common name (CN) and SAN DNS.
- If you configure domain-name but not hostname, the self-signed certificate uses `*.domain-name` for both CN and SAN DNS.
- If you configure hostname but not domain-name, the self-signed certificate uses `hostname.extremenetworks.com` for both CN and SAN DNS.
- If you do not configure either hostname or domain-name, the self-signed certificate uses `*.extremenetworks.com` for both CN and SAN DNS.
- All management IP addresses are added as SAN IP entries.

If you use default self-signed certificates generated in a release that did not include SAN entries in the certificate, generate the certificate again to avoid certificate errors in Web browsers that require these entries.



Note

As a best practice, use custom Public Key Infrastructure (PKI) certificates rather than the default self-signed certificates.

Two-Factor Authentication for SSH

Table 295: Two-Factor Authentication for SSH product support

Feature	Product	Release introduced
Two-Factor Authentication for SSH	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Use the following information to understand the Two-Factor Authentication feature.

Two-Factor Authentication uses smart card technology for remote device management. Two-Factor Authentication requires enhanced secure mode with Secure Shell (SSH) and X.509 V3 authorization enabled on the switch. You must provide the digital certificates to enable the identity management for the SSH client and SSH server. Two-Factor Authentication requires the following items:

- a Fabric Engine switch
- a computer with Secure CRT 8.3.2 or 8.3.3 as the SSH client
- a smart card reader
- a Common Access Card (CAC) or Personal Identity Verification (PIV) card for each configured user

You can also use a Windows Server 2008, or newer, configured with a Remote Access Dial-In User Services (RADIUS) server and Active Directory.

Digital certificates in the X.509 V3 format provide identity management. A chain of signatures by a trusted certificate authority (CA) and its intermediate certificate CAs binds a given public signing key to a given digital identity. For user authentication, the SSH client sends the user certificate stored on the CAC or PIV card to the SSH server

for verification. The SSH server validates the incoming user certificate using Public Key Infrastructure (PKI) trust-store.

After the switch validates the SSH certificate, the system parses for a user name to forward to the RADIUS server for authorization. The switch prompts you to enter a password for the user name. If the RADIUS server is unreachable or not configured, the authorization occurs locally on the switch for the user name and password.

Two-Factor Authentication on the switch uses SSH and the X.509 V3 certificates stored on the smart card. X.509 V3 digital certificates are documented in RFC5280.

Smart Card Authentication Process

The process for PIV or CAC card authentication is as follows:

1. The PIV Authentication or the Card Authentication certificate is read from the PIV Card Application.
2. The relying system validates the PIV Authentication certificate from the PIV Card Application using standards-compliant PKI path validation to ensure that the certificate is valid and from a trusted source.
3. The cardholder is prompted to submit a PIN to activate the card.
4. The relying system issues a challenge string to the card and requests an asymmetric operation in response.
5. The card responds to the previously issued challenge by signing using the PIV Authentication private key.
6. The relying system verifies that the response from the card is expected for the issued challenge.
7. A unique identifier from the PIV Authentication certificate is extracted and passed as input to the access control decision.

Security Configuration Using CLI

Configure security information used on the control and data paths to protect the network from uncontrolled access to network resources.

For more information about how to configure passwords and access policies, see [Managing the Switch](#) on page 246.

Enable hsecure

Enables High Secure mode. This mode provides the following password behavior:

- 10 character enforcement
- The password must contain a minimum of 2 uppercase characters, 2 lowercase characters, 2 numbers, and 2 special characters.
- Aging time
- Failed login attempt limitation

About This Task

If you enable the hsecure flag, you cannot enable the flags for the web server or SSH password-authentication.

When you upgrade from a previous release, if the password does not have at least 10 characters, you receive a prompt to change your password to the mandatory 10-character length.

If you enable hsecure for the first time and the password file does not exist, then the device creates a normal default user name (rwa) and password (rwa). In this case, the password does not meet the minimum requirements for hsecure and as a result the system prompts you to change the password.

The default value is disabled. If you enable High Secure mode, you must restart the switch to enforce secure passwords. If you operate the switch in High Secure mode, the switch prompts a password change if you enter invalid-length passwords.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable hsecure mode:

```
boot config flags hsecure
```

The system displays the following warning messages:

```
Warning: For security purposes, all unsecure services -
        TFTP, FTP, Telnet, SNMP and IQAgent are disabled.
        Individually enable the required services.
Warning: Please save configuration and reboot the switch
        for this to take effect.
```

3. Save the configuration and restart the device for the change to take effect.

Example

Enable hsecure mode. Save the configuration. Restart the switch.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#boot config flags hsecure
Warning: For security purposes, all unsecure services -
        TFTP, FTP, Telnet, SNMP and IQAgent are disabled.
        Individually enable the required services.
Warning: Please save configuration and reboot the switch
        for this to take effect.
Switch:1(config)#save config
Switch:1(config)#reset
Are you sure you want to reset the switch (y/n)?y
```

Change an Invalid-Length Password

Before You Begin



Important

When you enable hsecure, passwords must contain a minimum of 10 characters or numbers. The password must contain a minimum of: two uppercase characters, two lowercase characters, two numbers, and two special characters.

About This Task

After you enable hsecure mode and restart the system, change your password if you have an invalid-length password.

Procedure

1. At the CLI prompt, log on to the system.
2. Enter the password.

If you have an invalid-length password, the system displays the following message:

```
Your password is valid but less than mandatory 10 characters.  
Please change the password to continue.
```

3. When prompted, enter the new password.
4. When prompted, re-enter the new password.

Example

Log on to the switch:

```
Login: rwa
```

Enter the password:

```
Password: ***  
Your password is valid but less than mandatory 10 characters. Please change the password  
to continue.
```

Enter the new password:

```
Enter the new password: *****
```

Re-enter the new password:

```
Re-enter the new password: *****  
Password successfully changed.
```

Configure New Passwords

Configure new passwords for each access level, or change the logon or password for the different access levels of the switch. After you receive the switch, use default passwords to initially access CLI. If you use Simple Network Management Protocol version 3 (SNMPv3), you can change encrypted passwords.

Before You Begin

- You must use an account with read-write-all privileges to change passwords. For security, the switch saves passwords to a hidden file.

About This Task

If you enable the hsecure flag, after the aging time expires, the system prompts you to change your password. If you do not configure the aging time, the default is 90 days.

Procedure

- Enter Global Configuration mode:
`enable`

`configure terminal`
- Change a password:
`cli password WORD<1-20> {layer1|layer2|layer3|read-only|read-write|read-write-all}`
- Enter the old password.
- Enter the new password.
- Enter the new password a second time.
- Configure password options:
`password access-level WORD<2-8>`

`password aging-time day <1-365>`

`password default-lockout-retries <1-255>`

`password default-lockout-time <60-65000>`

`password lockout WORD<0-46> [time <60-65000>]`

`password min-passwd-len <10-20>`

`password password-history <3-32>`

Example

```
Switch:1>enable
Switch:1#configure terminal
```

Change a password:

```
Switch:1(config)# password smith read-write-all
```

Enter the old password:

```
Switch:1(config)#*****
```

Enter the new password:

```
Switch:1(config)#*****
```

Enter the new password a second time:

```
Switch:1(config)#*****
```

Set password to an access level of read-write-all and the expiration period for the password to 60 days:

```
Switch:1(config)#access-level rwa aging-time 60
```

Variable Definitions

The following table defines parameters for the **cli password** command.

Variable	Value
<i>layer1 layer2 layer3 read-only read-write read-write-all</i>	Changes the password for the specific access level.
<i>WORD<1-20></i>	Specifies the user logon name.

The following table defines parameters for the **password** command.

Variable	Value
<i>access level WORD<2-8></i>	Permits or blocks this access level. The available access level values are as follows: <ul style="list-style-type: none"> • l1 • l2 • l3 • ro • rw • rwa
<i>aging-time day <1-365></i>	Configures the expiration period for passwords in days, from 1–365. The default is 90 days.
<i>default-lockout-retries <1-255></i>	Configures the default number of login attempts. The default is 3.
<i>default-lockout-time <60-65000></i>	Changes the default lockout time after three invalid attempts. Configures the lockout time, in seconds, and is in the 60–65000 range. The default is 60 seconds.
<i>lockout WORD<0-46> time <60-65000></i>	Configures the host lockout time. <ul style="list-style-type: none"> • <i>WORD<0-46></i> is the host IP address in the format a.b.c.d. • <i><60-65000></i> is the lockout-out time, in seconds, in the 60–65000 range. The default is 60 seconds.
<i>min-passwd-len <10-20></i>	Configures the minimum length for passwords in high-secure mode. The default is 10 characters.
<i>password-history <3-32></i>	Specifies the number of previous passwords the switch stores. You cannot reuse a password that is stored in the password history. The default is 3.

Configure Password Hash

You can configure the Secure Hash Algorithm (SHA) level for password hash security. After a password hash level change, all custom users and passwords are deleted, and on first login each default user must change their default password.

Before You Begin

- You must use an account with read-write-all privileges to change password hash.

About This Task

Configure the password hash level as SHA1 or SHA2.

Procedure

- Enter Global Configuration mode:
`enable`
`configure terminal`
- Change the password hash level:
`password hash sha1|sha2`
- Read the warning prompt, then confirm the password hash change:
`y`

Example

Change password hash to SHA2 512-bit hash value:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#password hash sha2
Warning: Changing the password
        hash will reset all existing passwords and users to Factory defaults. All
custom Users
        will be deleted. Are you sure you want to proceed (y/n) ?y
```

Variable Definitions

The following table defines parameters for the **password hash** command.

Variable	Value
<code>sha1 sha2</code>	Specifies the password hash SHA value: <ul style="list-style-type: none">SHA1 - 160-bit password hashSHA2 - 512-bit password hash

Configure Directed Broadcast

About This Task

Configure directed broadcast on a VLAN to enable or disable directed broadcast traffic forwarding for an IP interface.

A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. When you disable (or suppress) directed broadcasts on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling directed broadcasts protects hosts from possible denial-of-service (DOS) attacks. By default, this feature is disabled on the device.



Note

Enabling this feature only applies to the primary IP address on the VLAN and not to Secondary IP Interfaces on the VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```
2. Configure the switch to forward directed broadcasts for a VLAN:

```
ip directed-broadcast enable
```
3. Verify the configuration:

```
show ip directed-broadcast vlan
```

Example

Enable directed broadcast on VLAN 2:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface vlan 2
Switch:1(config-if)#ip directed-broadcast enable
```

Display VLAN IDs with directed broadcast enabled:

```
Switch:1>show ip directed-broadcast vlan

                                Vlan Directed-Broadcast
-----
VLAN ID  DIRECTED-BROADCAST
-----
      2           true
```

Variable Definitions

The following table defines parameters for the **ip directed-broadcast** command.

Variable	Value
<i>enable</i>	Enables the device to forward directed broadcast frames to the specified VLAN. The default setting for this feature is disabled.

Prevent Certain Types of DOS Attacks

Protect the switch against IP packets with illegal IP addresses such as loopback addresses or a source IP address of ones, or Class D or Class E addresses from being routed. The switch supports high-secure configurable flag.

About This Task



Important

After you enable this flag, the desired behavior (not routing source packets with an IP address of 255.255.255.255) applies to all ports that belong to the same port.



Important

The setting to enable hsecure only takes effect for packets going to the CP; not to datapath traffic.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable high-secure mode:

```
high-secure [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] enable
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface GigabitEthernet 1/16
Switch:1(config-if)# high-secure enable
```

Variable Definitions

The following table defines parameters for the **high-secure** command.

Variable	Value
<code>port</code> <code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Specifies the port on which you want to enable high-secure mode.</p> <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>enable</code>	Enables the high-secure feature that blocks packets with illegal IP addresses. This flag is disabled by default.

Unicast Reverse Path Forwarding configuration using CLI

This section provides CLI procedures for Unicast Reverse Path Forwarding configuration.

Enable urpf-mode Boot Flag

To configure Unicast Reverse Path Forwarding on a port or VLAN, you are required to enable the urpf-mode boot flag. If you try to configure uRPF on an interface, that is, enable or change the urpf operating mode with the urpf-mode boot flag disabled, a consistency check error message is displayed: Unicast Reverse Path Forwarding configuration is not supported when urpf-mode boot flag is disabled.



Note

This procedure does not apply to 4220 Series.

About This Task

Use the following procedure to enable the urpf-mode boot flag. By default, urpf-mode is disabled.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the urpf-mode boot flag:
`boot config flags urpf-mode`

3. Enter `y` to reboot:

The new setting requires a reboot to take effect!

The configuration will be saved and rebooted.

Are you sure you want to re-boot the switch (y/n)?



Note

If you enter `n`, the following message displays: Warning: Please save the configuration and reboot the switch for this configuration to take effect.

4. Check the status of the urpf-mode boot flag:

```
show boot config flags
```

Example

Enable the urpf-mode boot flag:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# boot config flags urpf-mode
The new setting requires a reboot to take effect!
The configuration will be saved and rebooted.
Are you sure you want to re-boot the switch (y/n)? y
```

View the status of the urpf-boot flag:

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

Configure Unicast Reverse Path Forwarding on a Port

Before You Begin

- You must enable the urpf-mode boot flag. See [Enable urpf-mode Boot Flag](#) on page 3394.



Note

When you try to configure uRPF on an interface, that is, enable or change the urpf operating mode with the urpf-mode boot flag disabled, a consistency check error message is displayed: Unicast Reverse Path Forwarding configuration is not supported when urpf-mode boot flag is disabled.

- You must log on to the GigabitEthernet Interface Configuration mode in CLI.

About This Task

You can use the Unicast Reverse Path Forwarding (uRPF) feature to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. When you enable uRPF, the switch performs a check to determine if the source IP address of the packet is verifiable. If the address is not verifiable, the system drops the packet.

uRPF runs in two modes:

- strict mode
- loose mode (exist-only mode)

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Set or change the urpf operating mode on a port:

For IPv4, enter: `ip rvs-path-chk mode {strict|exist-only}`

For IPv6, enter: `ipv6 rvs-path-chk mode {strict|exist-only}`

3. To verify the configuration on the port, see [View the IPv4 and IPv6 Interface Configuration on a Port](#) on page 2125.

Example

Check whether the source IP address of the incoming packet exists in the FIB table.

Example for IPv4:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface GigabitEthernet 1/10
Switch:1(config-if)# ip rvs-path-chk mode strict
```

Example for IPv6:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface GigabitEthernet 1/16
Switch:1(config-if)# ipv6 rvs-path-chk mode strict
```

Variable Definitions

The following table defines parameters for the **ip rvs-path-chk mode** and **ipv6 rvs-path-chk mode** commands.

Variable	Value
mode{strict exist-only}	Specifies the mode for Unicast Reverse Path Forwarding (uRPF). In strict mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. If the incoming interface is not the best reverse path, the packet check fails and uRPF drops the packet. In exist-only mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. The packet is dropped only if the source address is not reachable via any interface on that router.

*Configure Unicast Reverse Path Forwarding on a VLAN***Before You Begin**

- You must enable the urpf-mode boot flag.

**Note**

When you try to configure uRPF on an interface, that is, enable or change the urpf operating mode with the urpf-mode boot flag disabled, a consistency check error message is displayed: Unicast Reverse Path Forwarding configuration is not supported when urpf-mode boot flag is disabled.

- You must log on to the VLAN Interface Configuration mode in CLI.

**Important**

You must assign a valid IP address to the selected port.

About This Task

Use the Unicast Reverse Path Forwarding (uRPF) feature to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses

into a network. When you enable uRPF, the switch performs a check to determine if the source IP address of the packet is verifiable. If the address is not verifiable, the system drops the packet.

uRPF runs in two modes:

- strict mode
- loose mode (exist-only mode)

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Set or change the urpf operating mode on a VLAN:

For IPv4, enter: `ip rvs-path-chk mode {strict|exist-only}`

For IPv6, enter: `ipv6 rvs-path-chk mode {strict|exist-only}`

3. To verify the configuration on the VLAN, see [View the IPv4 and IPv6 Interface Configuration on a VLAN](#) on page 4016.

Example

Example for IPv4:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface vlan 2
```

Check whether the source IP address of the incoming packet exists in the FIB table:

```
Switch:1(config-if)#ip rvs-path-chk mode exist-only
```

Example for IPv6:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface vlan 2
```

Check whether the source IP address of the incoming packet exists in the FIB table:

```
Switch:1(config-if)# ipv6 rvs-path-chk mode exist-only
```

Variable Definitions

The following table defines parameters for the **ip rvs-path-chk mode** and **ipv6 rvs-path-chk mode** commands.

Variable	Value
mode{strict exist-only}	Specifies the mode for Unicast Reverse Path Forwarding (uRPF). In strict mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. If the incoming interface is not the best reverse path, the packet check fails and uRPF drops the packet. In exist-only mode, uRPF checks whether the source IP address of the incoming packet exists in the FIB. The packet is dropped only if the source address is not reachable via an interface on that router.

Digital Certificate Configuration

The following section provides procedures to configure digital certificates using CLI.

For certificate configuration related to Syslog, see [Logs and Traps](#) on page 2535.

Configure Device Subject Parameters

About This Task

Use this procedure to configure the device subject parameters to identify the device, such as the name, Email ID, company, department, location, and subject name.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Configure the subject parameters of the device:

```
certificate subject {[common-name WORD<0-64>] [country WORD<2-2>] [e-mail WORD<0-254>] [locality WORD<0-128>] [organization WORD<0-64>] [province WORD<0-128>] [subject-name WORD<1-45>] [unit WORD<0-64>]}
```
3. Verify the subject names configured on the switch:

```
show certificate subject
```

Example

Configuring subject parameters:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# certificate subject common-name 822-pki e-mail
client1@extemenetworks.com unit Engineering
```

```
organization ExtremeNetworks locality Salem province Massachusetts country US subject-
name 822
```

View the configuration.

```
Switch:1>show certificate subject
Subject Name      : 822
Common Name       : 822-pki
Email Address     : client1@extemenetworks.com
Organizational Unit : Engineering
Organization      : ExtremeNetworks
Locality          : Salem
Province          : Massachusetts
Country           : US

Subject Name      : 823
Common Name       : 823-pki
Email Address     : client1@extemenetworks.com
Organizational Unit : Engineering
Organization      : ExtremeNetworks
Locality          : Salem
Province          : Massachusetts
Country           : US

Subject Name      : Global
Common Name       : 821
Email Address     : client1@extemenetworks.com
Organizational Unit : Engineering
Organization      : ExtremeNetworks
Locality          : Salem
Province          : Massachusetts
Country           : US
```

Variable Definitions

The following table defines parameters for the **certificate subject** command.

Variable	Value
<i>common-name</i> <i>WORD<0-64></i>	Specifies the name of the subject sending the Certificate Signing Request to the Certificate Authority.
<i>country</i> <i>WORD<2-2></i>	Specifies the country of the subject sending the Certificate Signing Request to the Certificate Authority.
<i>e-mail</i> <i>WORD<0-254></i>	Specifies the Email address of the subject sending the Certificate Signing Request to the Certificate Authority.
<i>locality</i> <i>WORD<0-128></i>	Specifies the locality of the subject sending the Certificate Signing Request to the Certificate Authority.
<i>organization</i> <i>WORD<0-64></i>	Specifies the organization of the subject sending the Certificate Signing Request to the Certificate Authority.
<i>province</i> <i>WORD<0-128></i>	Specifies the state or province of the subject sending the Certificate Signing Request to the Certificate Authority.

Variable	Value
<i>subject-name</i> <i>WORD<1-45></i>	Specifies the Subject Identity Label to be used in local digital certificate request. You can configure up to 10 subject DN identities. If a subject name is not specified, the default subject name is Global.
<i>unit WORD<0-64></i>	Specifies the organizational unit of the subject sending the Certificate Signing Request to the Certificate Authority.

Configure Subject Alternative Names

About This Task

Use this procedure to protect additional host names with a single certificate.



Note

The switch supports 100 subject alternative names.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a subject alternative name for the entity:

```
certificate subject-alternative-name {[dns WORD<1-255>] [e-mail  
WORD<1-255>] [ip> WORD<1-255>] [subject-name WORD<1-45>]}
```



Note



Note

You can configure up to 10 distinct subject names. The default subject name is Global.

3. View the subject alternative names configured on the switch:

```
show certificate subject-alternative-name
```

Examples

Configure a subject alternative name subject name 822:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#certificate subject-alternative-name subject-name 822 dns
822.extremenetworks.com
Switch:1(config)#certificate subject-alternative-name subject-name 822 e-mail
name@company.com
Switch:1(config)#certificate subject-alternative-name subject-name 822 ip 192.0.2.22
```

View the configuration:

```
Switch:1>show certificate subject-alternative-name
```

SAN Table		
TYPE	NAME	SUBJECT
E-MAIL	name@company.com	822
DNS	822.extremenetworks.com	822
IP	192.0.2.22	822

Configure a subject alternative name with the default subject name:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#certificate subject-alternative-name dns name@extrnetwk.com
Switch:1(config)#certificate subject-alternative-name e-mail j.smith@extremenetworks.com
Switch:1(config)#certificate subject-alternative-name ip 192.0.2.23
```

View the configuration:

```
Switch:1>show certificate subject-alternative-name
=====
SAN Table
=====
TYPE      NAME                                SUBJECT
-----
E-MAIL    j.smith@extremenetworks.com        Global
DNS       name@extrnetwk.com                 Global
IP        192.0.2.23                         Global
```

Variable Definitions

The following table defines parameters for the **certificate subject-alternative-name** command.

Variable	Value
<i>dns</i> WORD<1-255>	Specifies the DNS subject alternative name.
<i>e-mail</i> WORD<1-255>	Specifies the e-mail subject alternative name.
<i>ip</i> WORD<1-255>	Specifies the IP subject alternative name.
<i>subject-name</i> WORD<1-45>	Specifies the Subject Identity Label to be used in local digital certificate request. You can configure up to 10 subject DN identities. If the subject-name is not specified, the default subject name is Global.

Generate the Key Pair

About This Task

Use the following procedure to generate the private and public key pair for the specific cryptography type. By default, the switch generates a 2,048 RSA key when the system

starts. You can use this procedure to generate a new RSA key or to generate multiple RSA keys identified by a key-name. You can generate up to 10 RSA keys.

Procedure

1. Enter Global Configuration mode:

enable

```
configure terminal
```

2. Generate the key pair:

```
certificate generate-keypair {[type rsa size 2048] | [key-name WORD<1-45>1]}
```

Example

Generating the key pair identified by a key-name:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#certificate generate-keypair key-name pki_key
Switch:1(config)#1 2021-06-22T11:33:53.036Z 5520-24X-VOSS CP1 - 0x003a864c - 00000000
GlobalRouter DIGITALCERT INFO Digicert Key-pair generation is in progress
Switch:1(config):#1 2021-06-22T11:33:53.052Z 5520-24X-VOSS CP1 - 0x003a8604 - 00000000
GlobalRouter
DIGITALCERT INFO Generation of RSA key-pair for digital certificate is successful
1 2021-06-22T11:33:58.711Z Switch CP1 - 0x003a864d - 00000000 GlobalRouter DIGITALCERT
INFO Digicert Key-pair generation completed successfully for key-name: pki key
```

Display the configured key pairs:

```
Switch:1>show certificate key-name
Key Name: pki_key
Public Key Value:
00000000000000010000000102000000000301000100000100bdb1cf8382d66a2d2d0d24b4477908641c16423c
089d9131781a3ada005e
52074e1ff3561e29598f93c53dcb06e4d235335573419bb938b6ccf93d3e6767d0932e129ea2f556276efce2be
825df1f9dc661d3cafee7125f4f7126f5ba7e8
d9029623398b7d3fb00063ea0e4bedd56e276c52a6371b289de3ee4198ff2397b512b516604eac4e5f0f4a0621
d7ac42541491d368f21e17a440aa6130a825a2
a7ca6ab1d7a7868f93e4d0d83c7e4973cf204b4f5f654abbaa9aa6199247976488b0957e65b656a6d21a2a4ac4
d322a36c786d8a8deec763b6aec0d05b0f6bfe
87602caecb2cc71e2e4f9f4f8c4d4d4e9b25adf9c02eb44b763542f0449a326d0f3b

Key Name: rsa_2048
Public Key Value:
00000000000000010000000102000000000301000100000100c150b1851644aaef08060f3b3a7a0618758b841
84867ffd80b3e02ec306
76171fe36e99f5450656fc6e6db672b6239f760c97c3e49639cea5d503c0e478bf7a4d213d5698d09d63622ccb
279adbbaa34135c81d70660489b55b6babca59
4f17d8ed250cf917325df0f73a10896157e6e3a24a584bc713b2e6493d059c8efd53bbbf5db0aa95b43c1668ba
1053d0fe0e5c44dc889bd35bf11730e5827cb2
068048ab97e9f0757514f47332337376eed83a7cb95a53462639f5a47f026b0172cfa3ddffee7269e737a32d8f
2e5590a9ee07d3f329af4e4f2a73ed9de5991
6bc25e6ac51e482cbbb71f736ec0e396fc314e5eed3c438efff68d1a31bbed24d55
```

Variable Definition

The following table defines parameters for the **generate-keypair** command.

Variable	Value
<code>type rsa</code>	Specifies type of cryptography algorithm used to generate the key-pair. The switch uses only <code>rsa</code> as the cryptography algorithm type.
<code>size 2048</code>	Specifies the size or modulus of key-pair to be generated. The switch only supports 2048.
<code>key-name WORD<1-45></code>	Specifies the key label for RSA 2048 key to be generated. You can configure up to 10 RSA keys by specifying the key-name label. The default key-name label is <code>rsa_2048</code> .

Configure a Trustpoint CA

About This Task

Use this procedure to configure the certificate authority and perform related actions. You can configure up to eight CA trustpoints.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the trustpoint and associate it with the generated key pair:

```
certificate ca WORD<1-45> {[subject-name WORD<1-45>] [common-name WORD<0-64>] [key-name WORD<0-45>] [ca-url WORD<0-1000>] [use-post <true|false>]}
```

3. Configure an SHA-256 fingerprint to authenticate the received CA certificate:

```
certificate ca WORD<1-45> sha256-fingerprint WORD<64-64>
```

4. Configure the appropriate action:

- Configure trustpoint, authenticate the trustpoint CA by getting the certificate of the CA, and store the CA certificate locally:

```
certificate ca WORD<1-45> action caauth
```

- Generate certificate signing request to obtain identity certificate from configured trustpoint CA, get the digital certificate, and store it locally, associating with the trustpoint CA:

```
certificate ca WORD<1-45> action enroll [validity-days <7-1185>]
```

- Get the Certificate Revocation List from the CDP and store into a file:

```
certificate ca WORD<1-45> action get-crl
```

- Install the subject certificate obtained from the given trustpoint CA:

```
certificate ca WORD<1-45> action install
```

- Configure trustpoint and perform no other operation:

```
certificate ca WORD<1-45> action noop
```

- Release the locally stored certificate associated with the trustpoint CA post revocation:

```
certificate ca WORD<1-45> action remove
```

- Generate certificate renew request for given trustpoint CA, get the new digital certificate, and store it locally by replacing the old certificate with the new one:

```
certificate ca WORD<1-45> action renew [validity-days <7-1185>]
```

5. Install the Root Certificate Authority certificate obtained offline:

```
certificate ca WORD<1-45> install-file {root-ca-filename WORD<1-80>}
```

6. Set the HTTP request type to support the type of CA:

```
certificate ca WORD<1-45> use-post <false | true>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#certificate ca ej common-name CaA2 key-name pki_key subject-name 822
Switch:1(config)#certificate ca ej action enroll
Switch:1(config)#CP1 [07/21/16 12:22:11.992:CEST] 0x003a8604 00000000 GlobalRouter
DIGITALCERT
INFO Digital Certificate Module : Configuration Saved
CP1 [07/21/16 12:22:12.284:CEST] 0x003a8639 00000000 GlobalRouter DIGITALCERT INFO Sent
SCEP
Request To CA : ej
CP1 [07/21/16 12:22:12.504:CEST] 0x003a8615 00000000 GlobalRouter DIGITALCERT INFO
Received SCEP
Response With SUCCESS status!
CP1 [07/21/16 12:22:12.508:CEST] 0x003a8611 00000000 GlobalRouter DIGITALCERT INFO
Enroll
Certificate Successful!
CP1 [07/21/16 12:22:12.509:CEST] 0x003a8604 00000000 GlobalRouter DIGITALCERT INFO
Digital
Certificate Module : Configuration Saved
```

Display configured online CA trustpoints:

```
Switch:1(config)#show certificate ca

CA table entry
Name                : a1
CommonName          : CaA1
KeyName              : rsa_2048
SubjectName         :
CaUrl                : http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
UsePost              : 1
SubjectCertValidityDays : 365
Action               : (null)
LastActionStatus     : (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
bd9bb74b3f4d75e86113222a8d291b6349c7a42c457e487b9be0a48b4f09cc7c
UsedFor              :

CA table entry
Name                : a2
CommonName          : CaA2
KeyName              : pki_key
```

```

SubjectName      : 822
CaUrl            : http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
UsePost          : 1
SubjectCertValidityDays : 365
Action           : (null)
LastActionStatus : (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
0ccb8d0c38d36cf427187f0e1dd380536c078fd6fae39ec9872187327912056b
UsedFor          : Default

```

Variable Definitions

The following table defines parameters for the **certificate ca** command.

Variable	Value
<i>action caauth</i>	Authenticates the trustpoint CA by getting the certificate of the CA and stores the CA certificate locally.
<i>action enroll</i> <i>[validity-days <7-1185>]</i>	Generates certificate signing request to obtain identity certificate from configured trustpoint CA, gets the digital certificate, and stores it locally, associating with the trustpoint CA. The validity-days specifies the number of days for which the certificate will remain valid. The default value is 365 days.
<i>action get-crl</i>	Gets the Certificate Revocation List from the CDP and stores into a file.
<i>action install</i>	Installs the subject certificate obtained from the given trustpoint CA.
<i>action noop</i>	Specifies that no operation should be performed after configuring trustpoint.
<i>action remove</i>	Releases the locally stored certificate associated with the trustpoint CA post revocation.
<i>action renew</i> <i>[challenge-password</i> <i>WORD<0-128>]</i>	Specifies the password. This password is provided offline by the CA during the end entity registration.
<i>action renew</i> <i>[validity-days <7-1185>]</i>	Generates certificate renewal request for given trustpoint CA, gets the digital certificate, and stores it locally by replacing the old certificate with the new one. The validity-days specifies the number of days for which the certificate will remain valid. The default value is 365 days.
<i>ca WORD<1-45></i>	Specifies the name of the CA. You can configure up to 8 CA trustpoints by specifying the CA name. It should be alphanumeric and case-sensitive. The maximum length should be 45 characters.
<i>ca-url WORD<0-1000></i>	Specifies the trusted CA url.
<i>common-name</i> <i>WORD<0-64></i>	Specifies the name of the owner of the device or user.

Variable	Value
<code>key-name WORD<0-45></code>	Specifies the key pair generated by the command that was first associated with the CA trustpoint.
<code>install-file root-ca-filename WORD<1-80></code>	Installs the Root CA file obtained offline from the CA.
<code>sha256-fingerprint WORD<64-64></code>	Specifies an encrypted fingerprint of the expected certificate to match.
<code>subject-name WORD<1-45></code>	Specifies the configured Subject Identity label. The default is Global.
<code>use-post <false true></code>	Specify the HTTP request style. The default value is True. For example, True for EJBCA and False for Win2012 CA.

Install the Certificate

Before You Begin

To create a PKCS12 container with OpenSSL 3.0, use SHA-1. Certificate installation fails if you use SHA-256.

About This Task

Use this procedure to install the following:

- certificate authority (CA) certificate
- root CA certificates
- subject certificates
- Certificate Revocation List (CRL) file obtained offline from the CA

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Install the offline CA certificate:

```
certificate install-file offline-ca-filename WORD<1-80>
```
3. Install the CRL offline file:

```
certificate install-file offline-crl-filename WORD<1-80>
```
4. Install the root CA offline certificate:

```
certificate install-file offline-root-ca-filename WORD<1-80>
```

5. Install the subject offline certificate:

```
certificate install-file offline-subject-filename WORD<1-80> [relaxed]
[key-name WORD<1-45>] [subject-name WORD<1-45>]
```

**Note**

To obtain the offline subject certificate, you must first generate a certificate signing request (CSR).

6. (Optional) Install the subject offline certificate with PKCS12-format:

```
certificate install-file offline-subject-filename WORD<1-80> relaxed
pkcs12-password WORD<1-128>
```

Example

View the installed offline subject certificate:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1#certificate install-file offline-subject-filename 823pki.crt subject-name 823 key-
name pki
1 2021-02-02T14:19:01.587Z Switch CP1 - 0x003a864f - 00000000 GlobalRouter DIGITALCERT
INFO Performing OCSP Check For Certificate : 823-pki
1 2021-02-02T14:19:01.600Z Switch CP1 - 0x003a8603 - 00000000 GlobalRouter DIGITALCERT
INFO Subject Certificate obtained offline from CA successfully installed
1 2021-02-02T14:19:01.622Z Switch CP1 - 0x003a8604 - 00000000 GlobalRouter DIGITALCERT
INFO Digital Certificate Module : Configuration Saved
1 2021-02-02T14:19:01.666Z Switch CP1 - 0x003a8619 - 00000000 GlobalRouter DIGITALCERT
INFO Received OCSP Response with SUCCESS Status!
```

The following output displays the CA name derived from the subject name and the key name. You use this entry when you configure a specific application to use a specific CA identity.

```
#show certificate ca

CA table entry
Name           : 823-pki[auto-installed]
CommonName     : CaA2-1
KeyName        : pki
SubjectName    : 823
CaUrl          :
UsePost        : 0
SubjectCertValidityDays : 0
Action         : (null)
LastActionStatus : (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
UsedFor        :
```

Variable Definitions

The following table defines parameters for the **certificate install-file** command.

Variable	Definition
<code>offline-ca-filename WORD<1-80></code>	Specifies the certificate authority (CA) file name obtained from the CA.
<code>offline-crl-filename WORD<1-80></code>	Specifies the CRL file obtained from the CA.
<code>offline-root-ca-filename WORD<1-80></code>	Specifies the root CA file name obtained from the CA.
<code>offline-subject-filename WORD<1-80></code>	Specifies the subject certificate file name obtained from the CA.
<code>relaxed [pkcs12-password WORD<1-128>]</code>	Uses the relaxed mode for offline subject certificate installation for less restrictive consistency checks. You can also install a PKCS12 format certificate and secret key in relaxed mode. <code>WORD<1-128></code> is the password to extract the PKCS12 container. If you do not include this parameter, the supported format is Distinguished Encoding Rules (DER).
<code>key-name WORD<1-45></code>	Refers to the key name of the generated key-pair.
<code>subject-name WORD<1-45></code>	Refers to the subject identity name.

Generate the Certificate Signing Request

About This Task

Use this procedure to generate a certificate signing request (CSR) and store it into a file. This CSR is required to obtain the offline subject certificate. The generated file is located in the `/intflash/shared/certs` folder.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Generate certificate signing request:

```
certificate generate-csr [relaxed]
```
3. (Optional) Configure the subject-name and the key name:

```
certificate generate-csr subject-name WORD<1-45> key-name WORD<1-45>
```

Example

Display the generated CSR.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#certificate generate-crs subject-name 823 key-name pki
```

```
Switch:1(config)#1 2021-02-02T13:57:39.716z Switch CP1 - 0x003a8635 - 0000000
GlobalRouter DIGITALCERTIFICATE INFO Generate CSR for Digital Certificate successful!
```

Display Configured Key Pairs

About This Task

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Display the names and public keys of all the key-pairs:

```
show certificate key-name [WORD<1-64>]
```

Example

The following output example displays the names and public keys of all the key-pairs.

```
Switch:1(config)#show certificate key-name

Key Name: pki_key
Public Key Value:
0000000000000001020000000003010001000001009a8d66a1efaeac7bd2a4918ffdda59d8bebfd682
b59
4df64ce86df60cac8e81d93d88104e53dfa0e6446fdb91e34b83c3da6d4225ee0615e80e2e3e95c2ec8cac46db
a653f2365c6a11c892091
cb1a86746149100bdc2520e009c400c3fdb9d150735452ec08fb6b56efa23af0414126808b1ea1044740200b4
9074bafd5b10c9a48ee8b
5f79168a9403f0d33bb1997cf72040ddc43bfefc3343fe75757c6eb6a2cdbea9ddf2b91cac7bf91da9b3f55961
7587a94ef2cf0dc0d38ff
d9ad7a50991da234c6be8bcfec98f43b41e56704e969e794de9418f672ca6ec0e0451a8d318e704022c723ad7b
1171f72203a320b940340
88379128ccb35d2c5e217b19

Key Name: rsa_2048
Public Key Value:
0000000000000001020000000003010001000001009e5ffca4e1a22a33cb78ca489dec74e2c71cf5ba
e54
c97f6e656edf0fc6c547feeb7a6c66bd85b3ee2dceecadc92e4ab92be172b532a992d73fcfe606d292e9251689
3f591f651cd60443b3829
cb67f656c1d3ab65d46d16ff0a72c18710c84781b477c0830ab36dc304a2c8dcab93bcbf1c724cbb4f24ed7e8d
d76944d9c4289dd4a6af8
3c14bd266345ee20c24d610f630f66a74f355367235ade20288d85437a0f72990520f2ffab8ec700a13298e108
b4eddcdeedd7eceeef8e6
ea1fdb80db7caa931daff6ab643e93b2a9787c80db8c464946e85a7b8bd01d331ba1dd5302e3c8ec1f0dc56ea8
45b1f07afd0b476e5fe1
2c36a5c159182b769e800f44b
```

View the Certificate Details

About This Task

Use this procedure for the following tasks:

- Displaying the digital certificate for given certificate type or list all the certificate details from the local store for given certificate type.
- Displaying the CA details for a given trustpoint CA name or listing all the CA details from the local store if the CA name is not specified.


```

ExtendedKeyUsage      : TLS Web Client Authentication, OCSP Signing, TLS Web
Server Authentication,
CDPUrl                : http://192.51.100.9:8080/ejbca/publicweb/webdist/
certdist?cmd=crl&issuer=CN=rootCa2

OCSPUrl               : http://192.51.100.9:8080/ejbca/publicweb/status/ocsp
CertificateFileName    : root_ca_cert_dod.mil.der

```

Display the Intermediate CA certificate:

```

Switch:1>show certificate cert-type intermediate-ca-cert
CERT STORE table entry
Certificate Type      : Intermediate CA Certificate
CommonName           : subCa1
VersionNumber        : X.509 v3
SerialNumber         : 427c233041d7ef98
IssuerName           : CN:rootCa1, EM:, OU:, O:, L:, P:, C:
ValidityPeriodNotBefore : 03/23/2020 17:25:39
ValidityPeriodNotAfter  : 03/23/2023 17:25:39
CertificateSignatureAlgorithm : sha256withRSAEncryption
CertificateSignature   :
8f64d17ed47b467db2754e132cc9f1c9aed24aa9d279e769606caf2390cc21f23cb3d2ef943ed9577f82012589
3210b
ba2bd3d021c900bab94de8c24b6ba2538cd5843592bd4cefb966cb70ed36a73db0b38a297a85a04d8d021f37f4
983e77673b736cf86b37ec49447d44abbb71e6b91
b6b192939de0689d7e5e830753ab2b32bfb22f35dc571256a562c07a783b50f486a46855435c3b51f5ede16e86
7ff9b3de426e5fd103446a4c7bde5fca82ef273fe
8bdc88d66b5429366a1ad81c7c3bc9bc5b9adf5e60fd9559b679ccc97484f75406aa3669dc98982e74f5e2f382
be0700307e57fae1fc9f7a87abc76bf9a75e7678c
5cbea2cabdf80682adeca198
Subject              : CN:subCa1, EM:, OU:, O:, L:, P:, C:
SubjectPublicKeyAlgorithm : rsaEncryption
SubjectPublicKey      :
00000000000000200000000100000000300d06092a864886f70d0101010500000000000000010d020000000003
010001
00000100c12a373907e30965449febe3a25588ffebeld7cb452b75e14f43f27f5b1e2d0aa371123886c0bf6854
11d0b4bfac8472cde94c463a3ff15797236038d854
fde48efb4dfb9dd1ea1a16d5a9675c00d6df8e09a5200b3b6d93bfacbd6d56bc485bcd66f9b8a06fd10ba601108
ae6a4dd0e9ff32ffc5f69faca56ff49db781a5a13e
20ff616d918b81d0da97c823b6240903d18e641b0d35b510784e831a4fe44201162682ea1e82e0215b0ae88d7f
f876809e40ba69260da633242fb30e02e23eaf8147
7fb451a9474e3070b8e424b796af710a2b08a82ffaa90ac21bcf0759c8662e536bc170ad1fde3e913c7e28b2a7
06d40825ce3dbe5efe89958663c7e99e4b
HasBasicConstraint    : 1
HasKeyUsage           : 1
IsCa                  : 1
KeyUsage              : 117 digitalSignature keyEncipherment keyAgreement
keyCertSign cRLSign
ExtendedKeyUsage      : TLS Web Client Authentication, OCSP Signing, TLS Web
Server Authentication,
CDPUrl                : http://192.51.100.9:8080/ejbca/publicweb/webdist/
certdist?cmd=crl&issuer=CN=rootCa1

OCSPUrl               : http://192.51.100.9:8080/ejbca/publicweb/status/ocsp
CertificateFileName    : ca_cert_subCa1.der

```

Display the offline CA certificate:

```

Switch:1>show certificate cert-type offline-ca-cert
CERT table entry
Certificate Type      : Offline CA Certificate
VersionNumber        : X.509 v3
SerialNumber         : 427c233041d7ef98
IssuerName           : CN:rootCa1, EM:, OU:, O:, L:, P:, C:

```



```

3479437a69731287ab2a84091fe2b17625689
cc189fd513d240d42b1251c7e044740461c940ec49b64343d911dd89c3c24942b3dfc5585941833cfd157b3957
fc4f36304dde7d236b710eef09a97ab5bf184
355acfeddbe197ca93a4350568d3b6ac1edd4bafdlff223e6aca16bc3887582bd97101ca87bb678d48bcb6f0c
fbd217271a8c1b04c01baaf8aeb9803bbb625
e73693ca7477fc6916ae0ebe48f411f67f9785324fd9d3ab673bf1a35c5f75d3b05443254918b24aea55e94f14
ac3190ee24a22164e7696c778e5034e6ba7d5
2de100b0a8e65459b
HasBasicConstraint          : 1
HasKeyUsage                 : 1
IsCa                        : 1
KeyUsage                    : 117 digitalSignature keyEncipherment keyAgreement
keyCertSign cRLSign
ExtendedKeyUsage            : TLS Web Client Authentication, OCSP Signing, TLS Web
Server Authentication,
CDPUrl                      : http://192.51.100.9:8080/ejbca/publicweb/webdist/
certdist?cmd=crl&issuer=CN=rootCa2
OCSPUrl                     : http://192.51.100.9:8080/ejbca/publicweb/status/ocsp

```

Display the offline subject certificate:

```

Switch:1>show certificate cert-type offline-subject-cert
CERT table entry
Certificate Type             : Offline Subject Certificate
VersionNumber               : X.509 v3
SerialNumber                : 5de44b25394462b8
IssuerName                  : CN:subCa1, EM:, OU:, O:, L:, P:, C:
ValidityPeriodNotBefore    : 07/05/2021 12:24:45
ValidityPeriodNotAfter     : 07/05/2022 12:24:45
CertificateSignatureAlgorithm : sha256withRSAEncryption
CertificateSignature         :
34f5037b30b0332e15f504316be86afcc41ad0b93699bc8de1b5cbe97a8cc834593837032ab492e0c5eee9a1fe
8db
99e8ea7aeb41fdce86818e0c08b1ed9e79a43247383e88fd3ef504a28b1ee525be60cba78291be16f57fb54174
33ec9dce601c9b4e77986c5db9430ce6cece48b
3dc143d042614404bdc3c2df16f68bb1b0609e593636a2806b285cb8fa7e470b442b50e4d3c4a663ac99d5d3b4
29a9b4966ea5ce16da6b7d7c5607832cc6acaea
e578419ba52e11cbe30d2cbb53a05de58e374657fc5983a92c699ba6896160c9f32e6625bd6f71003259773e71
d7c89df3ddc0a8603c1a8c8f6e248002f2bd217
1a6e922abf2e8134b311d1897319bbb7
Subject                     : CN:s15, EM:demo, OU:demo, O:demo, L:demo, P:demo, C:RO
SubjectPublicKeyAlgorithm   : rsaEncryption
SubjectPublicKey            :
0000000000000000000020000000100000000300d06092a864886f70d0101010500000000000000010d020000000003
010
00100000100c150b1851644aaef08060f3b3a7a0618758b84184867ffd80b3e02ec30676171fe36e99f545065
6fc6e6db672b6239f760c97c3e49639cea5d503
c0e478bf7a4d213d5698d09d63622ccb279adbbaa34135c81d70660489b55b6babca594f17d8ed250cf917325d
f0f73a10896157e6e3a24a584bc713b2e6493d0
59c8efd53bbbf5db0aa95b43c1668ba1053d0fe0e5c44dc889bd35bf11730e5827cb2068048ab97e9f0757514f
47332337376eed83a7cb95a53462639f5a47f02
6b0172cfa3ddffee7269e737a32d8f2e5590a9ee07d3f329af4e4f2a73ed9de599916bc25e6ac51e482cbbb71f
736ec0e396fc314e5eed3c438efff68d1a31bdb
ed24d55
HasBasicConstraint          : 1
HasKeyUsage                 : 1
IsCa                        : 0
KeyUsage                    : 15 digitalSignature nonRepudiation keyEncipherment
dataEncipherment
ExtendedKeyUsage            : TLS Web Client Authentication, OCSP Signing, TLS Web
Server Authentication,
CDPUrl                      : http://192.51.100.9:8080/ejbca/publicweb/webdist/
certdist?cmd=crl&issuer=CN=subCa1
OCSPUrl                     : http://192.51.100.9:8080/ejbca/publicweb/status/ocsp

```

```

Revocation Status      : unknown
Status                 : offline-certificate
Installed              : 1
CertificateFileName     : self_cert_s15.der

```

Display the online subject certificate:

```

Switch:1>show certificate cert-type online-subject-cert

CERT table entry
Certificate Type      : Online Subject Certificate
VersionNumber        : X.509 v3
SerialNumber         : 22db2464ad9d616a
IssuerName           : CN:subCa2, EM:, OU:, O:, L:, P:, C:
ValidityPeriodNotBefore : 07/05/2021 12:35:22
ValidityPeriodNotAfter  : 07/05/2022 12:35:22
CertificateSignatureAlgorithm : sha256withRSAEncryption
CertificateSignature   :
42309f2e2072abb2fc3214da4b69cb8767375e2b899a01a84af126bca9b65abb3460b901997eb6dfc7687c6cf
611ffcd4bacee80a8be24886be1bbd2a18559cc3f01509c733729c6c22bc83310592d37eb545bd9ae209e8772
2234c12000acb0983119f9420816ee96a6e
52b028f6e8daa4641e192912f1c4972459d2de8e0f39759f0ca5f841e51460a24922eeb0ada602985666fa154f
6c7b088682e247df38f2161f2a57f4e281f
c8c56ea0e6afff6e249c326b62316a42df0a03250b85b9be1c8d3228cd387e63daba5650a2a3d6c1586fc91925
57aadbc60e71611b696b43957fa95295ffc
632d7d889e58086356834bd27e4e2b92ac81b1543c7e6993
Subject              : CN:clientr1.example.dod.mil, EM:demo@demo.com,
OU:demo, O:ExtremeNetworks, L:Salem, P:demo, C:US
SubjectPublicKeyAlgorithm : rsaEncryption
SubjectPublicKey      :
000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
301000100000100bdb1cf8382d66a2d2d0d24b4477908641c16423c089d9131781a3ada005e52074e1ff3561e2
9598f93c53dcb06e4d235335573419bb938
b6ccf93d3e6767d0932e129ea2f556276efce2be825df1f9dc661d3cafee7125f4f7126f5ba7e8d9029623398b
7d3fb00063ea0e4bedd56e276c52a6371b2
89de3ee4198ff2397b512b516604eac4e5f0f4a0621d7ac42541491d368f21e17a440aa6130a825a2a7ca6ab1d
7a7868f93e4d0d83c7e4973cf204b4f5f65
4abbaa9aa6199247976488b0957e65b656a6d21a2a4ac4d322a36c786d8a8deec763b6aec0d05b0f6bfe87602c
aeb2cc71e2e4f9f4f8c4d4d4e9b25adf9c
02eb44b763542f0449a326d0f3b
HasBasicConstraint    : 1
HasKeyUsage           : 1
IsCa                  : 0
KeyUsage              : 15 digitalSignature nonRepudiation keyEncipherment
dataEncipherment
ExtendedKeyUsage       : TLS Web Client Authentication, OCSP Signing, TLS Web
Server Authentication,
CDPUrl                 : http://192.51.100.9:8080/ejbca/publicweb/webdist/
certdist?cmd=crl&issuer=CN=subCa2
OCSPUrl                : http://192.51.100.9:8080/ejbca/publicweb/status/ocsp

Revocation Status     : active
Status                : active
Installed              : 1

```

Display the CA certificates details:

```

Switch:1(config)#show certificate ca

CA table entry
Name              : a1

```

```

CommonName      : CaA1
KeyName         : rsa_2048
SubjectName     :
CaUrl           : http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
UsePost         : 1
SubjectCertValidityDays : 365
Action          : (null)
LastActionStatus : (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
bd9bb74b3f4d75e86113222a8d291b6349c7a42c457e487b9be0a48b4f09cc7c
UsedFor         :

CA table entry
Name            : a2
CommonName     : CaA2
KeyName        : pki_key
SubjectName    : 822
CaUrl          : http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
UsePost        : 1
SubjectCertValidityDays : 365
Action         : (null)
LastActionStatus : (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
0ccb8d0c38d36cf427187f0e1dd380536c078fd6fae39ec9872187327912056b
UsedFor        : Default

```

Display offline subject certificate details:

```

#show certificate ca

CA table entry
Name            : 823-pki[auto-installed]
CommonName     : CaA2-1
KeyName        : pki
SubjectName    : 823
CaUrl          :
UsePost        : 0
SubjectCertValidityDays : 0
Action         : (null)
LastActionStatus : (null)
LastActionFailureReason :
CA-Auth Sha256Fingerprint :
UsedFor        :

```

Display the name and public key of all the key-pairs:

```

Switch:1>show certificate key-name
Key Name: pki_key
Public Key Value:
000000000000001000000000301000100000100bdb1cf8382d66a2d2d0d24b4477908641c16423c
089d9131781a3ada005e
52074e1ff3561e29598f93c53dcb06e4d235335573419bb938b6ccf93d3e6767d0932e129ea2f556276efce2be
825df1f9dc661d3cafee7125f4f7126f5ba7e8
d9029623398b7d3fb00063ea0e4bedd56e276c52a6371b289de3ee4198ff2397b512b516604eac4e5f0f4a0621
d7ac42541491d368f21e17a440aa6130a825a2
a7ca6ab1d7a7868f93e4d0d83c7e4973cf204b4f5f654abbaa9aa6199247976488b0957e65b656a6d21a2a4ac4
d322a36c786d8a8deec763b6aec0d05b0f6bfe
87602caecb2cc71e2e4f9f4f8c4d4d4e9b25adf9c02eb44b763542f0449a326d0f3b

```

```
Key Name: rsa_2048
Public Key Value:
0000000000000010000000102000000000301000100000100c150b1851644aaef08060f3b3a7a0618758b841
84867ffd80b3e02ec306
76171fe36e99f5450656fc6e6db672b6239f760c97c3e49639cea5d503c0e478bf7a4d213d5698d09d63622ccb
279adbbaa34135c81d70660489b55b6babca59
4f17d8ed250cf917325df0f73a10896157e6e3a24a584bc713b2e6493d059c8efd53bbbf5db0aa95b43c1668ba
1053d0fe0e5c44dc889bd35bf11730e5827cb2
068048ab97e9f0757514f47332337376eed83a7cb95a53462639f5a47f026b0172cfa3ddffee7269e737a32d8f
2e5590a9ee07d3f329af4e4f2a73ed9de59991
6bc25e6ac51e482cbbb71f736ec0e396fc314e5eed3c438efff68d1a31bdbed24d55
```

Display the details of the configured subject:

```
Switch:1>show certificate subject
Subject Name      : client1
Common Name       : client1.example.dod.mil
Email Address     : client1@extremenetworks.com
Organizational Unit : Engineering
Organization      : ExtremeNetworks
Locality          : Salem
Province          : Massachusetts
Country           : US

Subject Name      : Global
Common Name       : s15
Email Address     : enduser15@extremenetworks.com
Organizational Unit : Engineering
Organization      : ExtremeNetworks
Locality          : Salem
Province          : Massachusetts
Country           : US
```

The following example displays the self-signed certificate:

```
Switch:1>show certificate cert-type default-tls-certificate

Certificate Type           : Self-signed certificate
VersionNumber             : X.509 v3
SerialNumber              : 2feefda3f54ac70e3d0f1d2a7aa1b4eb23865820
IssuerName                : CN:*extremenetworks.com, EM:, OU:VOSS, O:Extreme
Networks, L:San Jose, P:CA, C:US
ValidityPeriodNotBefore   : 09/16/2021 10:32:29
ValidityPeriodNotAfter    : 09/15/2022 10:32:29
CertificateSignatureAlgorithm : sha256withRSAEncryption
CertificateSignature       :
814924d27cc95af2274e88d0d5174aced2c2eb721053ab8877c001c0433a38737085d3925cae6be
85f605f941814c16cf4f9fc122934d11726e852b1fca95012966f1672505fbfca2f451bab2c9c65ac8ecdffbf05
fa295d1076be55d7aef55be41
a9d44688dfcaa30e1c47dc3462c0dde59678d06fac4546175af12942aaee611b816a547e4de5618a96d29f75e5
be931c2fe0af1a003481d18b2
2fcabb39d54721e3acb2bc66ef19172b03067a4700911818fb028c680217a10bee365f1c415b9adefa46c4643e
bf884786154b6ec3f295fe956
f7fcf360e69308a5bcc6db444bec7a17154f1d1fb6e2e069696fd74692c51362d89ac8b77bc7701f2cf21505
Subject                   : CN:*extremenetworks.com, EM:, OU:VOSS, O:Extreme
Networks, L:San Jose, P:CA, C:US
SubjectPublicKeyAlgorithm : rsaEncryption
SubjectPublicKey          :
000000000000000020000000100000000300d06092a864886f70d010101050000000000000010d02
000000000301000100000100bfa70f688c6f0ceb8cb882940d73b2b1cee00767e46d6173ff265cd05bc2757438
6c302f08cc58f5dffbac40ca5a
1ac3e8c54c3114553759de2b5cf82ea265ec150ef32b4e8022c9db6427d70a34ef300f65f4f0d87cba3dbe20a2
248d7c5b785ef20286a48a6366
f494e669c9e414c629888586e4a6e8079ab0a21c535313026c002ad8e46dbdb55cd3bfc1c8a9738f05f666b9f7
```

```
d1a324a890e7798122f6058d6a
f4d22c6dc26274e7ab21e992667dba5e73c4f4a074e37d036482343c38e4e1aa6697f9644e2d6d40561459cb25
9ce8a55b7686d059b1116cabe4
a526c41284ce7369c6711b71ab1b31856119a844d9f6374c9e9e593b42aa4f903ee6f199
SubjectAlternativeNames      :   DNS:*.extremenetworks.com, IP:192.0.2.02,
IP:198.51.100.01
```

Variable Definitions

The following table defines parameters for the **show certificate** command.

Variable	Value
<code>cert-type default-tls-certificate</code>	Displays the default TLS certificate (self-signed).
<code>cert-type online-ca-cert</code>	Specifies Certificate Authority's Certificate obtained online from Certificate Authority.
<code>cert-type online-subject-cert</code>	Specifies subject certificate obtained online from Certificate Authority.
<code>cert-type offline-ca-cert</code>	Specifies Certificate Authority's certificate obtained offline from Certificate Authority.
<code>cert-type offline-subject-cert</code>	Specifies subject certificate obtained offline from Certificate Authority.
<code>cert-type intermediate-ca-cert [WORD<1-80>]</code>	Specifies the intermediate certificate obtained offline from Certificate Authority.
<code>cert-type root-ca-cert [WORD<1-80>]</code>	Specifies root certificate obtained offline from Root Certificate Authority.
<code>ca [WORD<1-45>]</code>	Specifies name of the Certificate Authority. If the name is not specified, the command displays the CA details of all configured CA.

Configure a Web Server Certificate

About This Task

Use this task to configure a web server certificate.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the Certificate Authority for the installed digital certificate:

```
web-server certificate ca-name <1-45>
```
3. Configure the subject identity label to be used

```
web-server certificate cert-subject-name <1-45>
```

Example

Display the web server certificate information:

```
Switch:1(config)#web-server certificate cert-subject-name 823
Switch:1(config)#1 2021--2--2T14:35:38.865z Switch CP1 - 0x003z8604 - 0000000
```



```
GlobalRouter DIGITALCERT INFO Digital Certificate Module: Configuration Saved
1 2021--2--2T14:35:38.865z Switch CP1 Switch CP1 - 0x00040601 - 0000000 Global Router
WEB INFO HTTPS: Using the CA Signed Server Cert
```

Variable Definitions

Use the data in the following table to use the **web-server certificate** command

Variable	Value
<i>ca-name</i> WORD<1-45>	Specifies the name of the Certificate Authority.
<i>ca-subject-label</i> WORD<1-45>	Specifies the name of the local certificate used.

Digital Certificate Configuration Examples

This section shows how to obtain an online CA signed certificate, remove the expired certificate, renew the certificate, and install an offline subject certificate.

Obtain an Online CA-signed Subject Certificate

Use the following procedure as an example to obtain an online CA signed subject certificate that the application can use.

About This Task

In the following commands, the variable *WORD<1-45>* refers to the name of the certificate authority and the variable *WORD<1-80>* refers to the certificate file name.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the subject:

```
certificate subject common-name 822-pki
certificate subject e-mail 822@extremenetworks.com
certificate subject unit Engineering
certificate subject organization ExtremeNetworks
certificate subject locality Salem
certificate subject country US
certificate subject province Massachusetts
certificate subject subject-name 822

certificate subject common-name 823-pki
certificate subject e-mail 823@extremenetworks.com
certificate subject unit Engineering
certificate subject organization ExtremeNetworks
certificate subject locality Salem
certificate subject country US
```

```
certificate subject province Massachusetts
certificate subject subject-name 823
```

**Note**

The values mentioned are for example only.

3. Generate the key pair:

```
certificate generate-keypair {[type rsa size 2048] | [key-name WORD<1-45>]}
```

4. Configure the certificate authority (CA):

```
certificate ca ej common-name subca5
certificate ca ej key-name rsa_2048
certificate ca ej ca-url http://192.51.100.9:8080/ejbca/publicweb/apply/scep/test/
pkiclient.exe
certificate ca ej use-post true
```

**Note**

The values mentioned are for example only.

5. Use SCP to upload the Root CA certificate to: /intflash/shared/certs.
6. Install the Root CA certificate:

```
certificate ca WORD<1-45> install-file root-ca-filename WORD<1-80>
```

7. Authenticate the CA:

```
certificate ca WORD<1-45> action caauth
```

8. Enroll the subject certificate by the CA:

```
certificate ca WORD<1-45> action enroll
```

9. Install the certificate:

```
certificate ca WORD<1-45> action install
```

10. (Optional) If the certificate expires, remove the enrolled subject certificate:

```
certificate ca WORD<1-45> action remove
```

11. (Optional) To obtain the new certificate before the old certificate expires, enter the following command to renew the certificate:

```
certificate ca WORD<1-45> action renew
```

The Certificate Authority generates a new certificate for the subject.

Install an Offline CA Certificate

Use the following procedure as an example to install an offline CA certificate.

About This Task

In the following commands, the variable *WORD<1-80>* refers to the certificate file name.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the subject:

```
certificate subject common-name 822-pki
certificate subject e-mail client1@extemenetworks.com
certificate subject unit Engineering
certificate subject organization ExtremeNetworks
certificate subject locality Salem
certificate subject province Massachusetts
certificate subject country US
certificate subject subject-name 822
```



Note

The values mentioned are for example only.

3. Generate the certificate signing request to support multiple subject identities on the switch:

```
certificate generate-csr subject-name WORD< 1-45> key-name WORD<1-64>
```

```
#certificate generate-csr subject-name 823 key-name mimi
Switch:1(config)#1 2021-02-02T13:57:39.716Z SwitchModel CP1 - 0x003a8635 - 00000000
GlobalRouter DIGITALCERT
INFO Generate CSR For Digital Certificate successful!
```

4. Use the generated CSR file to enroll the certificate on the server.

5. Use SCP to upload the enrolled certificate along with Root certificate and all intermediary certificates to:

```
/intflash/shared/certs/
```

6. Install the Root CA certificate:

```
certificate install-file offline-root-ca-filename WORD<1-80>
```



Note

If the subject certificate issuer is directly the Root, then Step 7 and 8 are optional. If the subject is issued by Intermediate CA, then Step 7 and 8 are mandatory, also in the certificate chain between Root and Subject, all the Intermediates must be installed using these steps.

7. Copy and paste the Intermediate CA certificate to:

```
/intflash/shared/certs/
```

8. Install the intermediate CA:

```
certificate install-file offline-ca-filename WORD<1-80>
```

9. Install the offline subject file name:

```
certificate install-file offline-subject-filename WORD<1-80>
```

```
#certificate install-file offline-subject-filename sd

Error: File Name Not Found in /intflash/shared/certs/ or /intflash/.cert/.offlineCert/
#certificate install-file offline-subject-filename 823mimi.crt subject-name 823 key-
name mimi
1 2021-02-02T14:19:01.587Z SwitchModel CP1 - 0x003a864f - 00000000 GlobalRouter
DIGITALCERT INFO
Performing OCSP Check For Certificate : 823-mimi
1 2021-02-02T14:19:01.600Z SwitchModel CP1 - 0x003a8603 - 00000000 GlobalRouter
DIGITALCERT INFO
Subject Certificate obtained offline from CA successfully installed
1 2021-02-02T14:19:01.622Z SwitchModel CP1 - 0x003a8604 - 00000000 GlobalRouter
```

```
DIGITALCERT INFO
Digital Certificate Module : Configuration Saved
1 2021-02-02T14:19:01.666Z SwitchModel CP1 - 0x003a8619 - 00000000 GlobalRouter
DIGITALCERT INFO
Received OCSP Response with SUCCESS Status!
```

Configure X.509 V3 Certificates for SSH Two Factor Authentication

Use the following procedure as an example to configure the SSH server on the switch, and the SSH client Secure CRT for two factor authentication using X.509 V3 certificates.

Before You Begin

The following certificates must be loaded on the SSH server and SSH client:

- For the Secure CRT (SSH client):
 - subject certificate from the PIV card.
- For the switch (SSH server):
 - CAC-server.pem - the subject certificate
 - ca.cert.pem - the root CA certificate
 - Self-signedTrustAnchorCertificate.cer - the root CA certificate that signed the intermediate certificate
 - RSA2048IssuingCACertificate.cer - the intermediate certificate signed by the previous root CA that signed the subject certificate.

About This Task

Use the following steps as an example to configure the SSH server on the switch, the RADIUS Windows server, and the SSH client Secure CRT.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Ensure the switch is running in Enhanced Secure Mode:

```
Switch:1(config)#show boot config flags
flags enhancedsecure-mode true
```



Note

This shows a partial output of only the relevant entry.

3. Ensure the switch clock is synchronized:

```
Switch:1#show clock
System Clock time : Fri Oct 12 19:36:36 2018 UTC
```

4. Provision PKI with certificates.

For information about provisioning PKI with certificates, see the following sections:

- [Obtain an Online CA-signed Subject Certificate](#) on page 3419
- [Install an Offline CA Certificate](#) on page 3420

X.509 Authentication User Name Option Example

Use the following procedure as an example to configure user name authentication options using X.509 V3 certificates.

Procedure

1. Enable X.509 V3 authentication user name override:

```
Switch:1(config)# ssh x509v3-auth username overwrite
```

The switch disregards the user name sent by the SSH client and uses the principal name from the client's certificate for authentication. If RADIUS authentication is configured, the user name is sent after you type the RADIUS password. For example, if you configure the SSH client with the user name "John" and enable x509v3-auth user name overwrite on the device, the switch sends the principal name 1403824387@mil to the RADIUS server for authorization.

```
Test CAC John Smith
Issuer: CN=DOD JITC ID CA-49, OU=PKI, OU=DoD, O=U.S. Government, C=US
NotBefore: 12/19/2017 7:00 PM
NotAfter: 12/19/2020 6:59 PM
Subject: CN=SMITH.JOHN.1403824387, OU=USAF, OU=PKI, OU=DoD, O=U.S.
Government, C=US
Serial: 019ab3
SubjectAltName: Other Name:Principal Name=1403824387@mil, Other
Name:2.16.840.1.101.3.6.6=04 19 d4 f8 10 da 08 26 6c 10 e7 22 e5 83 68 5a 04 0e 44 82
64 5c 85 78 10 93 ee
Cert: 8066aec3484d3740d7d99ec2f5ed1983365bb1
```

2. Enable X.509 V3 authentication user name strip:

```
Switch:1(config)#ssh x509v3-auth username strip-domain
```

If x509v3-auth user name strip-domain is configured, the switch sends the principal name without the domain to the RADIUS for authorization. The user name is sent after you type the RADIUS password. For example: If you select principal name 1403824387@mil, the switch sends the principal name 1403824387 without the domain to the RADIUS server for authorization.

```
Test CAC John Smith
Issuer: CN=DOD JITC ID CA-49, OU=PKI, OU=DoD, O=U.S. Government, C=US
NotBefore: 12/19/2017 7:00 PM
NotAfter: 12/19/2020 6:59
PM Subject: CN=SMITH.JOHN.1403824387, OU=USAF, OU=PKI, OU=DoD, O=U.S.
Government, C=US
Serial: 019ab3
SubjectAltName: Other Name:Principal Name=1403824387@mil, Other
Name:2.16.840.1.101.3.6.6=04 19 d4 f8 10 da 08 26 6c 10 e7 22 e5 83 68 5a 04 0e 44 82
64 5c 85 78 10 93 ee
Cert: 8066aec3484d3740d7d99ec2f5ed1983365bb129
```

3. Enable X.509 V3 authentication user name use-domain:

```
Switch:1(config)ssh x509v3-auth username use-domain extreme.com
```

If you select the user name as the principal name, the switch sends the principal name from the certificate with the domain configured on the switch to the RADIUS for authorization. The user name is sent after typing the RADIUS password. For example: If you configure use-domain "extreme.com" on the switch and you configure the user name to be the principal name 1403824387@mil, the

switch sends the user name 1403824387@extreme.com to the RADIUS server for authorization.

```
Test CAC John Smith
Issuer: CN=DOD JTC ID CA-49, OU=PKI, OU=DoD, O=U.S. Government, C=US
NotBefore: 12/19/2017 7:00 PM
NotAfter: 12/19/2020 6:59 PM
Subject: CN=SMITH.JOHN.1403824387, OU=USAF, OU=PKI, OU=DoD, O=U.S.
Government, C=US
Serial: 019ab3
SubjectAltName: Other Name:Principal Name=1403824387@mil, Other
Name:2.16.840.1.101.3.6.6=04 19 d4 f8 10 da 08 26 6c 10 e7 22 e5 83 68 5a 04 0e 44 82
64 5c 85 78 10 93 ee
Cert: 8066aec3484d3740d7d99ec2f5ed1983365bb129
```

4. RADIUS server is not configured:

If the RADIUS server is not configured, the authorization falls back locally on the switch, for the user name. You must configure the user names on the switch. You are prompted for the password. For example: If you select the principal name 1403824387@mil, the switch authorizes locally the user name as 1403824387@mil or 1403824387 if strip domain is enabled. You are prompted for the password.

```
Test CAC John Smith
Issuer: CN=DOD JTC ID CA-49, OU=PKI, OU=DoD, O=U.S. Government, C=US
NotBefore: 12/19/2017 7:00 PM
NotAfter: 12/19/2020 6:59 PM
Subject: CN=SMITH.JOHN.1403824387, OU=USAF, OU=PKI, OU=DoD, O=U.S.
Government, C=US
Serial: 019ab3
SubjectAltName: Other Name:Principal Name=1403824387@mil, Other
Name:2.16.840.1.101.3.6.6=04 19 d4 f8 10 da 08 26 6c 10 e7 22 e5 83 68 5a 04 0e 44 82
64 5c 85 78 10 93 ee
Cert: 8066aec3484d3740d7d99ec2f5ed1983365bb129
```

Configure TCP Keepalive and TCP Timestamp

About This Task

TCP Keepalive configures the system TCP keepalive interval, probes, and time.

TCP Timestamp option (RFC 1323) allows TCP to determine the order in which the packets are sent. The TCP Timestamp provides protection against Wrapped Sequence numbers. However, it is possible to calculate the system uptime when the Timestamp option is enabled. The analysis of timestamp behaviour can provide information on the system identity, which poses security threats and can cause a potential attack.

The TCP Timestamp option is enabled by default. You can disable the timestamp to avoid any security risks.



Note

The configuration will be applied only to the new TCP connections and the existing connections are not affected. You must perform a **config save** and **reboot** to apply the new configuration to all TCP connections.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the TCP keepalive interval:
`sys control tcp-keepalive-interval <1-3600>`
3. Configure the TCP keepalive probes:
`sys control tcp-keepalive-probes <1-50>`
4. Configure the TCP keepalive time:
`sys control tcp-keepalive-time <5-65535>`
5. Enable the TCP Timestamp:
`sys control tcp-timestamp`
6. Disable the TCP Timestamp:
`no sys control tcp-timestamp`
7. View the status of TCP Timestamp:
`show sys control`

Examples

Configure the TCP keepalive to interval 60 seconds, with 15 probes, and 120 seconds time:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#sys control tcp-keepalive-interval 60
Switch:1(config)#sys control tcp-keepalive-probes 15
Switch:1(config)#sys control tcp-keepalive-time 120
```

Disable TCP Timestamp:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#no sys control tcp-timestamp
```

Warning: Existing TCP connections won't be affected.

A config save and reboot is required to apply this configuration for all TCP connections

Display the status of the TCP Timestamp:

```
Switch:1>show sys control
=====
                        System Control Settings
=====
      tcp-timestamp : enable
      tcp-keepalive-time : 60
tcp-keepalive-interval : 10
      tcp-keepalive-probes : 5
      tcp-mtu-probing : disabled
      tcp-base-mss : 1024
      mac-move-protection : on
```

Variable Definitions

The following table defines parameters for the **sys control** command.

Variable	Value
<i>tcp-keepalive-interval</i> <1-3600>	Configure the TCP keepalive interval in seconds. The default is 10.
<i>tcp-keepalive-probes</i> <1-50>	Configure the TCP keepalive probes. The default is 5.
<i>tcp-keepalive-time</i> <5-65535>	Configure the TCP keepalive time in seconds. The default is 60.
<i>tcp-timestamp</i>	<p>Enable or disable tcp-timestamp.</p> <p>Note: The timestamp is enabled by default. The system displays the following warning message when a new configuration is applied: Warning: Existing TCP connections won't be affected. A config save and reboot is required to apply this configuration for all TCP connections.</p>

Enable Authentication for Privileged EXEC Command Mode

With authentication enabled, you must enter a user name and password to access Privileged EXEC command mode from User-EXEC command mode. The user name and password is the same user name and password you used to Telnet or SSH to the switch.



Note

When you enable authentication for the Privileged EXEC CLI command mode, the changes do not affect existing CLI sessions. For the changes to take effect, you must first log out from your current CLI session and log back in on a new session.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable authentication for Privileged EXEC CLI command mode:

```
sys priv-exec-password
```
3. Verify the configuration:

```
show sys priv-exec-password
```


Enable MAC Move Protection on vIST

About This Task

Perform this procedure to enable MAC move protection on virtual interswitch trunk (vIST).

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable MAC move protection:

```
sys control virtual-ist mac-move-protection
```
3. Confirm the configuration:

```
show sys control
```

Examples

Enable MAC move protection on vIST:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#sys control virtual-ist mac-move-protection
Switch:1>show sys control
=====
                        System Control Settings
=====
                        tcp-timestamp : enable
                        tcp-keepalive-time : 60
                        tcp-keepalive-interval : 10
                        tcp-keepalive-probes : 5
                        tcp-mtu-probing : disabled
                        tcp-base-mss : 1024
                        mac-move-protection : on
```

Variable Definitions

The following table defines parameters for the **sys control virtual-ist** command.

Variable	Value
<i>mac-move-protection</i>	Enable MAC move protection on vIST. The default value is disabled.

Security Configuration using EDM

Configure security information used on the control and data paths to protect the network from uncontrolled access to network resources.

For more information about how to configure passwords and access policies, see [System Access Configuration Using EDM](#) on page 309.

Change Passwords

About This Task

Configure new passwords for each access level, or change the logon or password for the different access levels of the system to prevent unauthorized access. After you receive the switch, use default passwords to initially access CLI. If you use Simple Network Management Protocol version 3 (SNMPv3), you can change passwords in encrypted format.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **CLI** tab.
4. Specify the user name and password for the appropriate access level.
5. Select **Apply**.

CLI Field Descriptions

The following table defines parameters for the **CLI** tab.

Name	Description
RWAUserName	Specifies the user name for the read-write-all CLI account.
RWAPassword	Specifies the password for the read-write-all CLI account.
RWEnable	Activates the read-write access. The default is enabled.
RWUserName	Specifies the user name for the read-write CLI account.
RWPassword	Specifies the password for the read-write CLI account.
RWL3Enable	Activates the read-write Layer 3 access. The default is enabled.
RWL3UserName	Specifies the user name for the Layer 3 read-write CLI account.
RWL3Password	Specifies the password for the Layer 3 read-write CLI account.
RWL2Enable	Activates the read-write Layer 2 access. The default is enabled.
RWL2UserName	Specifies the user name for the Layer 2 read-write CLI account.
RWL2Password	Specifies the password for the Layer 2 read-write CLI account.
RWL1Enable	Activates the read-write Layer 1 access. The default is enabled.

Name	Description
RWLUserName	Specifies the user name for the Layer 1 read-write CLI account.
RWLPassword	Specifies the password for the Layer 1 read-write CLI account.
ROEnable	Activates the read-only CLI account. The default is enabled.
ROUserName	Specifies the user name for the read-only CLI account.
ROPassword	Specifies the password for the read-only CLI account.
MaxTelnetSessions	Specifies the maximum number of concurrent Telnet sessions in a range from 0–8. The default is 8.
Timeout	Specifies the number of seconds of inactivity for a Telnet before the system initiates automatic timeout and disconnect, expressed in a range from 30–65535. The default is 900 seconds.
NumAccessViolations	Indicates the number of CLI access violations detected by the system. This variable is a read-only field.
CustomBannerText	Specifies the text message that is displayed to users on the CLI before authentication. The message can be company information, such as company name and contact, or a warning message for the users of CLI. With character limitation from 1-1800, the text box displays 79 characters per line.
CustomBannerEnable	Specifies whether custom logon banner is enabled or disabled. The default is enabled.

Configure Directed Broadcast on a VLAN

About This Task

Configure directed broadcast on a VLAN to enable or disable directed broadcast traffic forwarding for an IP interface.

A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. When you disable (or suppress) directed broadcasts on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling directed broadcasts protects hosts from possible denial-of-service (DOS) attacks. By default, this feature is disabled on the device.



Note

Enabling this feature only applies to the primary IP address on the VLAN and not to Secondary IP Interfaces on the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **Direct Broadcast** tab.
7. Select **DirectBroadcastEnable**.



Important

Configure multiple VLANs or IPs in the same subnet but in different systems simultaneously.

8. Select **Apply**.

Direct Broadcast Field Descriptions

Use the data in the following table to use the **Direct Broadcast** tab.

Name	Description
DirectBroadcastEnable	<p>Specifies that an Isolated Routing Port (IRP) can forward directed broadcast traffic. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcast on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling this function protects a host from possible denial of service (DoS) attacks.</p> <p>With the feature enabled, the Control Processor (CP) does not receive a copy of the directed broadcast. As a result, the system does not respond to a subnet broadcast ping sent from a remote subnet.</p> <p>The default is disabled.</p>

Configure TCP Keepalive and TCP Timestamp

About This Task

TCP Keepalive configures the system TCP keepalive interval, probes, and time.

TCP Timestamp option (RFC 1323) allows TCP to determine the order in which the packets are sent. The TCP Timestamp provides protection against Wrapped Sequence numbers. However, it is possible to calculate the system uptime when the Timestamp option is enabled. The analysis of timestamp behaviour can provide information on the system identity, which poses security threats and can cause a potential attack.

The TCP Timestamp option is enabled by default. You can disable the timestamp to avoid any security risks.

**Note**

The configuration will be applied only to the new TCP connections and the existing connections are not affected. You must save the configuration and reboot the switch to apply the new configuration to all TCP connections.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System Control** tab.
4. Configure the TCP timestamp by performing one of the following actions:
 - a. Select **TcpTimestamp** to enable the TCP timestamp.
 - b. Clear **TcpTimestamp** to disable the TCP timestamp.
5. In **KeepaliveTime**, enter a number in seconds.
6. In **KeepaliveInterval**, enter a number in seconds.
7. In **KeepaliveProbes**, enter a number of probes.
8. Select **Apply**.

System Control *Field Descriptions*

Use the data in the following table to use the **System Control** tab.

Name	Description
TcpTimestampEnable	Enables or disables the TCP timestamp. The timestamp is enabled by default. The system displays the following warning message after you apply a new configuration: Warning: Existing TCP connections won't be affected. A config save and reboot is required to apply this configuration for all TCP connections.
KeepaliveTime	Specify the TCP keepalive time in seconds. The default is 60.
KeepaliveInterval	Specify the TCP keepalive interval in seconds. The default is 10.
KeepaliveProbes	Specify the TCP keepalive probes. The default is 5.

Name	Description
PrivExecPasswordEnable	Enables authentication to access Privileged EXEC CLI command mode. Authentication is disabled by default.
RateLimitActionPollInterval	Specifies the interval, in seconds, at which this feature checks the hardware counters for packets dropped because they exceeded the configured port rate limit. The default is 3 seconds.

Enable Authentication for Privileged EXEC Command Mode

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **System Control** tab.
4. Select **PrivExecPasswordEnable**.

Unicast Reverse Path Forwarding configuration using EDM

This section provides EDM procedures for Unicast Reverse Path Forwarding configuration.

Configure Reverse Path Checking on a Port

Before You Begin

- The system supports reverse path checking only on ports that have a valid IP address.

About This Task

Configure reverse path checking on a port to determine if a packet IP address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.

4. Select the **Reverse Path Checking** tab.
5. Select **Enable** to enable reverse path checking.
6. Select **exist-only** or **strict**.
7. Select **Apply**.

Reverse Path Checking Field Descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
Enable	Enables reverse path checking on the selected port. The default is disabled.
Mode	<p>Specifies the mode for reverse path checking. The modes are</p> <ul style="list-style-type: none">• exist-only—reverse path checking checks whether the incoming packet source IP address exists in the routing table. If reverse path checking finds the source IP entry, the packet is forwarded; otherwise the packet is discarded.• strict—reverse path checking checks whether the incoming packet source IP address exists in routing table. If reverse path checking does not find the source IP entry, the packet is dropped; otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, the packet is forwarded; otherwise the packet is discarded. <p>The default is exist-only.</p>

Configure Reverse Path Checking on an IPv6 Port

Before You Begin

- The system supports reverse path checking only on ports that have a valid IP address.

About This Task

Configure reverse path checking on a port to determine if a packet IP address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Procedure

1. In the Device Physical View tab, select a port.

2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IPv6**.
4. Select the **Reverse Path Checking** tab.
5. Select **ReversePathCheckEnable** to enable reverse path checking.
6. Select **exist-only** or **strict**.
7. Select **Apply**.

Reverse Path Checking Field Descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
ReversePathCheckEnable	Enables reverse path checking on the selected port. The default is disabled.
ReversePathCheckMode	<p>Specifies the mode for reverse path checking. The modes are</p> <ul style="list-style-type: none">• exist-only—reverse path checking checks whether the incoming packet source IP address exists in the routing table. If reverse path checking finds the source IP entry, the packet is forwarded; otherwise the packet is discarded.• strict—reverse path checking checks whether the incoming packet source IP address exists in routing table. If reverse path checking does not find the source IP entry, the packet is dropped; otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, the packet is forwarded; otherwise the packet is discarded. <p>The default is exist-only.</p>

Configure Reverse Path Checking on a VLAN

Before You Begin

- Before you can configure reverse path checking on a VLAN, you must assign a valid IP address to the selected VLAN.

About This Task

Configure reverse path checking on a VLAN to determine if a packet IP address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN on which you want to configure reverse path checking.
5. Select **IP**.
6. Select the **Reverse Path Checking** tab.
7. Select **Enable** to enable reverse path checking.
8. Select **exist-only** or **strict**.
9. Select **Apply**.

Reverse Path Checking Field Descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
Enable	Enables reverse path checking on the selected VLAN.
Mode	<p>Specifies the mode for reverse path checking. The modes are</p> <ul style="list-style-type: none">• exist-only—reverse path checking checks whether the incoming packet source IP address exists in the routing table. If reverse path checking finds the source IP entry, the packet is forwarded; otherwise, the packet is discarded.• strict—reverse path checking checks whether the incoming packet source IP address exists in routing table. If reverse path checking does not find the source IP entry, then the packet is dropped. Otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, then the packet is forwarded. Otherwise, the packet is discarded. <p>The default is exist-only.</p>

Configure Reverse Path Checking on an IPv6 VLAN

Before You Begin

- Before you can configure reverse path checking on a VLAN, you must assign a valid IP address to the selected VLAN.

About This Task

Configure reverse path checking on a VLAN to determine if a packet IP address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select the VLAN on which you want to configure reverse path checking.
5. Select **IPv6**.
6. Select the **Reverse Path Checking** tab.
7. Select **ReversePathCheckEnable** to enable reverse path checking.
8. Select **exist-only** or **strict**.
9. Select **Apply**.

Reverse Path Checking Field Descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
ReversePathCheckEnable	Enables reverse path checking on the selected VLAN.
ReversePathCheckMode	<p>Specifies the mode for reverse path checking. The modes are</p> <ul style="list-style-type: none"> • exist-only—reverse path checking checks whether the incoming packet source IP address exists in the routing table. If reverse path checking finds the source IP entry, the packet is forwarded; otherwise, the packet is discarded. • strict—reverse path checking checks whether the incoming packet source IP address exists in routing table. If reverse path checking does not find the source IP entry, then the packet is dropped. Otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, then the packet is forwarded. Otherwise, the packet is discarded. <p>The default is exist-only.</p>

Digital Certificate Configuration Using EDM

The following section provides procedures to configure digital certificates using EDM.

*Configure Device Subject Parameters***About This Task**

Use this procedure to configure the device subject parameters to identify the device. The parameters include name, Email ID, company, department, and location of the subject.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Certificate**.
3. Select the **Global Subject** tab.
4. In **CommonName**, type the name of the subject.
5. Complete the remaining optional configuration to customize the policy.
6. Select **Apply**.

Subject Field Descriptions

Use the data in the following table to use the **Subject** tab.

Name	Description
CommonName	Specifies the name of the subject sending the Certificate Signing Request to the Certificate Authority.
EmailAddress	Specifies the Email address of the subject sending the Certificate Signing Request to the Certificate Authority.
OrganizationalUnit	Specifies the organizational unit of the subject sending the Certificate Signing Request to the Certificate Authority.
Organization	Specifies the organization of the subject sending the Certificate Signing Request to the Certificate Authority.
Locality	Specifies the locality of the subject sending the Certificate Signing Request to the Certificate Authority.
Province	Specifies the province of the subject sending the Certificate Signing Request to the Certificate Authority.
Country	Specifies the country of the subject sending the Certificate Signing Request to the Certificate Authority.
InstallFile	Installs the specific certificate file type obtained offline from the Certificate Authority.
InstallFileName	Specifies the certificate file name to install.
UninstallFile	Uninstalls the specific certificate file type obtained offline from the Certificate Authority.

Name	Description
UninstallFileName	Specifies the certificate file name to uninstall.
GenerateCsr	Generates the certificate signing request to obtain the offline subject certificate.

Configure Device Subject Entries Parameters

Use this procedure to configure the device subject parameters to identify the device. The parameters include name, Email ID, company, department, and location of the subject.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Certificate**.
3. Select the **Subject Entries** tab.
4. Select **Insert**.
5. In **Name**, type the name of the subject.
6. Complete the remaining optional configuration to customize the policy.
7. Select **Insert**.

Subject Entries Field Descriptions

Use the data in the following table to use the **Subject Entries** tab.

Name	Description
Name	Specifies the user defined name referring to the subject.
CommonName	Specifies the name of the subject sending the Certificate Signing Request to the Certificate Authority.
EmailAddress	Specifies the Email address of the subject sending the Certificate Signing Request to the Certificate Authority.
OrganizationalUnit	Specifies the organizational unit of the subject sending the Certificate Signing Request to the Certificate Authority.
Organization	Specifies the organization of the subject sending the Certificate Signing Request to the Certificate Authority.
Locality	Specifies the locality of the subject sending the Certificate Signing Request to the Certificate Authority.
Province	Specifies the province of the subject sending the Certificate Signing Request to the Certificate Authority.

Name	Description
Country	Specifies the country of the subject sending the Certificate Signing Request to the Certificate Authority.
InstallFile	Installs the specific certificate file type obtained offline from the Certificate Authority.
InstallFileName	Specifies the certificate file name to install.
UninstallFile	Uninstalls the specific certificate file type obtained offline from the Certificate Authority.
UninstallFileName	Specifies the certificate file name to uninstall.
GenerateCsr	Generates the certificate signing request to obtain the offline subject certificate.
RelaxedMode	Specifies relaxed mode for offline subject certificate installation for less restrictive consistency checks.
Pkcs1Password	Specifies the password to extract the PKCS12 container when in relaxed mode.
KeyName	Specifies the key name of the generated key-pair.

Generate the Key Pair

About This Task

Use the following procedure to generate the private and public key pair for the specific cryptography type.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Certificate**.
3. Select the **Key-Pair** tab.
4. Select **Insert**.
5. In **Type**, select the cryptography type.
6. In **Size**, type the size of the key.
7. Select **Insert**.

Key-Pair Field Description

Use the data in the following table to use the **Key-Pair** tab.

Name	Description
Type	Specifies the cryptography algorithm used to generate the key-pair. Only RSA is supported.
Size	Specifies the size of the key-pair to be generated. Only 2048 is supported.
Name	Specifies the name of the key-pair generated for the subject. This name is auto-generated as the combination of key-type and key-size.

Configure the Certificate Authority

Use this procedure to configure the certificate authority (CA) and perform related actions. You can configure only one CA in a device at a time.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Certificate**.
3. Select the **CA** tab.
4. Select **Insert**.
5. In **Name**, type a user-defined name of the CA.
6. In **CommonName**, type the common name of the CA.
7. In **KeyName**, type the name of the associated key pair.
8. Complete the remaining optional configuration to customize the policy.
9. Select **Insert**.
10. (Optional) Select **Retry Action** if the trustpoint CA certificate authentication fails or takes time for authentication. This can be done only when the selected Action is **caauth**.

CA Field Descriptions

Use the data in the following table to use the **CA** tab.

Name	Description
Name	Specifies the user-defined name referring to the Certificate Authority issuing the Digital Certificate.
CommonName	Specifies the Common Name of the Certificate Authority issuing the Digital Certificate.
KeyName	Specifies the name of the associated key pair.

Name	Description
CaUrl	Specifies the URL of the Certificate Authority issuing the Digital Certificate.
Action	Specifies the action the Certificate Authority can take: <ul style="list-style-type: none"> • noop — no operation • caauth — CA authentication • enroll — certificate enrolment request • renew — certificate renew request • remove — remove the subject certificate obtained online from the CA • install — install the subject certificate obtained online from the CA • getCrl — retrieve the Certificate Revocation List (CRL) from the CRL Distribution Point (CDP).
ActionChallengePassword	Specifies the challenge password required to perform the SCEP operation.
LastActionStatus	Specifies the status of the last action: <ul style="list-style-type: none"> • none - No action is performed yet • success - Execution of the action triggered is completed successfully • failed - Execution of the action triggered has failed • inProgress - Execution of the action triggered is in progress
LastActionFailureReason	Specifies the reason of failure for the last action performed by the Certificate Authority.
InstallRootCaFileName	Specifies the certificate file obtained offline from the Root Certificate Authority.
SubjectCertificateValidityDays	Specifies the number of days for which subject certificate will remain valid. The default value is 365 days.
UsePost	Specifies the HTTP request type: URL or POST. TRUE for EJBCA and FALSE for Win2012 CA
Sha256Fingerprint	Specifies an encrypted fingerprint of the expected certificate to match.
SubjectName	Specifies the Subject Name of the subject sending the Certificate Signing Request to the Certificate Authority.
UsedFor	Specifies the name of the application the certificate uses. The default is enabled if there is only 1 CA trustpoint configured.

View the Certificate Details

Use this procedure to:

- display the configured key details for given key name.
- display the digital certificate for the given certificate index or list all the certificate details from the local store if the certificate index is not specified.
- display the CA details for given trustpoint CA name or list all the CA details from the local store if the CA name is not specified.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Certificate**.
3. Select the **Certificate** tab.

Certificate Field Descriptions

Use the data in the following table to use the **Certificate** tab.

Name	Description
Type	Specifies the certificate type.
VersionNumber	Specifies the version number of the certificate for the subject as issued by the Certificate Authority.
SerialNumber	Specifies the serial number of the certificate for the subject as issued by the Certificate Authority.
IssuerName	Specifies the name of the issuer of the certificate for the subject as issued by the Certificate Authority.
ValidStartPeriod	Specifies the start date of the validation period of the certificate for the subject as issued by the Certificate Authority.
ValidEndPeriod	Specifies the last date of the validation period of the certificate for the subject as issued by the Certificate Authority.
CertificateSignatureAlgorithm	Specifies the algorithm used for the signature of the issuer on the certificate for the subject as issued by the Certificate Authority.
CertificateSignature	Specifies the signature of the issuer on the certificate for the subject as issued by the Certificate Authority.
Subject	Specifies the details of the subject on its certificate as issued by Certificate Authority.
SubjectPublicKeyAlgorithm	Specifies the algorithm used to generate the public key of the subject for the certificate issued by the Certificate Authority.

Name	Description
SubjectPublicKey	Specifies the public key of the subject used for the Certificate Signing Request.
HasBasicConstraint	Specifies whether the certificate contains any basic certificate constraint or not.
HasKeyUsage	Specifies whether the certificate contains basic key usage constraint or not.
IsCa	Specifies whether the certificate is a ca certificate or not.
KeyUsage	Specifies the purpose of the key used in the certificate. It is represented in the form of bits as follows: <ul style="list-style-type: none"> • bit 0 - digitalSignature • bit 1 - nonRepudiation • bit 2 - keyEncipherment • bit 3 - dataEncipherment • bit 4 - keyAgreement • bit 5 - keyCertSign • bit 6 - cRLSign • bit 7 - encipherOnly • bit 8 - decipherOnly
Status	Specifies the status of the certificate.
Installed	Specifies whether the certificate is installed or not.
CdpUrl	Specifies the CDP URL present in the Digital Certificate Extensions field.
OcspUrl	Specifies the OCSP URL present in the Digital Certificate AIA field.
ExtendedKeyUsage	Indicates the purpose for which the key is used in addition to or in place of the basic purposes indicated in the key-usage field of the certificate.
RevocationStatus	Specifies the revocation status of the certificate.

View the Certificate Store

About This Task

Use the following procedure to view the online, offline and root certificates in the local store.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Certificate**.
3. Select the **Certificate Store** tab.

Certificate Store Field Descriptions

Use the data in the following table to use the **Certificate Store** tab.

Name	Description
CommonName	Specifies the Common Name of the Certificate Authority issuing the Digital Certificate.
Type	Specifies the certificate type.
VersionNumber	Specifies the version number of the certificate for the subject as issued by the Certificate Authority.
SerialNumber	Specifies the serial number of the certificate for the subject as issued by the Certificate Authority.
IssuerName	Specifies the name of the issuer of the certificate for the subject as issued by the Certificate Authority.
ValidStartPeriod	Specifies the start date of the validation period of the certificate for the subject as issued by the Certificate Authority.
ValidEndPeriod	Specifies the last date of the validation period of the certificate for the subject as issued by the Certificate Authority.
CertificateSignatureAlgorithm	Specifies the algorithm used for the signature of the issuer on the certificate for the subject as issued by the Certificate Authority.
CertificateSignature	Specifies the signature of the issuer on the certificate for the subject as issued by the Certificate Authority.
Subject	Specifies the details of the subject on its certificate as issued by Certificate Authority.
SubjectPublicKeyAlgorithm	Specifies the algorithm used to generate the subject's public key for the certificate issued by the Certificate Authority.
SubjectPublicKey	Specifies the public key of the subject used for Certificate Signing Request.
HasBasicConstraint	Specifies whether certificate contains basic certificate constraint.
HasKeyUsage	Specifies whether certificate contains basic key usage constraint.
IsCa	Specifies if the certificate is a CA certificate or not.

Name	Description
KeyUsage	Specifies the purpose of the key used in the certificate. It is represented in the form of bits as follows: <ul style="list-style-type: none">• bit 0 - digitalSignature• bit 1 - nonRepudiation• bit 2 - keyEncipherment• bit 3 - dataEncipherment• bit 4 - keyAgreement• bit 5 - keyCertSign• bit 6 - cRLSign• bit 7 - encipherOnly• bit 8 - decipherOnly
Status	Specifies the status of the certificate.
Installed	Specifies whether the certificate is installed or not.
CdpUrl	Specifies the CDP URL present in the Digital Certificate Extensions field.
OscpUrl	Specifies the OCSP URL present in the Digital Certificate AIA field.
ExtendedKeyUsage	Indicates the purpose for which the key is used in addition to or in place of the basic purposes indicated in the key-usage field of the certificate.
CaFileName	Specifies the certificate file obtained offline from the Root Certificate Authority.



NEW! Segmented VRF

[Segmented VRF CLI Configuration](#) on page 3453

[Segmented VRF EDM Configuration](#) on page 3462

Table 296: Segmented VRF product support

Feature	Product	Release introduced
Segmented VRF Note: Only platforms that support multiple VRFs support this feature.	4220 Series	Not Supported The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 9.4 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.4
	5720 Series	Fabric Engine 9.4
	7520 Series	Fabric Engine 9.4
	7720 Series	Fabric Engine 9.4
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 9.4
	VSP 7400 Series	VOSS 9.4

Use a Segmented VRF (or Segmented Layer 3 VSN) to control, isolate, and secure traffic flows across the network by creating isolated routing domains within a single, traditional VRF or Layer 3 VSN. You can use a Segmented VRF to segment customers, internal services, functions, or security zones. You determine how, or if, the segments interconnect. Assign I-SIDs to each routing domain to advertise routes between the domains.

After you enable Segmented VRF on a traditional VRF, the switch separates the VRF into three access areas, or trust levels, called segments:

- Trusted segment—connects to secure and controlled networks
- Unrestricted segment—permits open communication without trust-based restrictions
- Untrusted segment—connects to potentially unsafe, or public networks

Assign VLAN routing interfaces to the three segments. The switch makes routing decisions based on the VRF context and access type. For more information, see [Routing Interface Types for Segmented VRF](#) on page 3448 and [Route Classification for Segmented VRF](#) on page 3449.

At a high-level, take the following actions to use Segmented VRF:

1. Enable Segmented VRF on a traditional VRF to create three segments within the VRF. The traditional VRF configuration becomes the trusted segment.
2. Associate an I-SID with the trusted segment to advertise trusted routes to the other segments.
3. (Optional) Associate I-SIDs with the unrestricted and untrusted segments. If you do not assign an I-SID to the unrestricted and untrusted segments, the switch automatically derives a value based on the trusted I-SID.
4. (Optional) Specify BGP neighbors in the unrestricted segment. By default, BGP neighbors are configured in the trusted segment.
5. (Optional) Create a CLIP interface in the unrestricted or trusted segment.
6. (Optional) Enable IP VPN on the trusted segmented.
7. Associate a VLAN routing interface with each segment.

For detailed configuration steps, see the following tasks:

- [Segmented VRF CLI Configuration](#) on page 3453
- [Segmented VRF EDM Configuration](#) on page 3462

The sections that follow provide conceptual information about this feature.

Routing Interface Types for Segmented VRF

Assign VLAN routing interfaces to the different segments. The following table defines how the switch handles traffic for the different routing interface types.



Note

If you enable Segmented VRF on an existing configured VRF, the existing routing interfaces associated with the VRF automatically become trusted interfaces.

Table 297: Routing interfaces by access segment

Interface Type	Trust level	Traffic handling	Use case	Example
Trusted	High	Access to all trusted segments	Internal communication between known and authenticated devices	An interface on a corporate LAN that connects to internal servers and employee devices
Unrestricted	Full	Open access	Development, testing, or public-facing services with broad access requirements	A network interface on a public web server that accepts connections from any IP address
Untrusted	Low	Limited access	Isolate external or unknown traffic from internal systems	A WAN interface on a firewall that connects to the public Internet

Route Classification for Segmented VRF

The following table defines the behavior of routes associated with the three access segments.

Table 298: Route behavior by type

Route Type	Trust level	Behavior	Use case	Example
Trusted	High	Allowed, prioritized, monitored lightly	Internal corporate network, VPN, known secure endpoints	Routing traffic to a known internal subnet using a secure gateway
Unrestricted	Full	Open access, no filtering or inspection	Open networks, default routes in less secure environments	A default route that forwards all traffic to the Internet without firewall rules
Untrusted	Low	Restricted, monitored, or blocked	Public Internet, guest network, unknown external source	Routing traffic from an external IP range through a firewall with strict inspection rules

You do not need to configure communication between the segments of the VRF. Communication between segments is based on rules, such as the following list:

- route availability algorithm, see [Figure 220](#) for an illustration of this concept
- unicast protocol configuration
- global configuration, for example, ECMP

The following figure illustrates route accessibility across segments.

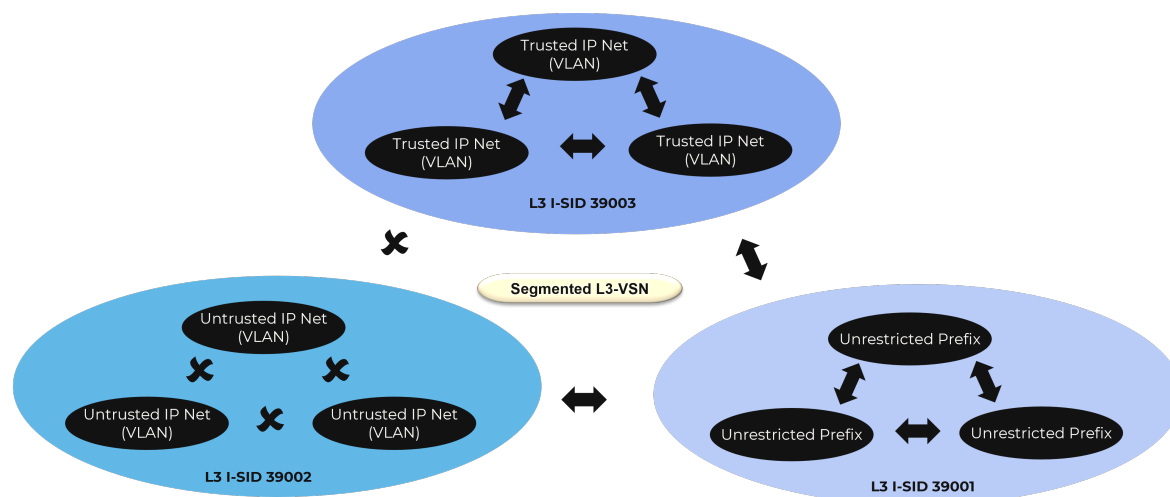


Figure 220: Segmented VRF behavior

Operations and Restrictions

The following list identifies supported operations and restrictions for Segmented VRF:

- You cannot directly configure the untrusted or unrestricted segments, nor use their IDs or names in command syntax. The switch replicates the trusted configuration to the other segments, with some exceptions as documented in this section and in [Table 299](#) on page 3451. Feature configurations that replicate across all segments are not listed.
- After you configure Segmented VRF on a switch, several CLI **show** commands, especially those that display route information or IP statistics, display the output separated by applicable segment, with the trusted segment first, followed by the unrestricted segment, and finally, the untrusted segment. For an example of this output, see the **show ip route verbose vrf** output in [Segmented VRF CLI Example](#) on page 3460.

Several **IP** folder tabs in EDM can also separate information based on applicable segment, such as the **IP > IP > ARP** and **IP > BGP > Peer Groups** tabs. You must launch the VRF context view for the Segmented VRF to see this information.

- Some CLI **show** command outputs and EDM tabs include an ACCESS or L3 ACCESS column to specify the associated Segmented VRF access type for each entry in the output.
- BGP and OSPF:
 - You can only enable OSPF on a VLAN associated with a trusted or unrestricted segment. OSPF adjacencies form only in trusted and unrestricted segments.
 - You cannot configure OSPF static neighbors on interfaces associated with an untrusted segment.
 - You can configure BGP peers or peer groups for trusted or unrestricted segments.
 - OSPF adjacencies and BGP peers redistribute routes based on owner and not route-access level. The trust level is lost during redistribution.
 - If the source VRF for BGP or OSPF inter-VRF routing is a Segmented VRF, only trusted routes are redistributed. If the destination VRF is a Segmented VRF, routes are installed only for the trusted segment and are available in the unrestricted routing table.
- After you segment an existing VRF configured for DHCP Relay, the configuration becomes part of the trusted segment. The switch associates new configurations with a segment based on the interface on which the agent IP address resides. All trust-level interfaces can use the agent IP address.
- After you segment an existing VRF configured for VRRP or Anycast IP Gateway, the configuration becomes part of the trusted segment. The switch associates new configurations with a segment based on the VLAN access level.
- To enable both Segmented VRF and IP VPN on the same VRF, you must enable Segmented VRF on the VRF before you enable IP VPN.
- Ensure the I-SIDs you configure for Segmented VRF, or that the switch automatically derives, do not conflict with I-SIDs used for other purposes in the network.

- The Layer 3 VSN I-SID you associate with an access segment must be used with the same access segment across the Fabric network. For example, do not configure I-SID 1011 for a trusted segment on BEB A and for an untrusted segment on BEB B.
- To use Segmented VRF in a Multi-area SPB deployment, you must create the Segmented VRFs, with IP VPN enabled and IP VPN I-SIDs, on all boundary nodes.
- VLANs:
 - The Segmented Management Instance VLAN cannot use an untrusted or unrestricted segment VLAN.
 - The Layer 2 VSN must have the same access type on all routers.
 - To change the VLAN ownership, you must remove the original configuration first.
 - The VLAN associated with a segment can be a VLAN with Secondary IP Interfaces.

The following table identifies feature replication exceptions across the three segments.

Table 299: Feature exceptions across segments

Feature	Trusted	Unrestricted	Untrusted	Notes
<ul style="list-style-type: none"> • BFD • Brouter interfaces • DvR, including redistribution • IPv6 • Multicast, including IP Multicast config-lite for Fabric Connect • OSPF tags in IS-IS • RIP 	Y	N	N	Configuration not replicated to unrestricted or untrusted RIP routes learned on trusted segments are synchronized to unrestricted
IS-IS Accept Policies	Y	N	N	Accepted I-SID can only be a trusted I-SID Configuration not replicated, with the following exceptions: <ul style="list-style-type: none"> • default and adv-rtr policy configuration replicated to all segments • I-SID and I-SID list configuration replicated to unrestricted
I-SID Lists	Y	Y	N	Configuration not replicated to untrusted

Table 299: Feature exceptions across segments (continued)

Feature	Trusted	Unrestricted	Untrusted	Notes
CLIP interfaces	Y	Y	N	Can add in either trusted or unrestricted Accessible from trusted and unrestricted if added in trusted Accessible from all segments if added in unrestricted
Inter-VRF Routing	Y	Y	N	Configuration not replicated to untrusted

Use Case Examples

Typically, position a firewall on an unrestricted VLAN, which means it can respond to untrusted as well as trusted VLANs. Position unsecure IoT devices on untrusted VLANs, which means they are restricted to only reach a subset (unrestricted) of destinations.

Use untrusted segments to build a larger number of segmented communication groups, using VLAN and I-SID, with devices that can only communicate with each other, and route to the unrestricted segment or firewall. For example, a multi-device, third-party IoT infrastructure with remote access for management.

Use untrusted segments to build many isolated IP subnets, using Private VLAN and E-Tree I-SID, that can only communicate with, and route to, the unrestricted segment or firewall. For example, a single-device, third-party IoT infrastructure with remote access for management.

The following image illustrates that you can combine Segmented VRF untrusted segments with VLANs, and E-Tree/Private VLANs, for full Layer 2 isolation.

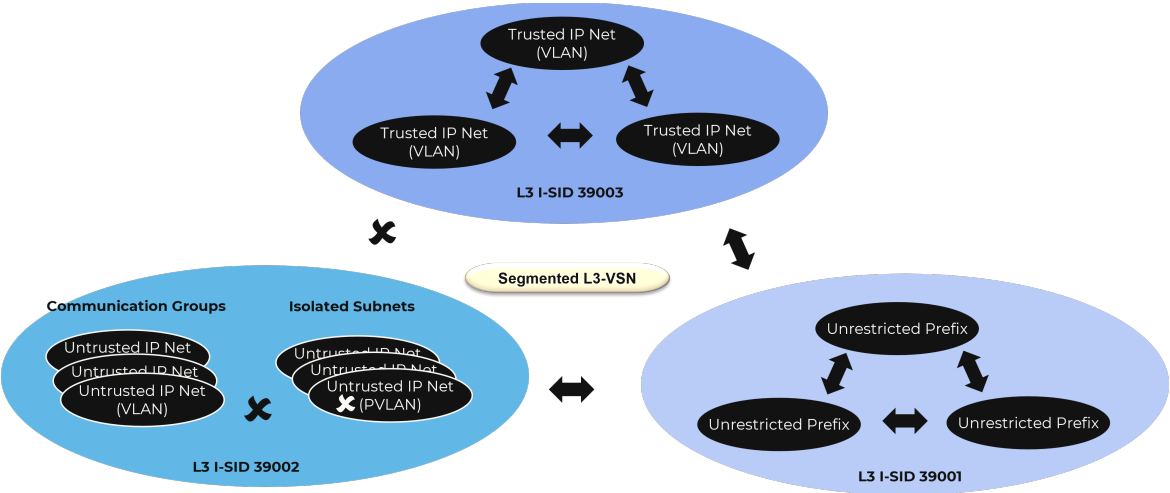


Figure 221: Communication groups and isolated networks

NEW! Segmented VRF CLI Configuration

Because the switch replicates most of the trusted configuration to the other segments, this task flow focuses on the steps that are specific to Segmented VRF. This task flow does not include traditional VLAN, CLIP, or VRF configuration.

For an example configuration, see [Segmented VRF CLI Example](#) on page 3460.



Note

After you configure Segmented VRF on a switch, several CLI **show** commands, especially those that display route information or IP statistics, display the output separated by applicable segment, with the trusted segment first, followed by the unrestricted segment, and finally, the untrusted segment. For an example of this output, see the **show ip route verbose vrf** output in the example mentioned in the preceding paragraph.

Table 300: Segmented VRF CLI Task Flow

Steps	Procedure	Description
1	Enable Segmented VRF on page 3454	Enable Segmented VRF on the traditional VRF.
2	Associate I-SIDs With VRF Segments on page 3455	Associate I-SIDs with the VRF segments to advertise routes between segments. The trusted I-SID is a mandatory configuration. If you do not assign an I-SID to the unrestricted or untrusted segment, the switch automatically derives a value based on the trusted I-SID.
3	Configure a BGP Neighbor in an Unrestricted Segment on page 3456	Optional. Configure a BGP neighbour in the unrestricted segment. By default, BGP neighbors configured on the Segmented VRF are part of the trusted segment.
4	Create a CLIP Interface in an Unrestricted Segment on page 3457	Optional. Create a CLIP interface in the unrestricted segment if it needs to be accessible from hosts in all three segments. If you do not specify the interface as unrestricted, a CLIP interface associated with a Segmented VRF is part of the trusted segment and is only accessible from hosts in the trusted and unrestricted segments.
5	Associate a VLAN Routing Interface With a VRF Segment on page 3458	Associate VLANs with each of the three segments.

Table 300: Segmented VRF CLI Task Flow (continued)

Steps	Procedure	Description
6	Create an IP VPN Instance on a VRF on page 4086	Optional. Create an IP VPN instance to advertise IP routes from the Segmented VRF.
7	Verify the Segmented VRF Configuration on page 3459	Verify the configuration.

NEW Enable Segmented VRF



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Create and configure the VRF.
- Confirm IP VPN is disabled on the VRF. You can enable IP VPN after you enable Segmented VRF.

About This Task

The traditional VRF configuration becomes the trusted segment for trusted access and the switch creates segments for unrestricted and untrusted access. The new segments use the name of the traditional VRF and append either -unr for unrestricted or -unt for untrusted. For non-trusted segments, the switch truncates the traditional VRF name to accommodate the suffix and remain within allowed character limits.

If you enable Segmented VRF on an existing configured VRF, the existing routing interfaces associated with the VRF automatically become trusted interfaces.



Note

Segmented VRFs use more system resources than traditional VRFs (Layer 3 VSNs). For scaling information, see *Fabric Engine Release Notes*.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable Segmented VRF on the traditional VRF:

```
segmented
```

NEW! Associate I-SIDs With VRF Segments



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF. This task assumes an I-SID was not previously assigned to the VRF. If the VRF has an I-SID, it becomes the trusted I-SID after you segment the VRF.
- Create three I-SIDs to use for the segmented access levels.

About This Task



Important

Ensure the I-SIDs you configure for Segmented VRF, or that the switch automatically derives, do not conflict with I-SIDs used for other purposes in the network.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Associate an I-SID with the trusted segment to advertise routes learned on that routing interface type. The trusted I-SID configuration is mandatory:

```
i-sid <0-16777215> [trusted]
```
3. (Optional) Associate I-SIDs with the unrestricted or untrusted segment to advertise routes learned on that routing interface type:

```
i-sid <0-16777215> <unrestricted | untrusted>
```



Important

If you do not assign an I-SID to the unrestricted and untrusted segments, the switch automatically derives a value based on the trusted I-SID.

NEW! Variable Definitions

The following table defines parameters for the **i-sid <0-16777215>** command.

Variable	Value
<0-16777215>	Specifies the service instance identifier (I-SID).
trusted	Assigns an I-SID to advertise trusted routes to trusted and unrestricted segments.
unrestricted	Assigns an I-SID to advertise unrestricted routes to all segments. If you do not assign an I-SID to the unrestricted segment, the switch automatically derives a value based on the trusted I-SID, (trusted I-SID + 1).
untrusted	Assigns an I-SID to advertise untrusted routes to restricted segments. If you do not assign an I-SID to the untrusted segment, the switch automatically derives a value based on the trusted I-SID, (trusted I-SID + 2).

NEW! Configure a BGP Neighbor in an Unrestricted Segment**Note**

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF.

About This Task**Note**

If you do not enter this command with the *unrestricted* parameter, BGP neighbors configured on the Segmented VRF are part of the trusted segment.

The switch adds existing BGP neighbors in the VRF context to the trusted segment automatically.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure a BGP neighbor in the unrestricted segment:

```
ip bgp neighbor WORD<0-1536> unrestricted
```

NEW! Variable Definitions

The following table defines parameters for the **ip bgp neighbor WORD<0-1536>** command.

Variable	Value
unrestricted	Configures a BGP neighbor as part of the unrestricted segment. If you do not enter this parameter, new BGP neighbors that you configure in the VRF context are part of the trusted segment.
WORD<0-1536>	Specifies the peer IP address or the peer group name.

NEW! Create a CLIP Interface in an Unrestricted Segment



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF.

About This Task

If you create the interface in the trusted segment, it is only accessible from hosts in the trusted and unrestricted segments. Create a CLIP interface in the unrestricted segment if it needs to be accessible from hosts in all three segments.

Procedure

1. Enter Loopback Interface Configuration mode

```
enable
```

```
configure terminal
```

```
interface Loopback <1-256>
```

2. Create a CLIP interface in the unrestricted segment:

```
ip address {<A.B.C.D/X> | <A.B.C.D> <A.B.C.D>} vrf WORD<1-16>  
unrestricted
```

NEW! Variable Definitions

The following table defines parameters for the **ip address** command.

Variable	Value
<A.B.C.D/X> <A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask for the CLIP interface in either of the following notation formats: A.B.C.D/X or A.B.C.D A.B.C.D.
unrestricted	Specifies that the interface belongs to the unrestricted segment of a Segmented VRF. If you omit this parameter, the interface is part of the trusted segment and is only accessible from hosts in the trusted and unrestricted segments. Use the <i>unrestricted</i> parameter to make the interface accessible from hosts in all three segments. The default is trusted.
vrf WORD<1-16>	Specifies the Segmented VRF by name.

NEW! Associate a VLAN Routing Interface With a VRF Segment



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF.
- Create and configure a VLAN routing interface to associate with each segment.

About This Task

If you enable Segmented VRF on an existing configured VRF, the existing routing interfaces associated with the traditional VRF automatically become trusted interfaces.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Associate the VLAN with the appropriate segment:

```
vrf WORD<1-16> <trusted | unrestricted | untrusted>
```

NEW! Variable Definitions

The following table defines parameters for the **vrf WORD<1-16>** command.

Variable	Value
trusted	Associates the VLAN with the trusted segment. If you do not use this parameter, the VLAN is associated with the trusted segment by default.
untrusted	Associates the VLAN with the untrusted segment. If you do not use this parameter, the VLAN is associated with the trusted segment by default.
unrestricted	Associates the VLAN with the unrestricted segment. If you do not use this parameter, the VLAN is associated with the trusted segment by default.
WORD<1-16>	Specifies the name of the traditional VRF. You cannot use untrusted or unrestricted segment names in command syntax.

NEW! Verify the Segmented VRF Configuration



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Confirm the traditional VRF is now a Segmented VRF:

```
show ip vrf [detail]
```

3. Verify the I-SID configuration for each segment:

```
show ip ipvpn
```

4. View route information for each segment:

```
show ip route [verbose] vrf WORD<1-16>
```

```
show ip route unrestricted-segment {vrf WORD<1-16> | vrfids  
WORD<0-512>}
```

```
show ip route untrusted-segment {vrf WORD<1-16> | vrfids WORD<0-512>}
```

```
show ip route trusted-segment {vrf WORD<1-16> | vrfids WORD<0-512>}
```

NEW Segmented VRF CLI Example

Create a VRF named sales, and enable Segmented VRF. Configure a Layer 3 I-SID for the VRF, which becomes the trusted I-SID, but permit the switch to derive the untrusted and unrestricted I-SIDs automatically. Identify BGP neighbors for both the trusted segment and the unrestricted segment. Associate VLAN 100 with the trusted segment, VLAN 200 with the unrestricted segment, and VLAN 300 with the untrusted segment.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip vrf sales
Switch:1(config)#router vrf sales
Switch:1(router-vrf)#segmented
Switch:1(router-vrf)#i-sid 1010
Switch:1(router-vrf)#ipvpn
Switch:1(router-vrf)#ipvpn enable
Switch:1(router-vrf)#ip bgp neighbor 192.0.2.1
Switch:1(router-vrf)#ip bgp neighbor trPeerGroup
Switch:1(router-vrf)#ip bgp neighbor 198.51.100.1 unrestricted
Switch:1(router-vrf)#ip bgp neighbor unresPeerGroup unrestricted
Switch:1(router-vrf)#interface vlan 100
Switch:1(config-if)#vrf sales
Switch:1(router-vrf)#interface vlan 200
Switch:1(config-if)#vrf sales unrestricted
Switch:1(config-if)#interface vlan 300
Switch:1(config-if)#vrf sales untrusted
### VRF and VLAN configuration not shown###
```

Verify the configuration.

Notice that the **show ip vrf** output indicates there are 3 VRFs on the switch but 5 IDs are allocated. After you segment the sales VRF, the switch uses the existing VRF ID for the trusted segment and automatically allocates VRF IDs to the new untrusted and unrestricted segments, starting with the highest available VRF ID. You can see the IDs assigned to each segment in the **show ip vrf detail** output.

Also notice the different origins between the segment types.

```
Switch:1#show ip vrf
```

VRF INFORMATION											
VRF	VLAN	ARP	RIP	OSPF	BGP	PIM	NBRv6	RIPng	OSPFv3	PIM6	VRF IDs
COUNT	COUNT	COUNT	COUNT	COUNT	COUNT	COUNT	COUNT	COUNT	COUNT	COUNT	ALLOCATED
3	8	4	1	4	1	1	0	1	1	1	5

```

VRF          VRF  VLAN  ARP          NBRv6          UNICAST  SD-WAN LOCAL  VRF
NAME         ID   COUNT COUNT  RIP    OSPF  BGP    PIM    COUNT  RIPng  OSPFv3  PIM6    ACTIVE  BREAKOUT  ORIGIN  SEGMENTED
-----
GlobalRouter 0    3    1    TRUE  TRUE  TRUE  TRUE  0    TRUE  TRUE  TRUE  TRUE  TRUE    FALSE    DYNAMIC  FALSE
sales        1    2    3    FALSE TRUE  TRUE  FALSE 0    FALSE FALSE FALSE TRUE  FALSE    CONFIG   TRUE
MgmtRouter   512  1    0    FALSE FALSE FALSE FALSE 0    FALSE FALSE FALSE TRUE  FALSE    DYNAMIC  FALSE

Switch:1#show ip vrf detail
=====
VRF INFORMATION
=====
VRF          VRF  VLAN  ARP          NBRv6          UNICAST  SD-WAN LOCAL  VRF
NAME         ID   COUNT COUNT  RIP    OSPF  BGP    PIM    COUNT  RIPng  OSPFv3  PIM6    ACTIVE  BREAKOUT  ORIGIN  SEGMENTED
-----
3            8    4    1    4    1    1    0    1    1    1    1    5
COUNT      COUNT COUNT COUNT COUNT COUNT COUNT COUNT COUNT COUNT COUNT COUNT ALLOCATED
-----
3            8    4    1    4    1    1    0    1    1    1    1    5

VRF          VRF  VLAN  ARP          NBRv6          UNICAST  SD-WAN LOCAL  VRF
NAME         ID   COUNT COUNT  RIP    OSPF  BGP    PIM    COUNT  RIPng  OSPFv3  PIM6    ACTIVE  BREAKOUT  ORIGIN  SEGMENTED
-----
GlobalRouter 0    3    1    TRUE  TRUE  TRUE  TRUE  0    TRUE  TRUE  TRUE  TRUE  TRUE    FALSE    DYNAMIC  FALSE
sales        1    2    3    FALSE TRUE  TRUE  FALSE 0    FALSE FALSE FALSE TRUE  FALSE    CONFIG   TRUE
sales-unr    510  2    0    FALSE TRUE  FALSE FALSE 0    FALSE FALSE FALSE TRUE  FALSE    SEGMENTED TRUE
sales-unr    511  0    0    FALSE TRUE  FALSE FALSE 0    FALSE FALSE FALSE TRUE  FALSE    SEGMENTED TRUE
MgmtRouter   512  1    0    FALSE FALSE FALSE FALSE 0    FALSE FALSE FALSE TRUE  FALSE    DYNAMIC  FALSE

5 out of 5 Total Num of VRF Entries displayed.

Switch:1#show ip ipvpn
=====
IPv4 IPVPN
=====
VRF Name      VRF ID  IPv4 IPVPN  IPv6 IPVPN  Trusted  Unrestricted  Untrusted  I-SID Name
-----
sales         1       enabled disabled 1010      1011      1012      ISID-1010

1 out of 1 Total IPv4 L3 VSN, 1 active IPv4 and 0 active IPv6 displayed.

Switch:1#show ip route verbose vrf sales
=====
IP Route - VRF sales
=====
DST          MASK          NEXT          NH          INTER          COST  FACE  PROT  AGE  TYPE  PRF  ACCESS
VRF/ISID
-----
[VRF Segment Lookup for: Trusted]:
0.0.0.0      0.0.0.0      2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
10.1.1.0     255.255.255.0 100.1.1.4     sales        1    100  OSPF  0    IB    125 Trusted
20.1.1.0     255.255.255.0 20.1.1.6     -            1    20   LOC   0    DB    0    Trusted
50.1.1.0     255.255.255.0 2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
60.1.1.0     255.255.255.0 2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
80.1.1.0     255.255.255.0 80.1.1.6     -            1    80   LOC   0    IB    0    Unrestricted
100.1.1.0    255.255.255.0 100.1.1.6    -            1    100  LOC   0    DB    0    Trusted

[VRF Segment Lookup for: Unrestricted]:
0.0.0.0      0.0.0.0      2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
10.1.1.0     255.255.255.0 100.1.1.4     sales        10    100  OSPF  0    IB    125 Trusted
20.1.1.0     255.255.255.0 20.1.1.6     -            1    20   LOC   0    IB    0    Trusted
30.1.1.0     255.255.255.0 30.1.1.6     -            1    30   LOC   0    DB    0    Untrusted
40.1.1.0     255.255.255.0 40.1.1.6     -            1    40   LOC   0    DB    0    Untrusted
50.1.1.0     255.255.255.0 2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
60.1.1.0     255.255.255.0 2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
80.1.1.0     255.255.255.0 80.1.1.6     -            1    80   LOC   0    DB    0    Unrestricted
100.1.1.0    255.255.255.0 100.1.1.6    -            1    100  LOC   0    IB    0    Trusted

[VRF Segment Lookup for: Untrusted]:
0.0.0.0      0.0.0.0      2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
50.1.1.0     255.255.255.0 2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
60.1.1.0     255.255.255.0 2407          sales-unr    10    4000  ISIS  0    IBSV  7    Unrestricted
80.1.1.0     255.255.255.0 80.1.1.6     -            1    80   LOC   0    IB    0    Unrestricted

20 out of 20 Total Num of Route Entries, 20 Total Num of Dest Networks displayed.

TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route, R=Remote SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

```

NEW! Segmented VRF EDM Configuration

Because the switch replicates most of the trusted configuration to the other segments, this task flow focuses on the steps that are specific to Segmented VRF. This task flow does not include traditional VLAN, CLIP, or VRF configuration.



Note

After you configure Segmented VRF on a switch, several **IP** folder tabs in EDM can separate information based on applicable segment, such as the **IP > IP > ARP** and **IP > BGP > Peer Groups** tabs. You must launch the VRF context view for the Segmented VRF to see this information.

Table 301: Segmented VRF EDM Task Flow

Steps	Procedure	Description
1	Enable Segmented VRF on page 3463	Enable Segmented VRF on the traditional VRF.
2	Associate I-SIDs With VRF Segments on page 3463	Associate I-SIDs with the VRF segments to advertise routes between segments. The trusted I-SID is a mandatory configuration. If you do not assign an I-SID to the unrestricted or untrusted segment, the switch automatically derives a value based on the trusted I-SID.
3	Configure a BGP Neighbor in an Unrestricted Segment on page 3464	Optional. Configure a BGP neighbour in the unrestricted segment. By default, BGP neighbors configured on the Segmented VRF are part of the trusted segment.
4	Create a CLIP Interface in a VRF Segment on page 3465	Optional. Create a CLIP interface in the unrestricted segment if it needs to be accessible from hosts in all three segments. If you do not specify the interface as unrestricted, a CLIP interface associated with a Segmented VRF is part of the trusted segment and is only accessible from hosts in the trusted and unrestricted segments.
5	Create an IP VPN Instance on a VRF on page 4101	Optional. Create an IP VPN instance to advertise IP routes from the Segmented VRF.
6	Associate a VLAN Routing Interface With a VRF Segment on page 3466	Associate VLANs with each of the three segments.

NEW Enable Segmented VRF



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Create and configure the VRF.
- Confirm IP VPN is disabled on the VRF. You can enable IP VPN after you enable Segmented VRF.

About This Task

The traditional VRF configuration becomes the trusted segment for trusted access and the switch creates segments for unrestricted and untrusted access. The new segments use the name of the traditional VRF and append either -unr for unrestricted or -unt for untrusted. For non-trusted segments, the switch truncates the traditional VRF name to accommodate the suffix and remain within allowed character limits.

If you enable Segmented VRF on an existing configured VRF, the existing routing interfaces associated with the VRF automatically become trusted interfaces.



Note

Segmented VRFs use more system resources than traditional VRFs (Layer 3 VSNs). For scaling information, see *Fabric Engine Release Notes*.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **VRF** tab.
4. In **Segmented**, select **true** in the row for the traditional VRF you want to segment.
5. Select **Apply**.



Note

After you segment the VRF, the switch uses the existing VRF ID for the trusted segment and automatically allocates VRF IDs to the new untrusted and unrestricted segments, starting with the highest available VRF ID.

NEW Associate I-SIDs With VRF Segments



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF. This task assumes an I-SID was not previously assigned to the VRF. If the VRF has an I-SID, it becomes the trusted I-SID after you segment the VRF.
- Create three I-SIDs to use for the segmented access levels.

About This Task



Important

Ensure the I-SIDs you configure for Segmented VRF, or that the switch automatically derives, do not conflict with I-SIDs used for other purposes in the network.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP-VPN**.
3. Select the **VPN** tab.
4. In **Trusted Isid** in the row for the VRF ID for the trusted segment, type the I-SID.
5. (Optional) In **Unrestricted Isid** or **Untrusted Isid** in the row for the VRF ID of the trusted segment, type the I-SID.



Important

If you do not assign an I-SID to the unrestricted and untrusted segments, the switch automatically derives a value based on the trusted I-SID.

6. Select **Apply**.

NEW! Configure a BGP Neighbor in an Unrestricted Segment



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF.
- Select and launch the VRF instance associated with the Segmented VRF. For more information about how to use EDM to configure a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100.

About This Task

By default, BGP neighbors configured on the Segmented VRF are part of the trusted segment.

The switch adds existing BGP neighbors in the VRF context to the trusted segment automatically.

Procedure

1. In the navigation pane of the VRF context tab, expand **Configuration > IP**.
2. Select **BGP**.
3. Select the **Peers** tab or the **Peer Groups** tab.
4. Select **Insert**.
5. Configure the necessary parameters for the peer or peer group.
6. In **AccessType**, select **unrestricted**.
7. Select **Insert**.
8. Select **Apply**.

NEW Create a CLIP Interface in a VRF Segment



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF.
- Select and launch the VRF instance associated with the Segmented VRF. For more information about how to use EDM to configure a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100.

About This Task

If you create the interface in the trusted segment, it is only accessible from hosts in the trusted and unrestricted segments. Create a CLIP interface in the unrestricted segment if it needs to be accessible from hosts in all three segments.

Procedure

1. In the navigation pane of the VRF context tab, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select **Insert**.
5. In **Interface**, assign a CLIP interface number.
6. Type the IP address.
7. Type the network mask.
8. In **AccessType**, select **unrestricted** or **trusted** to specify the segment.
9. Select **Insert**.
10. Select **Apply**.

NEW! Associate a VLAN Routing Interface With a VRF Segment



Note

This procedure does not apply to 7830 Series or to switch models that support a single active VRF. For more information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- Enable Segmented VRF on the traditional VRF.
- Create and configure a VLAN routing interface to associate with each segmented access level.

About This Task

If you enable Segmented VRF on an existing configured VRF, the existing routing interfaces associated with the VRF automatically become trusted interfaces.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**
2. Select **VLANs**.
3. Select the **Basic** tab.
4. In **Vrfid** in the row for the VLAN, double-click to display the **Vrfid** dialog box.
5. Select the VRF ID for the appropriate segment to associate with the VLAN.
6. Select **Ok**.
7. Select **Apply**.



Simple Network Management Protocol (SNMP)

[SNMPv3](#) on page 3468

[SNMP Community Strings](#) on page 3474

[SNMPv3 support for VRF](#) on page 3476

[SNMPv3 Remote Engine ID Discovery](#) on page 3476

[SNMP Configuration Using CLI](#) on page 3477

[SNMP Configuration Using EDM](#) on page 3487

Table 302: Simple Network Management Protocol product support

Feature	Product	Release introduced
Simple Network Management Protocol (SNMP) v1/2/3 (IPv4)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 302: Simple Network Management Protocol product support (continued)

Feature	Product	Release introduced
SNMP (IPv6)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

You can use the Simple Network Management Protocol (SNMP) to remotely collect management data and configure devices.

An SNMP agent is a software process that monitors the UDP port 161 for SNMP messages. Each SNMP message sent to the agent contains a list of management objects to retrieve or modify.

SNMPv3

The SNMP version 3 (v3) is the third version of the Internet Standard Management Framework and is derived from and builds upon both the original Internet Standard Management Framework SNMP version 1 (v1) and the second Internet Standard Management Framework SNMP version 2 (v2).

The SNMPv3 is not a stand-alone replacement for SNMPv1 or SNMPv2. The SNMPv3 defines security capabilities you must use in conjunction with SNMPv2 (preferred) or SNMPv1. The following figure shows how SNMPv3 specifies a user-based security model (USM) that uses a payload of either an SNMPv1 or an SNMPv2 Protocol Data Unit (PDU).

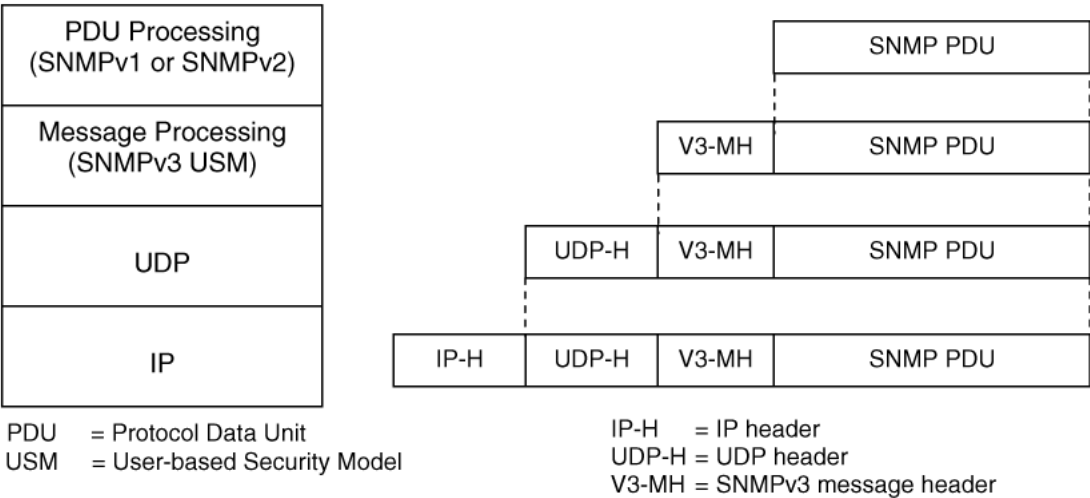


Figure 222: SNMPv3 USM

SNMPv3 is an SNMP framework that supplements SNMPv2 by supporting the following:

- New SNMP message formats
- Security for messages
- Access control
- Remote configuration of SNMP parameters

The recipient of a message can use authentication within the USM to verify the message sender and to detect if the message is altered. According to RFC2574, if you use authentication, the USM checks the entire message for integrity.

An SNMP entity is an implementation of this architecture. Each SNMP entity consists of an SNMP engine and one or more associated applications.

SNMP Engine

An SNMP engine provides services for sending and receiving messages, authenticating and encrypting messages, and controlling access to managed objects. A one-to-one association exists between an SNMP engine and the SNMP entity, which contains the SNMP engine.

EngineID

Within an administrative domain, an EngineID is the unique identifier of an SNMP engine. Because there is a one-to-one association between SNMP engines and SNMP entities, the ID also uniquely and unambiguously identifies the SNMP entity within that

administrative domain. The system generates an EngineID during the startup process. The SNMP engine contains a:

- [Dispatcher](#) on page 3470.
- [Message Processing Subsystem](#) on page 3470.
- [Security Subsystem](#) on page 3470.
- [Access Control Subsystem](#) on page 3471.

Automatic discovery of a Remote EngineID is performed so that Inform messages are populated properly and not rejected by the SNMP Manager. This automatic discovery is performed by the switch using a request to the Manager for a response that contains the EngineID. The switch then uses this Remote EngineID value to populate all subsequent Inform messages.

Dispatcher

The dispatcher is part of an SNMP engine. You can use the dispatcher for concurrent support of multiple versions of SNMP messages in the SNMP engine through the following ways:

- To send and receive SNMP messages to and from the network.
- To determine the SNMP message version and interact with the corresponding message processing model.
- To provide an abstract interface to SNMP applications for delivery of a PDU to an application.
- To provide an abstract interface for SNMP applications to send a PDU to a remote SNMP entity.

Message Processing Subsystem

The message processing subsystem prepares messages for sending and extracts data from received messages. The subsystem can contain multiple message processing models.

Security Subsystem

The security subsystem provides the following features:

- Authentication
- Privacy
- Security

Authentication

You can use authentication within the SNMPv3 to verify the message sender and whether the message is altered. If you use authentication, the integrity of the message is verified. The supported SNMPv3 authentication protocols are HMAC-MD5 and HMAC-SHA-96. By default, the switch uses HMAC-SHA1-96 with 160-bit key length.

Privacy

SNMPv3 is an encryption protocol for privacy. Only the data portion of a message is encrypted; the header and the security parameters are not. The privacy protocol that SNMPv3 supports is CBC-DES Symmetric Encryption Protocol and Advanced Encryption Standard (AES).

Security

The SNMPv3 security protects against:

- Modification of information—protects against altering information in transit.
- Masquerade—protects against an unauthorized entity assuming the identity of an authorized entity.
- Message stream modification—protects against delaying or replaying messages.
- Disclosure—protects against eavesdropping.

The SNMPv3 security also offers:

- Discovery procedure—finds the EngineID of an SNMP entity for a given transport address or transport endpoint address.
- Time synchronization procedure—facilitates authenticated communication between entities

The SNMPv3 does not protect against the following:

- Denial-of-service—prevention of exchanges between manager and agent.
- Traffic analysis—general pattern of traffic between managers and agents.

Access Control Subsystem

SNMPv3 provides a group option for access policies.

The access policy feature in the switch determines the access level for the users connecting to the device with different services like File Transfer Protocol (FTP), Trivial FTP (TFTP), and Telnet. The system access policy feature is based on the user access levels and network address. This feature covers services, such as TFTP, HTTP, SSH, and SNMP. However, with the SNMPv3 engine, the community names do not map to an access level. The View-based Access Control Model (VACM) determines the access privileges.

Use the configuration feature to specify groups for the SNMP access policy. You can use the access policy services to cover SNMP. Because the access restriction is based on groups defined through the VACM, the synchronization is made using the SNMPv3 VACM configuration. The administrator uses this feature to create SNMP users (USM community) and associate them to groups. You can configure the access policy for each group and network.

The following are feature specifications for the group options:

- After you enable SNMP service, this policy covers all users associated with the groups configured under the access policy. The access privileges are based on access allow

or deny. If you select allow, the VACM configuration determines the management information base (MIB)-views for access.

- The SNMP service is disabled by default for all access policies.
- The access level configured under **access-policy policy <id>** does not affect SNMP service. The VACM configuration determines the SNMP access rights.

User-Based Security Model

In a USM system, the security model uses a defined set of user identities for any authorized user on a particular SNMP engine. A user with authority on one SNMP engine must also have authorization on all SNMP engines with which the original SNMP engine communicates.

The USM provides the following levels of communication:

- NoAuthNoPriv—communication without authentication and privacy.
- AuthNoPriv—communication with authentication and without privacy.
- AuthPriv—communication with authentication and privacy.

The following figure shows the relationship between USM and VACM.




Figure 223: USM association with VACM

View-Based Access Control

View-based Access Control Model (VACM) provides group access, group security levels, and context based on a predefined subset of MIB objects. These MIB objects define a set of managed objects and instances.

VACM is the standard access control mechanism for SNMPv3, and it provides:

- Authorization service to control access to MIB objects at the PDU level.
- Alternative access control subsystems.

The access is based on principal, security level, MIB context, object instance, and type of access requested (read or write). You can use the VACM MIB to define the policy and control remote management.

SNMPv3 Encryption

A user-based security port for SNMPv3 is defined as a security subsystem within an SNMP engine. The switch USM uses HMAC-MD5-96 and HMAC-SHA-96 as the authentication protocols, and CBC-DES as the privacy protocol. Use USM to use other protocols instead of, or concurrently with, these protocols. CFB128-AES-128, an AES-based Symmetric Encryption Protocol, is an alternative privacy protocol for the USM.

The AES standard is the current encryption standard, Federal Information Processing Standard 140-2 (FIPS 140-2), intended to be used by the U.S. Government organizations to protect sensitive information. The AES standard is also becoming a global standard for commercial software and hardware that uses encryption or other security features.

The AES-Based Symmetric Encryption Protocol

This symmetric encryption protocol provides support for data confidentiality. The system encrypts the designated portion of the SNMP message and includes it as part of the transmitted message.

The USM specifies that the scoped PDU is the portion of the message that requires encryption. An SNMP engine that can legitimately originate messages on behalf of the appropriate user shares a secret value, in combination with a timeliness value and a 64-bit integer, used to create the (localized) encryption/decryption key and the initialization vector.

The AES Encryption Key and Initialization Vector

The AES encryption key uses the first 128 bits of the localized key. The 128-bit Initialization Vector (IV) is the combination of the authoritative SNMP engine 32-bit `snmpEngineBoot`, the SNMP engine 32-bit `snmpEngineTime`, and a local 64-bit integer. The system initializes the 64-bit integer to a pseudo-random value at startup time.

Data Encryption

The switch handles data encryption in the following manner:

1. The system treats data as a sequence of octets.
2. The system divides the plaintext into 128-bit blocks.

The first input block is the IV, and the forward cipher operation is applied to the IV to produce the first output block.

3. The system produces the first cipher text block by executing an exclusive-OR function on the first plaintext block with the first output block.
4. The system uses the cipher text block as the input block for the subsequent forward cipher operation.
5. The system repeats the forward cipher operation with the successive input blocks until it produces a cipher text segment from every plaintext segment.
6. The system produces the last cipher text block by executing an exclusive-OR function on the last plaintext segment of r bits (r is less than or equal to 128) with the segment of the r most significant bits of the last output block.

Data Decryption

The switch handles data decryption in the following manner:

1. In CFB decryption, the IV is the first input block, the system uses the first cipher text for the second input block, the second cipher text for the third input block, and this continues until the system runs out of blocks to decrypt.
2. The system applies the forward cipher function to each input block to produce the output blocks.
3. The system passes the output blocks through an exclusive-OR function with the corresponding cipher text blocks to recover the plaintext blocks.
4. The system sends the last cipher text block (whose size r is less than or equal to 128) through an exclusive-OR function with the segment of the r most significant bits of the last output block to recover the last plaintext block of r bits.

Trap Notifications

You configure traps by creating SNMP trap notifications, creating a target address to which you want to send the notifications, and specifying target parameters. For more information about how to configure trap notifications, see [Logs and Traps](#) on page 2535.

SNMP Community Strings

For security reasons for SNMPv1 and SNMPv2, the SNMP agent validates each request from an SNMP manager before responding to the request by verifying that the manager belongs to a valid SNMP community. An SNMP community is a logical relationship between an SNMP agent and one or more SNMP managers (the manager software implements the protocols used to exchange data with SNMP agents). You define communities locally at the agent level.

The agent establishes one community for each combination of authentication and access control characteristics that you choose. You assign each community a unique name (community string), and all members of a community have the same access privileges, either read-only or read-write:

- Read-only: members can view configuration and performance information.
- Read-write: members can view configuration and performance information, and change the configuration.

By defining a community, an agent limits access to its MIB to a selected set of management stations. By using more than one community, the agent can provide different levels of MIB access to different management stations.

SNMP community strings are used when a user logs on to the device over SNMP, for example, using an SNMP-based management software. You set the SNMP community strings using CLI . If you have read/write/all access authority, you can modify the SNMP community strings for access to the device through Enterprise Device Manager (EDM).

You are provided with community strings for SNMPv1 and SNMPv2. If you want to use SNMPv3 only, you must disable SNMPv1 and SNMPv2 access by deleting the default community string entries and create the SNMPv3 user and group.[SNMPv3](#).

**Note**

If you enable enhanced secure mode, the switch does not support the default SNMPv1 and default SNMPv2 community strings, and default SNMPv3 user name. The individual in the administrator access level role can configure a non-default value for the community strings, and the switch can continue to support SNMPv1 and SNMPv2. The individual in the administrator access level role can also configure a non-default value for the SNMPv3 user name and the switch can continue to support SNMPv3.

If you disable enhanced secure mode, the SNMPv1 and SNMPv2 support for community strings remains the same, and the default SNMPv3 user name remains the same. Enhanced secure mode is disabled by default.

For more information on enhanced secure mode, see [Enhanced Secure Mode](#) on page 266.

The following table lists the default community strings for SNMPv1 and SNMPv2.

VRF	Default community string	Access
GlobalRouter VRF	public	Read access
	private	Write access
ManagementRouter VRF	public:512	Read access
	private:512	Write access

Community strings are encrypted using the AES encryption algorithm. The system does not display community strings on the device and are not stored in the configuration file.

**Caution**

Security risk

For security reasons, as a best practice, set the community strings to values other than the factory defaults.

The switch handles community string encryption in the following manner:

- When the device starts up, community strings are restored from the hidden file.

- When the SNMP community strings are modified, the modifications are updated to the hidden file.
- Stale snmp-server community entries for different VRFs that the system displays it after reboot with no VRFs . On an node with any valid config file saved with more than the default vrf0 , snmp_community entries for that VRF are created and maintained in a separate txt file, snmp_comm.txt, on every boot. The node reads this file and updates the snmp communities available on the node. As a result for a boot with config having no VRFs, you can still see snmp_community entries for VRFs other than the globalRouter vrf0.

Hsecure with SNMP

If you enable hsecure, the system disables SNMPv1, SNMPv2 and SNMPv3. If you want to use SNMP, you must use the command **no boot config flag block-snmp** to re-enable SNMP.

SNMPv3 support for VRF

Use Virtual Router Forwarding (VRF) to offer networking capabilities and traffic isolation to customers that operate over the same node (switch). Each virtual router emulates the behavior of a dedicated hardware router and is treated by the network as a separate physical router. You can use VRF Lite to perform the functions of many routers using a single router running VRF Lite. This substantially reduces the cost associated with providing routing and traffic isolation for multiple clients.

SNMPv3 Remote Engine ID Discovery

The following sections contain information about SNMPv3 remote engine ID discovery.

Traps and Informs

SNMPv3 supports two types of notifications, traps and informs.

Traps are unacknowledged notifications sent by agents to managers. Traps are generated by the agent, and the authoritative SNMP engine for a trap packet is the sending SNMP agent. Because the generator of the message and the authoritative engine are the same, there is no need for the SNMPv3 discovery process.

Informs are acknowledged notifications. An agent sends an inform notification and waits for acknowledgement. If the agent does not receive the acknowledgement within the configured timeout period, it resends the inform notification. The agent continues to resend the inform notification until a reply is received or until the maximum retry value is reached. For information about configuring an inform timeout period or configuring a maximum retry value, see [Configure SNMP settings](#) on page 3477.

Remote Engine ID Discovery

Inform packets must contain the management (remote) SNMP engine ID. The agent generates informs, but the authoritative SNMP engine is the manager. To generate valid inform packets and avoid manual configuration of the manager SNMP engine ID, the agent must discover the SNMP engine ID of the manager.

The agent discovers the management SNMP Engine ID by sending a probe message to the manager. The manager response to the probe message is a report message that contains the SNMP engine ID of the authoritative SNMP engine. The agent stores the SNMP engine ID received from the manager in the engine table, and sends inform packets using the manager SNMP engine ID. If the manager SNMP engine ID changes, the discovery process updates the manager SNMP engine ID value in the engine table.



Note

Config files including local user with specified security level (noAuthNoPriv or authNoPriv or authPriv) must be saved to be compatible with remote engine ID discovery.

SNMP Configuration Using CLI

Configure the SNMP engine to provide services to send and receive messages, authenticate and encrypt messages, and control access to managed objects. A one-to-one association exists between an SNMP engine and the SNMP entity.

Configure SNMP settings

Configure Simple Network Management Protocol (SNMP) to define or modify the SNMP settings, and specify the security level for SNMP communications.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the generation of authentication traps:
`snmp-server authentication-trap enable`
3. Configure the contact information for the system:
`snmp-server contact WORD<0-255>`
4. Create an SNMPv1 server host:
`snmp-server host WORD<1-256> [port <1-65535>] v1 WORD<1-32> [filter WORD<1-32>]`
5. Create an SNMPv2 server host:
`snmp-server host WORD<1-256> [port <1-65535>] v2c WORD<1-32> [inform] [timeout <1-2147483647>] [retries <0-255>] [mms <0-2147483647>] [filter WORD<1-32>]`

6. Create an SNMPv3 server host:

```
snmp-server host WORD<1-256> [port <1-65535>] v3 {noAuthNoPriv|
authNoPriv|authPriv WORD<1-32> [inform] [timeout <1-2147483647>]
[retries <0-255>]] [filter WORD<1-32>]
```

7. Configure the system location:

```
snmp-server location WORD<0-255>
```

8. Enable login-success traps:

```
snmp-server login-success-trap enable
```

9. Configure the SNMP system name:

```
snmp-server name WORD<0-255>
```

10. Create a new entry in the notify filter table:

```
snmp-server notify-filter WORD<1-32> WORD<1-32>
```

Example

Enable the generation of SNMP traps. Configure the contact information for the system. Configure hosts to receive SNMP notifications:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server authentication-trap enable
Switch:1(config)#snmp-server contact xxxx@company.com
Switch:1(config)#snmp-server host 192.0.2.16 port 1 v1 SNMPv1 filter SNMPfilterv1
```

Variable Definitions

The following table defines parameters for the `snmp-server` command.

Variable	Value
<code>authentication-trap {enable}</code>	Enables generation of SNMP server authentication traps.
<code>contact WORD<0-255></code>	Changes the sysContact information for the switch. WORD<0-255> is an ASCII string from 0–255 characters (for example, a phone extension or email address).

Variable	Value
<pre>host WORD<1-256> [port <1-65535>] {v1 WORD<1-32> v2c WORD<1-32> [inform] [timeout <1-2147483647>] [retries <0-255>] [mms <0-2147483647>]] v3 {noAuthPriv authNoPriv authPriv} WORD<1-32> [inform [timeout <1-2147483647>] [retries <0-255>]]} [filter WORD<1-32>]</pre>	<p>Configures hosts to receive SNMP notifications.</p> <ul style="list-style-type: none"> <code>host WORD<1-256></code> specifies the IPv4 or IPv6 host address <code>port <1-65535></code> specifies the port number <code>v1 WORD<1-32></code> specifies the SNMPv1 security name <code>v2c WORD<1-32></code> specifies the SNMPv2 security name <code>inform</code> specifies the notify type <code>timeout <1-2147483647></code> specifies the timeout value <code>retries <0-255></code> specifies the number of retries <code>mms < 1-2147483647></code> specifies the maximum message size <code>v3</code> specifies SNMPv3 <code>noAuthPriv authNoPriv authPriv</code> specifies the security level <code>WORD<1-32></code> specifies the user name <code>filter</code> specifies a filter profile name
<pre>location WORD<0-255></pre>	Configures the sysLocation information for the system. <WORD 0-255> is an ASCII string from 0–255 characters.
<pre>login-success-trap {enable}</pre>	Enables generation of SNMP server login-success traps.
<pre>name WORD<0-255></pre>	Configures the sysName information for the system. <WORD 0-255> is an ASCII string from 0–255 characters. SNMP system name is derived from the global system name that is synchronized with SNMP system name, until the SNMP system name is manually configured.
<pre>notify-filter WORD<1-32> WORD<1-32></pre>	Creates a new entry in the notify filter table. The first WORD<1-32> specifies the filter profile name, and the second WORD<1-32> specifies the subtree object identifier (OID).

Creating a user

Create a new user in the USM table to authorize a user on a particular SNMP engine.

Procedure

- Enter Global Configuration mode:

```
enable

configure terminal
```
- Create a user on a remote system:

```
snmp-server user engine-id WORD<16-97> WORD<1-32> {md5 | sha }[aes |
des ]
```
- Enter and confirm your password.

4. Create a user on the local system:

```
snmp-server user WORD<1-32> [notify-view WORD<0-32>] [read-view WORD<0-32>] [write-view WORD<0-32>] {md5 | sha} [aes | des ]
```

5. Enter and confirm your password.

6. Add the user to a group:

```
snmp-server user WORD<1-32> group WORD<1-32> {md5 | sha} [aes | des ]
```

7. Enter and confirm your password.

8. Verify the configuration:

```
show snmp-server user
```

Example

Create a user named test1 on a remote system with MD5:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server user engine-id 0x80:00:08:E0:03:10:CD:AE:6B:D0:00 test1 md5
aes
Enter the authentication protocol password : *****
Re-enter the authentication protocol password : *****
Enter the privacy protocol password : *****
Re-enter the privacy protocol password : *****

WARNING: For best security practices avoid the use
         of repeated patterns in passwords.

Switch:1(config)#show snmp-server user
*****
Engine ID = 0x80:00:08:E0:03:10:CD:AE:6B:D0:00

=====
                        USM Configuration
=====
User/Security Name      Engine Id                Protocol
-----
User2                   0x80:00:08:E0:03:10:CD:AE:6B:D0:00 HMAC_MD5, AES PRIVACY,
test1                   0x80:00:08:E0:03:10:CD:AE:6B:D0:00 HMAC_MD5, AES PRIVACY,

2 out of 2 Total entries displayed
=====
```

Variable Definitions

The following table defines parameters for the `snmp-server user` command.

Variable	Value
<code>{aes des}</code>	Specifies a privacy protocol. If no value is entered, no authentication capability exists. The choices are aes or des. Important: You must set authentication before you can set the privacy option.
<code>engine-id WORD<16-97></code>	Assigns an SNMPv3 engine ID.

Variable	Value
<code>group WORD<1-32></code>	Specifies the group access name.
<code>{md5 sha}</code>	Specifies an authentication protocol. If no value is entered, no authentication capability exists. The protocol choices are: MD5 and SHA.
<code>notify-view WORD<0-32></code>	Specifies the view name. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view.
<code>read-view WORD<0-32></code>	Specifies the view name. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view.
<code>write-view WORD<0-32></code>	Specifies the view name. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view.
<code>user WORD<1-32></code>	Creates the new entry with this security name. The name is used as an index to the table.

Create a New User Group

Create a new user group to logically group users who require the same level of access. Create new access for a group in the View-based Access Control Model (VACM) table to provide access to managed objects.



Note

There are several default groups (public and private) created that you can use. To see the list of default groups and their associated security names (secnames), enter **show snmp-server group**. If you use one of these groups, there is no need to create a new group.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a new user group:

```
snmp-server group WORD <1-32> WORD<1-32> {auth-no-priv|auth-priv|no-  
auth-no-priv} [notify-view WORD<1-32>] [read-view WORD<1-32>] [write-  
view WORD<1-32>]
```

Examples

The following example uses the following variable names:

- The new group name is `lan6grp`.
- The context of the group is `"`, which represents the Global Router (VRF 0).

- The security level is *no-auth-no-priv*.
- The access view name is *v1v2only* for all three views: **notify-view**, **read-view**, and **write-view**.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server group lan6grp "" no-auth-no-priv notify-view v1v2only read-view v1v2only write-view v1v2only
```

Variable Definitions

The following table defines parameters for the `snmp-server group` command.

Variable	Value
<i>auth-no-priv</i>	Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the <i>auth-no-priv</i> parameter is included, it creates one entry for SNMPv3 access.
<i>auth-priv</i>	Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the <i>auth-priv</i> parameter is included, it creates one entry for SNMPv3 access.
<i>group</i> WORD<1-32> WORD<1-32>	The first WORD<1-32> specifies the group name for data access. The range is 1-32 characters. The second WORD<1-32> specifies the context name. The range is 1-32 characters. If you use a particular group name value but with different context names, you create multiple entries for different contexts for the same group. You can omit the context name and use the default. If the context name value ends in the wildcard character (*), the resulting entries match a context name that begins with that context. For example, a context name value of foo* matches contexts starting with foo, such as foo6 and foofofum.
<i>no-auth-no-priv</i>	Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the <i>no-auth-no-priv</i> parameter is included, it creates 3 entries, one for SNMPv1 access, one for SNMPv2c access, and one for SNMPv3 access.
<i>notify-view</i> WORD<1-32>	Specifies the view name.
<i>read-view</i> WORD<1-32>	Specifies the view name.
<i>write-view</i> WORD<1-32>	Specifies the view name.

Creating a new entry for the MIB in the view table

Create a new entry in the MIB view table. The default Layer 2 MIB view cannot modify SNMP settings. However, a new MIB view created with Layer 2 permission can modify SNMP settings.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Create a new entry:
`snmp-server view WORD<1-32> WORD<1-32>`

Example

```
Switch:1>enable
Switch:1#configure terminal
```

Create MIB views:

```
Switch:1(config)#snmp-server view 2 1.3.8.7.1.4
```

Variable Definitions

The following table defines parameters for the `snmp-server view` command.

Variable	Value
The first <i>WORD<1-32></i>	Specifies the prefix that defines the set of MIB objects accessible by this SNMP entity. The range is 1–32 characters.
The second <i>WORD<1-32></i>	Specifies a new entry with this group name. The range is 1–32 characters.

Creating a community

Create a community to use in forming a relationship between an SNMP agent and one or more SNMP managers. You require SNMP community strings to access the system using an SNMP-based management software.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`

2. Create a community:

```
snmp-server community WORD<1-32> [group WORD<1-32>] [index WORD<1-32>]
[secname WORD<1-32>]
```

**Important**

- The **group** parameter is only required if you created a new user group using the procedure in [Creating a new user group](#). If you use any of the default groups, the **secname** automatically links the community to its associated group so there is no need specify the group in this command.
- If you do create a new group, use the **snmp-server community** command to create an SNMP community with a new security name and link it to the new group you created. There is no separate command to create a security name (secname). You use the **snmp-server community** command. The security name is the key to link the community name to a group.
- You cannot use the @ character or the string :: when you create community strings.

Example

In the following example, the community name is *anewcommunity*, the index is *third*, and the secname is *readview*. There is no group specified because this is a default public/read only group.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server community anewcommunity index third secname readview
```

Variable Definitions

The following table defines parameters for the `snmp-server community` command.

Variable	Value
<i>community</i> <i>WORD<1-32></i>	Specifies a community string. The range is 1–32 characters.
<i>group</i> <i>WORD<1-32></i>	Specifies the group name. The range is 1–32 characters.
<i>index</i> <i>WORD<1-32></i>	Specifies the unique index value of a row in this table. The range is 1–32 characters.
<i>secname</i> <i>WORD<1-32></i>	Maps the community string to the security name in the VACM Group Member Table. The range is 1–32 characters.

Add a User to a Group

Add a user to a group to logically group users who require the same level of access.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Create a new user group:
snmp-server user *WORD*<1-32> group *WORD*<1-32> {md5 | sha} [aes | des]
3. Enter and confirm your password.
4. Verify the configuration:
show snmp-server group

Example

Add a user to a group to logically group users who require the same level of access:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#snmp-server user test4 group grouptest4 md5 aes
Enter the authentication protocol password : *****
Re-enter the authentication protocol password : *****
Enter the privacy protocol password : *****
Re-enter the privacy protocol password : *****
```

```
WARNING: For best security practices avoid the use
         of repeated patterns in passwords.
```

```
Switch:1(config)#show snmp-server group
```

```
*****
```

VACM Group Membership Configuration

Sec Model	Security Name	Group Name
snmpv1	readview	readgrp
snmpv1	initialview	vlv2grp
snmpv2c	readview	readgrp
snmpv2c	initialview	vlv2grp
usm	test1	Grouptest1
usm	test2	geet1
usm	test4	grouptest4

```
7 out of 7 Total entries displayed
```

VACM Group Access Configuration

Group	Prefix Model	Level	ReadV	WriteV	NotifyV
initial	usm	noAuthNoPriv	root	root	root
initial	usm	authPriv	root	root	root
initial	vrf512 usm	noAuthNoPriv	vrf	vrf	vrf
initial	vrf512 usm	authPriv	vrf	vrf	vrf

```

readgrp          snmpv1  noAuthNoPriv vlv2only          org
readgrp          snmpv2c noAuthNoPriv vlv2only          org
readgrp  vrf512  snmpv1  noAuthNoPriv vrf              vrf
readgrp  vrf512  snmpv2c noAuthNoPriv vrf              vrf
v1v2grp          snmpv1  noAuthNoPriv vlv2only  vlv2only  vlv2only
v1v2grp          snmpv2c noAuthNoPriv vlv2only  vlv2only  vlv2only
v1v2grp  vrf512  snmpv1  noAuthNoPriv vrf          vrf      vrf
v1v2grp  vrf512  snmpv2c noAuthNoPriv vrf          vrf      vrf

12 out of 12 Total entries displayed
-----

```

Variable Definitions

The following table defines parameters for the `snmp-server user` command.

Variable	Value
<code>{aes des}</code>	Specifies a privacy protocol. If no value is entered, no authentication capability exists. The choices are aes or des. Important: You must set authentication before you can set the privacy option.
<code>engine-id WORD<16-97></code>	Assigns an SNMPv3 engine ID.
<code>group WORD<1-32></code>	Specifies the group access name.
<code>{md5 sha}</code>	Specifies an authentication protocol. If no value is entered, no authentication capability exists. The protocol choices are: MD5 and SHA.
<code>notify-view WORD<0-32></code>	Specifies the view name. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view.
<code>read-view WORD<0-32></code>	Specifies the view name. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view.
<code>write-view WORD<0-32></code>	Specifies the view name. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view.
<code>user WORD<1-32></code>	Creates the new entry with this security name. The name is used as an index to the table.

Block SNMP Access

Disable SNMP by using the SNMP block flag. By default, SNMP access is enabled.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable SNMP access:

```
boot config flags block-snmp
```

Variable Definitions

The following table defines parameters for the `boot config flags` command.

Variable	Value
<i>block-snmp</i>	Configures the block SNMP flag as active. The default is off.

View SNMP System Information

View SNMP system information to view trap and authentication profiles.

For a comprehensive set of SNMP-related **show** commands, see *Fabric Engine Command References*.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View SNMP system information:

```
show snmp-server
```

Example

```
Switch:1>show snmp-server

      contact : http://www.extremenetworks.com/contact/
      location :
        name : Switch-BEB
AuthenticationTrap : false
LoginSuccessTrap : false
```

SNMP Configuration Using EDM

Configure SNMP to provide services to send and receive messages, authenticate and encrypt messages, and control access to managed objects with EDM.

Create a User

About This Task

Create a new user in the USM table to authorize a user on a particular SNMP engine.



Note

In EDM, to create new SNMPv3 users you must use the **CloneFromUser** option. However, you cannot clone the default user, named initial. As a result, you must first use CLI to configure at least one user, and then you can use EDM to create subsequent users with the **CloneFromUser** option.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **USM Table**.
3. Select **Insert**.
4. In **EngineID**, use the default Engine ID provided or type an administratively-unique identifier to an SNMP engine.
5. In **User Name**, type a name.
6. From **CloneFromUser**, select a security name from which the new entry copies authentication data and private data, if required.
7. From **Auth Protocol**, select an authentication protocol.
8. In **Cloned User's Auth Password**, type the authentication password of the cloned user.
9. In **New User's Auth Password**, type an authentication password for the new user.
10. From **Priv Protocol**, select a privacy protocol.
11. In **Cloned User's Priv Password**, type the privacy password of the cloned user.
12. In **New User's Priv Password**, type a privacy password for the new user.
13. Select **Insert**.



Caution

To ensure security, change the GroupAccess table default view after you set up a new user in the USM table. This prevents unauthorized people from accessing the system using the default user logon. Also, change the Community table defaults, because the community name is used as a community string in SNMPv1/v2 PDU.

USM Table Field Descriptions

Use the data in the following table to use the **USM Table** tab and the **Insert USM Table** dialog box. The system displays some fields only on the Insert USM Table dialog box.

Name	Description
EngineID	Specifies an administratively-unique identifier to an SNMP engine.
UserName	Creates the new entry with this security name. The name is used as an index to the table. The range is 1–32 characters.

Name	Description
SecurityName	Identifies the name on whose behalf SNMP messages are generated.
Clone From User	Specifies the security name from which the new entry must copy privacy and authentication parameters. The range is 1–32 characters. The system displays this option only in the Insert USM Table dialog box.
Auth Protocol (Optional)	Assigns an authentication protocol (or no authentication) from a list. If you select an authentication protocol, you must enter an old AuthPass and a new AuthPass.
Cloned User's Auth Password	Specifies the current authentication password of the cloned user. The system displays this option only in the Insert USM Table dialog box.
New User's Auth Password	Specifies the authentication password of the new user. The system displays this option only in the Insert USM Table dialog box.
Priv Protocol (Optional)	Assigns a privacy protocol (or no privacy) from a list. If you select a privacy protocol, you must enter an old PrivPass and a new PrivPass.
Cloned User's Priv Password	Specifies the current privacy password of the cloned user. The system displays this option only in the Insert USM Table dialog box.
New User's Priv Password	Specifies the privacy password of the new user. The system displays this option only in the Insert USM Table dialog box.

Create a New Group Membership

About This Task

Create a new group membership to logically group users who require the same level of access.



Note

There are several default groups (public and private) created that you can use. To see the list of default groups and their associated security names (secnames), enter **show snmp-server group**. If you use one of these groups, there is no need to create a new group.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **VACM Table**.
3. Select the **Group Membership** tab.
4. Select **Insert**.
5. From the **SecurityModel** options, select a security model.
6. In **SecurityName**, type a security name.
7. In **GroupName**, type a group name.

8. Select **Insert**.

Group Membership *Field Descriptions*

Use the data in the following table to use the **Group Membership** tab.

Name	Description
SecurityModel	Specifies the security model to use with this group membership.
SecurityName	Specifies the security name assigned to this entry in the View-based Access Control Model (VACM) table. The range is 1–32 characters.
GroupName	Specifies the name assigned to this group in the VACM table. The range is 1–32 characters.

Create Access for a Group

About This Task

Create access for a group in the View-based Access Control Model (VACM) table to provide access to managed objects.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **VACM Table**.
3. Select the **Group Access Right** tab.
4. Select **Insert**.
5. In **GroupName**, type a VACM group name.
6. (Optional) In **ContextPrefix**, select a VRF instance.
7. From the **SecurityModel** options, select a model.
8. From the **SecurityLevel** options, select a security level.
9. In **ContextMatch**, select a value to match the context name.
10. (Optional) In **ReadViewName**, type the name of the MIB view that forms the basis of authorization when reading objects.
11. (Optional) In **WriteViewName**, type the name of the MIB view that forms the basis of authorization when writing objects.
12. (Optional) In **NotifyViewName**, type MIB view that forms the basis of authorization for notifications.
13. Select **Insert**.

Group Access Right *Field Descriptions*

Use the data in the following table to use the **Group Access Right** tab.

Name	Description
GroupName	Specifies the name of the new group in the VACM table. The range is 1–32 characters.
ContextPrefix	Specifies if the contextName must match the value of the instance of this object exactly or partially. The range is an SnmpAdminString, 1–32 characters.
SecurityModel	Specifies the authentication checking to communicate to the switch. The security models are: <ul style="list-style-type: none"> • SNMPv1 • SNMPv2 • USM
SecurityLevel	Specifies the minimum level of security required to gain the access rights allowed. The security levels are: <ul style="list-style-type: none"> • noAuthNoPriv • authNoPriv • authpriv
ContextMatch	Specifies if the prefix and the context name must match. If the value is exact, all rows where the contextName exactly matches vacmAccessContextPrefix are selected. If you do not select exact, all rows where the contextName with starting octets that exactly match vacmAccessContextPrefix are selected. This default is exact .
ReadViewName	Identifies the MIB view of the SNMP context to which this conceptual row authorizes read access. The default is the empty string.
WriteViewName	Identifies the MIB view of the SNMP context to which this conceptual row authorizes write access. The default is the empty string.
NotifyViewName	Identifies the MIB view of the SNMP context to which this conceptual row authorizes access for notifications. The default is the empty string.

Create Access Policies for SNMP Groups

Create an access policy to determine the access level for the users who connect to the switch with different services like File Transfer Protocol (FTP), Trivial FTP (TFTP), and Telnet.

About This Task

You only need to create access policies for SNMP groups if you have the access policy feature enabled.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **Access Policies**.
3. Select the **Access Policies-SNMP Groups** tab.
4. Select **Insert**.
5. Enter an **ID**.
6. In the **Name** box, type a name.
7. From the **Model** options, select a security model.
8. Select **Insert**.

Access Policies — SNMP Groups *Field Descriptions*

The following table defines parameters for the **Access Policies-SNMP Groups** tab.

Name	Description
Id	Specifies the ID of the group policy.
Name	Specifies the name assigned to the group policy. The range is 1–32 characters.
Model	Specifies the security model {SNMPv1 SNMPv2c USM}.

Assign MIB View Access for an Object

About This Task

Create a new entry in the MIB View table.

You cannot modify SNMP settings with the default Layer 2 MIB view. However, you can modify SNMP settings with a new MIB view created with Layer 2 permissions.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **VACM Table**.
3. Select the **MIB View** tab.
4. Select **Insert**.
5. In **ViewName**, type a view name.
6. In **Subtree**, type a subtree.
7. In **Mask**, type a mask.
8. From the **Type** options, select whether access to the MIB object is granted.
9. Select **Insert**.

MIB View *Field Descriptions*

Use the data in the following table to use the **MIB View** tab.

Name	Description
ViewName	Creates a new entry with this group name. The range is 1–32 characters.
Subtree	Specifies a valid object identifier that defines the set of MIB objects accessible by this SNMP entity, for example, 1.3.6.1.1.5.
Mask (optional)	Specifies a bit mask with vacmViewTreeFamilySubtree to determine whether an OID falls under a view subtree.
Type	Determines whether access to a MIB object is granted (included) or denied (excluded). The default is included.

Create a Community

About This Task

Create a community to use in forming a relationship between an SNMP agent and one or more SNMP managers. You require SNMP community strings for access to the switch using an SNMP-based management software.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **Community Table**.
3. Select **Insert**.
4. In **Index**, type an index.
5. In **Name**, type a name that is a community string.
6. In **SecurityName**, type a security name.
7. In **ContextName**, type the context name.
8. Select **Insert**.

Community Table *Field Descriptions*

Use the data in the following table to use the **Community Table** tab.

Name	Description
Index	Specifies the unique index value of a row in this table. The range is 1–32 characters.
Name	Specifies the community string for which a row in this table represents a configuration.
SecurityName	Specifies the security name in the VACM group member table to which the community string is mapped. The range is 1–32 characters.

Name	Description
ContextEngineID	Indicates the location of the context in which management information is accessed when using the community string specified in Name .
ContextName	Specifies the context in which management information is accessed when you use the specified community string.

View All Contexts for an SNMP Entity

About This Task

View contexts to see the contents of the context table in the View-based Access Control Model (VACM). This table provides information to SNMP command generator applications so that they can properly configure the VACM access table to control access to all contexts at the SNMP entity.

Procedure

1. In the navigation pane, expand **Configuration > Edit > SnmpV3**.
2. Select **VACM Table**.
3. Select the **Contexts** tab.

Contexts Field Descriptions

Use the data in the following table to use the **Contexts** tab.

Variable	Value
ContextName	Shows the name identifying a particular context at a particular SNMP entity. The empty contextName (zero length) represents the default context.

Show SNMP Statistics

About This Task

Display SNMP statistics to monitor the number of specific error messages, such as the number of messages that were delivered to SNMP but were not allowed.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **General**.
3. Select the **SNMP** tab.

SNMP Field Descriptions

Use the data in the following table to display **SNMP** statistics.

Name	Description
OutTooBigs	Shows the number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status field is tooBig.
OutNoSuchNames	Shows the number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status is noSuchName.
OutBadValues	Shows the number of SNMP PDUs that SNMP protocol entity generated and for which the value of the error-status field is badValue.
OutGenErrors	Shows the number of SNMP PDUs that the SNMP protocol entity generated and for which the value of the error-status field is genErr.
InBadVersions	Shows the number of SNMP messages that were delivered to the SNMP protocol entity and were for an unsupported SNMP version.
InBadCommunityNames	Shows the number of SNMP messages delivered to the SNMP protocol entity that used an SNMP community name not known to the entity.
InBadCommunityUsers	Shows the number of SNMP messages delivered to the SNMP protocol entity that represented an SNMP operation not allowed by the SNMP community named in the message.
InASNParseErrs	Shows the number of ASN.1 or BER errors encountered by the SNMP protocol entity when decoding received SNMP messages.
InTooBigs	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is tooBig.
InNoSuchNames	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is noSuchName.
InBadValues	Shows the number of SNMP PDUs that were delivered to the SNMP protocol entity and for which the value of the error-status field is badValue.
InReadOnlys	Shows the number of valid SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is "read-only". It is a protocol error to generate an SNMP PDU that contains the value "read-only" in the error-status field; this object is provided as a means of detecting incorrect implementations of the SNMP.
InGenErrors	Shows the number of SNMP PDUs delivered to the SNMP protocol entity and for which the value of the error-status field is "genErr."



Spanning Tree Configuration

[Spanning Tree Fundamentals](#) on page 3496

[Spanning Tree Configuration Using CLI](#) on page 3506

[Spanning Tree Configuration Using EDM](#) on page 3532

This section provides conceptual and procedural information about Spanning Tree.

Spanning Tree Fundamentals

Table 303: Spanning Tree Protocol product support

Feature	Product	Release introduced
Spanning Tree Protocol (STP): <ul style="list-style-type: none">• Multiple STP (MSTP)• Rapid STP (RSTP)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3 Note: 7830 Series does not support RSTP.
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 303: Spanning Tree Protocol product support (continued)

Feature	Product	Release introduced
Multiple Spanning Tree Protocol (MSTP) Restricted Role (Root Guard) and Restricted Topology Change Notification (TCN)	4220 Series	Fabric Engine 9.4
	5320 Series	Fabric Engine 9.4
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.4
	5720 Series	Fabric Engine 9.4
	7520 Series	Fabric Engine 9.4
	7720 Series	Fabric Engine 9.4
	7830 Series	Fabric Engine 9.4
	VSP 4900	VOSS 9.4
	VSP 7400	VOSS 9.4

The switch supports Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP).

Spanning Tree

Spanning Tree protocols detect and eliminate logical loops in a bridged or switched network. If multiple paths exist, the spanning Tree algorithm configures the network so that a bridge or device uses the root bridge path based on hop counts. Although link speed is taken into account, the path is based on the root bridge rather than on an optimized path. If that path fails, the protocol automatically reconfigures the network and makes another path active, thereby sustaining network operations.

The switch supports RSTP and MSTP but can downgrade a port automatically if it receives an STP Bridge Protocol Data Unit (BPDU) from a switch that runs STP. While in MSTP mode, you can configure the switch to temporarily operate in RSTP mode and send RSTP BPDUs on a specific port. This RSTP mode remains active until the port receives an MSTP BPDU, at which point it reverts automatically to MSTP mode.



Note

Spanning Tree is disabled on all Switched UNI (S-UNI) ports. The ports move into forwarding state as soon as the physical port or VLACP or LACP comes up on the port. If the platform VLAN is associated to the S-UNI Service Instance Identifier (I-SID), then the S-UNI ports added to the platform VLAN become the member of MSTP instances associated with the platform VLAN. To enable SLPP on the S-UNI ports, the platform VLAN must be associated with the S-UNI I-SID.

Spanning Tree Groups

Spanning Tree Groups (STGs) represent logical topologies. A topology is created based on bridge configuration values such as root bridge priority. In the case of multiple STGs, you can map a VLAN to the most appropriate logical topology in the physical network.

The switch supports Spanning Tree modes RSTP and MSTP. The default Spanning Tree mode is MSTP. The default STG is 0. In RSTP mode, all VLANs run in the default STG. In MSTP mode, you can create additional STGs by using the VLAN create command. The switch supports up to 64 STGs.

Although STP and MSTP are variations of the same Spanning Tree protocol, they communicate information differently. A switch in MSTP mode cannot recognize the Spanning Tree groups running on a chassis configured with STP. MSTP Spanning Tree groups are not the same as STP Spanning Tree groups. Using a switch in MSTP mode with a chassis in STP mode can create a loop in the network.

The root bridge for Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP) is determined by comparing attributes of each bridge in the network.

The protocol considers bridge priority first. If more than one bridge has the same priority, then the protocol must consider the bridge ID. The bridge with the lowest ID becomes the root bridge. For MSTP, this bridge is called the Common and Internal Spanning Tree (CIST) Root because it is the root of the entire physical network.

In MSTP mode, you can create additional Spanning Tree instances, by using the VLAN command. These instances, known as Multiple Spanning Tree Instances (MSTIs), can assign different priorities to switches. The MSTIs have different link costs or port priorities and as a result create separate logical topologies.

MSTP also allows the creation of MSTP regions. A region is a collection of switches sharing the same view of physical and logical topologies. For switches to belong to the same region, the following attributes must match:

- MSTP configuration ID selector
- MSTP configuration name
- MSTP configuration revision number
- VLAN instance mapping

Links connecting sections are called boundary ports. In a region, the boundary switch that contains the boundary port providing the shortest external path cost to the CIST Root is the CIST Regional Root.

STGs and VLANs

When you map VLANs to STGs, be aware that all links on the bridge belong to all STGs. Because each Spanning Tree group can differ in its decision to make a link forwarding or blocking, you must ensure that the ports you add to a VLAN are in the expected state.

Untagged ports can only belong to one VLAN and therefore can only belong to one STG. Tagged ports can belong to multiple VLANs and therefore to multiple STGs.

BPDU handling on S-UNI port/MLT

The switch handles Bridge Protocol Data Units (BPDUs) according to whether or not you configure a platform VLAN.

- When you configure a platform VLAN:
 - BPDUs are forwarded to the CPU by default.
 - For both the ingress and egress ports, BPDUs are not flooded in the S-UNI I-SID associated with the platform VLAN.



Note

If the platform VLAN is configured for the S-UNI port, you cannot enable BPDU forwarding.

- When you DO NOT configure a platform VLAN:
 - BPDUs received on untagged-traffic ports are dropped by default.
 - To flood BPDUs in its I-SID, enable BPDU forwarding under S-UNI I-SID using the command **untagged-traffic port <port no> bpdu enable**.

BPDU Guard

Table 304: Bridge Protocol Data Unit (BPDU) Guard product support

Feature	Product	Release introduced
Bridge Protocol Data Unit (BPDU) Guard	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The switch supports Bridge Protocol Data Unit (BPDU) Guard for STGs, RSTP, and MSTP.

Overview

Spanning Tree eliminates loops in a network. A bridge that participates in spanning tree uses BPDUs to exchange information with other bridges. The bridges select a single bridge as the root bridge based on the BPDU information exchange. The bridge with the lowest priority becomes the root bridge. If all bridges share the same priority,

the bridge with the lowest bridge ID becomes the root bridge. This process is the root selection process.

After you add a new bridge to the network, or remove an existing bridge, the bridges repeat the root selection process, and then select a new root bridge.

To ensure the correct operation of Spanning Tree in the network, BPDU Guard protects the stability of the Root Bridge by dropping stray, unexpected, or unwanted BPDU packets entering a port, and immediately shutting down those ports for a specified time period. BPDU Guard is normally enabled on access ports connecting to end user devices such as servers that are not expected to operate Spanning Tree.

Use BPDU Guard to achieve the following results:

- Block the root selection process after an edge device, such as a laptop that uses Linux with STP enabled, is added to the network. Blocking the root selection process prevents unknown devices from influencing the spanning tree topology.
- Block BPDU flooding of the switch from an unknown device.

Operation

You can enable or disable BPDU Guard on an individual port basis, regardless of the spanning tree state. Each port uses a timer to determine port-state recovery.

After you enable BPDU Guard on a port and the port receives a BPDU, the following actions occur:

1. The guard disables the port.
2. The switch generates an SNMP trap and alarm, and the following log message:

```
BPDU Guard - Port <slot/port> is being shutdown by BPDU Guard, timeout  
<time_seconds>
```

3. The port timer begins.
4. The port remains in the disabled state until the timer expires.

If you disable BPDU Guard before the timer expires, the timer stops and the port remains in the disabled state. You must manually enable the port.

BPDU Guard is enabled at the interface level. You can configure the BPDU Guard timer for each port. If you set the port timer to zero, it does not expire.

Rapid Spanning Tree Protocol and Multiple Spanning Tree Protocol

The Rapid Spanning Tree Protocol (RSTP or IEEE 802.1w) reduces the recovery time after a network breakdown. It also maintains backward compatibility with IEEE 802.1d (the spanning tree implementation prior to RSTP). In certain configurations, the recovery time of RSTP can be reduced to less than 1 second. RSTP also reduces the amount of flooding in the network by enhancing the way Topology Change Notification (TCN) packets are generated.

With Multiple Spanning Tree Protocol (MSTP or IEEE 802.1s), you can configure multiple instances or Spanning Tree groups on the same device. Each instance or Spanning Tree group can include one or more VLANs.

By using RSTP and MSTP, the switch achieves the following:

- reduces convergence time after a topology change (from 30 seconds to less than 1 or 2 seconds)
- eliminates unnecessary flushing of the MAC database and the flooding of traffic to the network
- creates backward compatibility with classic 802.1d switches
- creates support for 64 instances of spanning tree in MSTP mode

RSTP interoperability with STP

RSTP provides a parameter called ForceVersion to provide backward compatibility with standard STP. A user can configure a port in either STP-compatible mode or RSTP mode:

- An STP-compatible port transmits and receives only STP Bridge Protocol Data Units (BPDUs). An RSTP BPDU that the port receives in this mode is discarded.
- An RSTP-compatible port transmits and receives only RSTP BPDUs. If an RSTP port receives an STP BPDU, it becomes an STP port. User intervention is required to change this port back to RSTP mode. This process is called Port Protocol Migration.



Note

You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.

Before implement MSTP or RSTP you must be aware of the following:

- The default mode is MSTP. A special boot configuration flag identifies the mode.
- You can lose your configuration if you change the spanning tree mode from MSTP to RSTP and the configuration file contains VLANs configured with MSTI greater than 0. RSTP only supports VLANs configured with the default instance 0.
- For best interoperability results, contact your vendor representative.

Differences in port roles for STP and RSTP

RSTP is an enhanced version of STP. These two protocols have almost the same parameters.

The following table lists the differences in port roles for STP and RSTP. STP supports two port roles, while RSTP supports four port roles.

Table 305: Differences in port roles for STP and RSTP

Port Role	STP	RSTP	Description
Root	Yes	Yes	This port receives a better BPDU than its own and has the best path to reach the Root. The root port is in Forwarding state. The root port and designated ports can be in the Discarding state before they go to root forwarding.
Designated	Yes	Yes	This port has the best BPDU on the segment. The designated port is in the Forwarding state.
Alternate	No	Yes	This port receives a better BPDU than its own BPDU, and a root port exists within the same device. The alternate port is in the Discarding state.
Backup	No	Yes	This port receives a better BPDU than its own BPDU, and this BPDU is from another port within the same device. The backup port is in the Discarding state.

Port roles: root forwarding role

MSTP and RSTP root forwarding roles are as follows:

- The port that receives the best path BPDU on a device is the root port, and is referred to as a Root Forwarding (RF) port. This is the port that is the closest to the root bridge in terms of path cost.
- The spanning tree algorithm elects a single root bridge in a bridged network. With MSTP, a root bridge is selected for the Common and Internal Spanning Tree (CIST). A root bridge is selected for the region, and a root bridge is selected for each spanning tree instance.
- The root bridge is the only bridge in a network that does not have root ports; all ports on a root bridge are Designated Forwarding (DF).
- Only one path towards a root bridge can exist on a given segment; otherwise, loops can occur.

Port roles: designated forwarding role

MSTP and RSTP designated forwarding roles are as follows:

- All bridges connected on a segment monitor the BPDUs of all other bridges. The bridge that sends the best BPDU is the root bridge for the segment.
- The corresponding port on the bridge is referred to as a Designated Forwarding Port.

Port roles: alternate blocking role

MSTP and RSTP alternate blocking roles are as follows:

- A blocked port is defined as not being the designated or root port. An alternate port provides an alternate path to the root and can replace the root port if it fails.
- An alternate blocked port is a port that is blocked because it received better path cost BPDUs from another bridge.

Port roles: backup blocking role

MSTP and RSTP backup blocking roles are as follows:

- A backup port receives the more useful BPDUs from the bridge on which the port exists.

Edge port

RSTP uses a parameter called the edge port. After a port connects to a nonswitch device, such as a PC or a workstation, it must be configured as an edge port. An active edge port enters the forwarding state without delay. An edge port becomes a nonedge port if it receives a BPDU.

Path cost values

RSTP and MSTP path cost values support a wide range of link speeds. The following table lists the path cost values.

Table 306: Path cost values

Link speed	Value
Less than or equal to 100 Kbps	200 000 000
1 Mbps	20 000 000
10 Mbps	2 000 000
100 Mbps	200 000
1 Gbps	20 000
10 Gbps	2000
100 Gbps	200
1 Tbps	20
10 Tbps	2

RSTP negotiation process

The following section describes the negotiation process between switches that takes place before PCs can exchange data (see the following figure).




Figure 224: RSTP negotiation process

After turning on, all ports assume the role of designated ports. All ports are in the discarding state except edge ports. Edge ports go directly into the forwarding state without delay.

Unit A port 1 and Unit B port 1 exchange BPDUs. Unit A knows that it is the root and that Unit A port 1 is the designated port. Unit B learns that Unit A has higher priority. Unit B port 1 becomes the root port. Both Unit A port 1 and Unit B port 1 are still in the discarding state.

Unit A starts the negotiation process by sending a BPDU with the proposal bit set.

Unit B receives the proposal BPDU and configures its nonedge ports to discarding state. This operation occurs during the synchronization process.

Unit B sends a BPDU to Unit A with the agreement bit set.

Unit A configures port 1 to the forwarding state, and Unit B configures port 1 to the forwarding state. PC 1 and PC 2 can now communicate. The negotiation process now moves on to Unit B port 3 and its partner port. PC 3 cannot exchange data with either PC 1 or PC 2 until the negotiation process between Unit B and Unit C finishes.

The RSTP convergence time depends on how quickly the switch can exchange BPDUs during the negotiation process, and on the number of switches in the network.

NEW Multiple Spanning Tree Protocol Restricted Role and Restricted TCN

Multiple Spanning Tree Protocol (MSTP) Restricted Role (Root Guard) prevents a port from accepting superior BPDUs from non-root bridges. When triggered, it puts the port in a root-inconsistent state. Restricted Topology Change Notification (TCN) prevents a switch port from propagating topology-change messages to other ports. When enabled, the port blocks and ignores all received TCNs.

NEW! Examples

The following examples use a simplified topology with two Backbone Edge Bridges (BEB-1 and BEB-2) and, where noted, a remote BEB-3 to describe how MSTP Restricted Role (Root Guard) and Restricted Topology Change Notification (TCN) operate in common deployment scenarios.

MSTP Restricted Role (Root Guard)

BEB-1 and BEB-2 provide an uplink toward the core and run MSTP Restricted Role on those uplink ports. With Restricted Role active, each BEB refuses any superior BPDU that attempts to advertise a better path to the root bridge. Each BEB treats itself as the root bridge and ignores the other device in spanning tree processing, which causes one uplink to block. This behavior removes the need for vIST between the BEBs and avoids the requirement for SPB STP-Multihoming when the VLAN operates as a routed Layer 2 VSN at the access layer. Restricted Role therefore delivers a straightforward method to maintain predictable spanning tree behavior when dual edge devices connect to a shared Layer 2 segment.

MSTP Restricted Role (Root Guard) Protection Against External Bridges

When an external bridge tries to become the root bridge—for example, by advertising a bridge priority of zero—BEB-1 and BEB-2 immediately trigger Restricted Role on the ports connected to that device. Restricted Role rejects the superior BPDUs that the external bridge sends, which prevents that device from influencing the spanning tree topology and helps the BEBs maintain a stable network. This behavior highlights the security advantage of running Restricted Role on ports that connect to untrusted or unmanaged Layer 2 equipment, because it stops external devices from overriding the intended root bridge position within the BEB network.

MSTP Restricted TCN Behavior

When BEB-1 and BEB-2 run Restricted TCN on the ports that connect to an external spanning tree network, they block any topology change events that originate from that external network. The BEBs drop incoming TCNs from the external bridges, which stops unnecessary MAC address table flushes and avoids unwanted reconvergence events within the BEBs and the rest of the SPB fabric when STP-Multihoming is in use. By enabling Restricted Role and Restricted TCN on the same ports, the BEBs isolate the core network and prevent both root bridge manipulation attempts and disruptive topology change signals from the external Layer 2 domain.

Uplink A and Uplink B Switchover with Remote BEB-3

In this scenario, BEB-1 and BEB-2 provide dual uplinks—uplink A (active) and uplink B (standby)—toward the core, and BEB-3 extends the same access VLANs across a Layer 2 VSN with STP-Multihoming. When uplink A on BEB-1 fails and uplink B unlocks, the fabric does not propagate that STP uplink change to BEB-3. BEB-3 continues to use the existing I-SID destination MAC for host H1—the B-MAC of BEB-1—so if H1 remains silent, BEB-3 cannot reach H1 during this interval. After BEB-3 ages out the MAC entry for H1, the fabric updates the destination MAC to reflect the new path, and H1 becomes reachable again. This behavior highlights the MAC-aging dependency during uplink

transitions when you extend access VLANs across a Layer 2 VSN to a remote BEB that uses STP-Multihoming.

Spanning Tree Configuration Using CLI



Important

The switch supports up to 64 STGs on a device, however, SPBM uses STG 63 and MSTI 62 for internal use. STG 63 or MSTI 62 cannot be used by other VLANs or MSTIs.

Configure the Spanning Tree Mode

About This Task

Specifies the Multiple Spanning Tree Protocol (MSTP) or Rapid Spanning Tree Protocol (RSTP) mode. If you do not specify a protocol, the switch uses the default mode, which is mstp. If you change the spanning tree mode, you must save the current configuration and restart the switch.



Important

Only MSTP mode is compatible with Fabric SPBM. If you configure RSTP mode, the switch cannot participate in an SPBM network.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the STP mode:
`boot config flags spanning-tree-mode {mstp | rstp}`

Example

Configure the STP mode:

```
Switch:1(config)#boot config flags spanning-tree-mode rstp
Warning: RSTP Spanning Tree mode is incompatible with Fabric SPBM
         and will result in the switch no longer being able to
         participate in an SPBM network.
Warning: Please save the configuration and reboot the switch
         for this to take effect.
Warning: Please carefully save your configuration files before
         starting configuring the switch in RSTP or MSTP mode.
```

Configure BPDU Guard

Configure BPDU Guard to block the root selection process or to prevent BPDU flooding from unknown devices.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable BPDU Guard for the port:

```
spanning-tree bpduguard enable
```

3. (Optional) Configure the timer for port-state recovery:

```
spanning-tree bpduguard timeout <0, 10-65535>
```

4. (Optional) Enable BPDU Guard on an additional port or group of ports:

```
spanning-tree bpduguard port {slot/port[-slot/port] [,...][slot/all]
[all]} enable
```

5. (Optional) Configure the timer for port-state recovery for an additional port or group of ports:

```
spanning-tree bpduguard port {slot/port[-slot/port] [,...][slot/all]
[all]} timeout <0-65535>
```

6. Verify the configuration:

```
show spanning-tree bpduguard [GigabitEthernet {slot/port[-slot/port]
[,...][slot/all][all]}] [{slot/port[-slot/port] [,...][slot/all][all]}]
```

Example

Enable BPDU Guard on port 1/8, and specify a timer value of 200 seconds. Verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/8
Switch:1(config-if)#spanning-tree bpduguard enable
Switch:1(config-if)#spanning-tree bpduguard timeout 200
Switch:1(config-if)#show spanning-tree bpduguard 1/8
```

Bpdu Guard							
Port	PORT	PORT		TIMER	BPDUGUARD	BPDUGUARD	
NUM	MLTID	ADMIN_STATE	OPER_STATE	TIMEOUT	COUNT	ADMIN_STATE	ORIGIN
1/8	Up	Up	200	0	Enabled	CONFIG	

Variable Definitions

Use the data in the following table to use the **spanning-tree bpduguard** commands.

Variable	Value
<code>enable</code>	Enables BPDU Guard on the port. The default is disabled.
<code>port {slot/ port[-slot/port] [,...][slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all)
<code>timeout <0, 10-65535></code>	Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires. You can configure a value from 10 to 65535. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.

Use the data in the following table to use the **show spanning-tree bpduguard** command.

Variable	Value
<code>{slot/port[/ sub-port]}[- slot/port[/sub- port]][,...] [slot/all] [all]}</code>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Troubleshoot No BPDU Packets Received on the Port

For BPDU Guard to work on a port, the port must receive BPDU packets. Perform the following procedure to troubleshoot cases when the port does not receive packets.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Show the BPDU Guard status for the port:

```
show spanning-tree bpduguard {slot/port[-slot/port]][,...][slot/all]  
[all]}
```

3. Use the following command to verify that the port receives packets:

```
show interfaces gigabitEthernet statistics verbose {slot/port[-slot/  
port]][,...][slot/all][all]}
```

4. Verify that the remote port is sending packets:

```
show spanning-tree {mstp|rstp} port role [{slot/port[-slot/port]}[,...]  
[slot/all][all]]
```

```
show spanning-tree {mstp|rstp} port statistics [{slot/port[-slot/port]}  
[,...][slot/all][all]]
```

Example

Port 1/8 receives packets. The remote port is disabled and does not send BPDU packets.

The following example shows that BPDU Guard is enabled for port 1/8. The BPDU Guard administrative state for the port is enabled but the timer counter is 0.

```
Switch:1>enable
Switch:1#show spanning-tree bpduguard 1/8
=====
Bpdu Guard
=====
Port      PORT      PORT      TIMER  BPDUGUARD  BPDUGUARD
NUM MLTID  ADMIN_STATE  OPER_STATE  TIMEOUT  COUNT      ADMIN_STATE  ORIGIN
-----
1/8      Up        Up          120       0         Enabled     CONFIG
Switch:1#show interfaces gigabitEthernet statistics verbose 1/8
=====
Port Stats Interface Extended
=====
PORT_NUM  IN_UNICST  OUT_UNICST  IN_MULTICST  OUT_MULTICST  IN_BRDCST  OUT_BRDCST  IN_LSM  OUT_LSM
-----
1/8      201        0           160062       60943         4          72          0       0
Switch:1#show spanning-tree mstp port role 1/8
=====
CIST Port Roles and States
=====
Port-Index  Port-Role  Port-State  PortSTPStatus  PortOperStatus
-----
1/8        Disabled   Forwarding   Disabled        Disabled
Switch:1#show spanning-tree mstp port statistics 1/8
=====
MSTP Cist Port Statistics
=====
Port Number                : 1/8
Cist Port Fwd Transitions   : 0
Cist Port Rx MST BPDUs Count : 0
Cist Port Rx RST BPDUs Count : 0
Cist Port Rx Config BPDUs Count : 0
Cist Port Rx TCN BPDUs Count : 0
Cist Port Tx MST BPDUs Count : 0
Cist Port Tx RST BPDUs Count : 0
Cist Port Tx Config BPDUs Count : 0
Cist Port Tx TCN BPDUs Count : 0
Cist Port Invalid MSTP BPDUs Rx : 0
Cist Port Invalid RST BPDUs Rx : 0
Cist Port Invalid Config BPDUs Rx : 0
Cist Port Invalid TCN BPDUs Rx : 0
Cist Port Proto Migr Count   : 0
```

Configure Rapid Spanning Tree Protocol

Configure Rapid Spanning Tree Protocol (RSTP) to reduce the recovery time after a network breakdown.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure RSTP:

```
spanning-tree rstp [forward-time <400-3000>] [group-stp enable]
[hello-time <100-1000>] [max-age <600-4000>] [pathcost-type <bits16|
bits32>] [priority <0-61440>] [tx-holdcount <1-10>] [version <rstp|
stp-compatible>]
```

Example

Configure RSTP:

```
Switch:1(config)# spanning-tree rstp forward-time 1000 hello-time 200 max-age 4000
pathcost-type bits16 priority 4096 tx-holdcount 10 version rstp group-stp enable
```

Variable Definitions

Use the data in the following table to use the **spanning-tree rstp** command.

Variable	Value
<i>forward-time <400-3000></i>	Configures the RSTP forward delay for the bridge in hundredths of a second.
<i>group-stp enable</i>	Enables RSTP for a specific STG.
<i>hello-time <100-1000></i>	Assigns the RSTP hello time delay for the bridge in hundredths of a second.
<i>max-age <600-4000></i>	Assigns the RSTP maximum age time for the bridge in hundredths of a second.
<i>pathcost-type {bits16 bits32}</i>	Assigns the RSTP default pathcost version. The default is 32 bits.
<i>priority <0-61440></i>	Assigns the RSTP bridge priority.
<i>tx-holdcount <1-10></i>	Assigns the RSTP transmit hold count. The default value is 6.
<i>version {rstp/stp-compatible}</i>	Sets the version to RSTP or STP compatible.

Configure Rapid Spanning Tree Protocol for a Port

Configure RSTP to reduce the recovery time after a network breakdown.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure RSTP:

```
spanning-tree rstp cost <1-200000000> edge-port <false|true> p2p  
<auto|force-false|force-true> priority <0-240> protocol-migration  
<false|true> stp enable
```

Example

Configure RSTP:

```
Switch:1(config-if)#spanning-tree rstp cost 100 edge-port true p2p auto priority 32  
protocol-migration true stp enable
```

Variable Definitions

Use the data in the following table to use the **spanning-tree rstp** command.

Variable	Value
<i>cost</i> <1-200000000>	Specifies the contribution of this port to the path cost.
<i>edge-port</i> <false true>	Configures the edge-port value for the port. A value of true indicates that this port is an edge-port, and a value of false indicates that this port is a nonedge-port.
<i>p2p</i> <auto force-false force-true>	Specifies the point-to-point status of the LAN segment attached to this port. A value of force-true indicates that this port is treated as if it connects to a point-to-point link. A value of force-false indicates that this port is treated as having a shared media connection. A value of auto indicates that this port is considered to have a point-to-point link if it is an aggregator and all of its members are aggregatable, or if the MAC entity is configured for full-duplex operation, either through autonegotiation or by management means.

Variable	Value
<code>priority <0-240></code>	Assigns the RSTP bridge priority in a range of 0–240. The value has to increment in steps of 16.
<code>protocol-migration <false true></code>	<p>If you chose true, the option initiates protocol migration for a port. If you chose false, the option terminates protocol migration for a port.</p> <p>An RSTP-compatible port transmits and receives only RSTP BPDUs. If an RSTP port receives an STP BPDU, it becomes an STP port. User intervention is required to change this port back to RSTP mode. This process is called Port Protocol Migration.</p>
<code>stp enable</code>	Configures STP for the port.

Configure the Rapid Spanning Tree Protocol Version

Perform this procedure to specify the RSTP mode.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure Rapid Spanning Tree Protocol version:

```
spanning-tree rstp version {rstp|stp-compatible}
```

Example

Configure Rapid Spanning Tree Protocol version:

```
Switch:1(config)#spanning-tree rstp version rstp
```

Variable Definitions

Use the data in the following table to use the **spanning-tree rstp version** command.

Variable	Value
<code>rstp version {rstp stp-compatible}</code>	<p>Sets the version to RSTP or to STP compatible.</p> <p>The default is RSTP.</p>

Display Global RSTP Configuration Information

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display global RSTP configuration information:

```
show spanning-tree rstp config
```

Example

Display global RSTP configuration information:

```
Switch:1>show spanning-tree rstp config
=====
                        RSTP Configuration
=====
Rstp Module Status      : Enabled
Priority                 : 32768 (0x8000)
Stp Version              : rstp Mode
Bridge Max Age           : 20 seconds
Bridge Hello Time        : 2 seconds
Bridge Forward Delay Time : 15 seconds
Tx Hold Count            : 6
PathCost Default Type    : 32-bit
```

Display RSTP statistics

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display RSTP statistics:

```
show spanning-tree rstp statistics
```

Example

View RSTP statistics:

```
Switch:1>show spanning-tree rstp statistics
=====
                        RSTP Statistics
=====
Rstp UP Count            : 1
Rstp Down Count          : 0
Count of Root Bridge Changes : 0
Stp Time since Topology change: 0 day(s), 00H:00M:00S
Total No. of topology changes : 0
```

Display the RSTP Status

Perform the following procedure to display the RSTP related status information for the selected bridge.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the RSTP status:

```
show spanning-tree rstp status
```

Example

View the RSTP status:

```
Switch:1>show spanning-tree rstp status
=====
                        RSTP Status Information
=====
Designated Root          : 80:00:00:24:7f:9f:60:00
Stp Root Cost            : 0
Stp Root Port            : cpp
Stp Max Age              : 20 seconds
Stp Hello Time           : 2 seconds
Stp Forward Delay Time   : 15 seconds
```

Display the RSTP Configuration Information

Perform the following procedure to display the RSTP-related port level configuration details.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display RSTP configuration information:

```
show spanning-tree rstp port config {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}
```



Note

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

Display RSTP configuration information:

```
Switch:1>show spanning-tree rstp port config 1/1
=====
                        RSTP Port Configurations
=====
Port Number              : 1/1
Port Priority             : 128 (0x80)
Port PathCost            : 200000000
Port Protocol Migration  : False
Port Admin Edge Status   : False
Port Oper Edge Status    : False
Port Admin P2P Status    : Auto
Port Oper P2P Status     : False
Port Oper Protocol Version : Rstp
```


Variable Definitions

Use the data in the following table to use optional parameters with the **show spanning-tree rstp port config** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display the RSTP Status for a Port

Perform the following procedure to display the RSTP-related status information for a selected port.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the RSTP status for a port:

```
show spanning-tree rstp port status {slot/port[-slot/port][,...][slot/all][all]}
```



Note

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

View the RSTP status for a port:

```
Switch:1>show spanning-tree rstp port status 1/2
```

```

=====
                        RSTP Port Status
                      (Port Priority Vector)
=====
Port Number           : 1/2
Port Designated Root   : 80:00:00:24:7f:9f:60:00
Port Designated Cost   : 0
Port Designated Bridge : 80:00:00:24:7f:9f:60:00
Port Designated Port   : 80:c1

```

Variable Definitions

Use the data in the following table to use optional parameters with the **show spanning-tree rstp port status** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display RSTP Information for a Port

Perform the following procedure to display the RSTP-related configuration information for the selected port.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the RSTP information for a selected port:
show spanning-tree rstp port statistics [{slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]]

Example

Display the RSTP information for a selected port:

```
Switch:1# show spanning-tree rstp port statistics 1/4
```

```
=====
                        RSTP Port Statistics
=====
Port Number                : 1/4
Number of Fwd Transitions  : 0
Rx RST BPDUs Count        : 0
Rx Config BPDU Count       : 0
Rx TCN BPDU Count         : 0
Tx RST BPDUs Count        : 9
Tx Config BPDU Count       : 0
Tx TCN BPDU Count         : 0
Invalid RST BPDUs Rx Count : 0
Invalid Config BPDU Rx Count : 0
Invalid TCN BPDU Rx Count  : 0
Protocol Migration Count   : 0
```

Variable definitions

Use the data in the following table to use optional parameters with the `show spanning-tree rstp port statistics` command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display the RSTP Role

Perform the following procedure to display the RSTP information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the RSTP role:

```
show spanning-tree rstp port role [{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}
```

Example

Display the RSTP role:

```
Switch:1> show spanning-tree rstp port role 1/3
```

```
=====
                        RSTP Port Roles and States
=====
```

Port-Index	Port-Role	Port-State	PortSTPStatus	PortOperStatus
1/3	Designated	Forwarding	Enabled	Enabled

```
=====
```

Variable definitions

Use the data in the following table to use optional parameters with the **show spanning-tree rstp port role** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display Spanning Tree Configuration

Perform this procedure to view configuration and status information for spanning tree in your network.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display spanning tree configuration information:

```
show spanning-tree config
```
3. Display spanning tree status information:

```
show spanning-tree status
```

Example

Display spanning tree configuration information:

```
Switch:1> show spanning-tree config
```

```

=====
                        Spanning Tree Config
=====
      BRIDGE      BRIDGE      FORWARD
ID  PRIORITY  MAX_AGE  HELLO_TIME  DELAY  STATE
-----
0   32768     20       0           15     Enabled
1   32768     20       0           15     Enabled

      TAGGBPDU
ID  ADDRESS      TYPE      PORT
-----
0   01:80:c2:00:00:00  mstp    1/1-1/9,1/11-1/48
1   01:80:c2:00:00:00  mstp    1/10

Total number of Spanning Tree IDs : 2

```

View spanning tree status information:

```
Switch:1>show spanning-tree status
```

Spanning Tree Status								
STG ID	BRIDGE ADDRESS	NUM PORTS	PROTOCOL SPECIFICATION	TOP CHANGES				
0	00:24:7f:a1:70:00	47	ieee8021s	1				
1	00:24:7f:a1:70:00	1	ieee8021s	1				
STG ID	DESIGNATED ROOT	ROOT COST	ROOT PORT	MAX AGE	HELLO TIME	HOLD TIME	FORWARD DELAY	
0	80:00:00:24:7f:a1:70:00	0	cpp	20	0	1	15	
1	80:00:00:24:7f:a1:70:00	0	cpp	20	0	1	15	
Total number of Spanning Tree IDs : 2								

Configure Multiple Spanning Tree Protocol

Use the following procedure to configure the Multiple Spanning Tree Protocol.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Configure the MSTP forward delay for the bridge:
spanning-tree mstp forward-time <400-3000>
3. Configure the maximum age time for the bridge:
spanning-tree mstp max-age <600-4000>
4. Configure the maximum hop count for the bridge:
spanning-tree mstp max-hop <100-4000>
5. Configure the MSTP MSTI:
spanning-tree mstp msti <1-63> priority <0-65535>
6. Configure the default pathcost version:
spanning-tree mstp mstp pathcost-type <bits16 | bits32>
7. Configure the bridge priority:
spanning-tree mstp priority <0-61440>
8. Configure the region:
spanning-tree mstp region {config-id-sel <0-255> | region-name WORD<1-32> | region-version <0-65535>}
9. Configure the transmit hold count:
spanning-tree mstp tx-holdcount <1-10>
10. Configure the bridge version:
spanning-tree mstp version <mstp | rstp | stp-compatible>

Example

Configure Multiple Spanning Tree Protocol to configure the MSTP configuration version.

```
Switch>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#spanning-tree mstp forward-time 500 max-age 3000 max-hop 200 pathcost-
type bits32 priority 8192 tx-holdcount 10 version mstp
```

Variable Definitions

Use the data in the following table to use the **spanning-tree mstp** command.

Variable	Value
<i>forward-time</i> <400-3000>	Configures the MSTP forward delay for the bridge in hundredths of a second.
<i>max-age</i> <600-4000>	Assigns the MSTP maximum age time for the bridge from 600 to 4000 one hundredths of a second.
<i>max-hop</i> <100-4000>	Assigns the MSTP bridge maximum hop count. The range is 100 to 4000 one hundredths of a second. The original MIB erroneously designated the value in hundredths of a second, when it should have been in hops. The replacement MIB kept the range at 100-4000 to remain backwards compatible. To convert this value to hops, divide by 100 so 100-4000 equals 1-40 hops.
<i>msti</i> <1-63> <i>priority</i> <0-65535>	Assigns the MSTP MSTI instance parameter.
<i>pathcost-type</i> {bits16 bits32}	Assigns the MSTP default pathcost type to either 16 bits or 32 bits. The default is 32 bits.
<i>priority</i> <0-61440>	Assigns the MSTP bridge priority in a range of 0 to 61440 in steps of 4096.
<i>region</i> [<i>config-id-sel</i> <0-255>] [<i>region-name</i> <WORD 1-32>] [<i>region-version</i> <0-65535>]	Assigns the MSTP region commands: <ul style="list-style-type: none"> <i>config-id-sel</i>—Assigns the MSTP region configuration ID number. The range is 0 to 255. <i>region-name</i>—Assigns the MSTP region name. The character string can be a range of 1 to 32 characters <i>region-version</i>—Assigns the MSTP region version. The range is 0 to 65535.

Variable	Value
<code>tx-holdcount <1-10></code>	Assigns the MSTP transmit hold count. The range is 1 to 10. The default value is 3.
<code>version {mstp rstp stp-compatible}</code>	<p>Assigns the bridge version.</p> <p>Although STP and MSTP are variations of the same spanning tree protocol, they communicate information differently. A switch in MSTP mode cannot recognize the spanning tree groups running on a chassis configured with STP. MSTP spanning tree groups are not the same as STP spanning tree groups. Using a switch in MSTP mode with another chassis in STP mode can create a loop in the network.</p> <p>You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.</p> <p>Note: You cannot configure the MSTP version when MSTP-Fabric Connect Multi Homing is enabled on the switch.</p>

NEW! Configure MSTP Restricted Role

About This Task

Perform the following procedure to ensure ports do not accept superior BPDUs or claims from other non-root bridges trying to become the root bridge.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure MSTP Restricted Role (Root Guard):

```
spanning-tree mstp restricted-role enable
```

NEW! Configure Restricted Topology Change Notification

About This Task

Perform the following procedure to stop a switch port from propagating Topology Change Notifications (TCN) to other ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...] [slot/all] [all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Restricted TCN:

```
spanning-tree mstp restricted-tcn enable
```

Configure MSTP MSTI Options

Use the following procedure to configure MSTP multiple spanning tree instance (MSTI) options.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure MSTP MSTI:

```
spanning-tree mstp msti <1-63> priority <0-65535>
```

Example

Configure MSTP MSTI:

```
Switch:1(config)#spanning-tree mstp msti 62 priority 4096
```


Variable Definitions

Use the data in the following table to use the **spanning-tree mstp msti <1-63> priority <0-65535>** command.

Variable	Value
<1-63>	Specifies the instance ID.
<0-65535>	Specifies the priority value. Enter values in increments of 4096: <ul style="list-style-type: none"> • 4096 • 8192 • 12288 • 16384 • 20480 • 24576 • 28672 • 32768 • 36864 • 40960 • 45056 • 49152 • 53248 • 57344 • 61440

Configure Ethernet MSTP

Configure Ethernet MSTP on a port to enable this feature.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Ethernet MSTP:

```
spanning-tree mstp [cost <1-2000000000>] [edge-port <false|true>]
[force-port-state enable] [hello-time <100-1000>] [msti <1-63>] [p2p
{auto|force-false|force-true}] [port {slot/port[/sub-port]}] [priority
<0-240>] [protocol-migration <false|true>]
```

3. (Optional) Enable RSTP migration:

```
spanning-tree mstp allow-rstp-migration enable
```

Example

Configure Ethernet MSTP:

```
Switch:1(config)#spanning-tree mstp cost 1 edge-port true force-port-state enable hello-time 100 p2p auto priority 2 protocol-migration true allow-rstp-migration enable
```

Variable Definitions

Use the data in the following table to use the **spanning-tree mstp** command.

Variable	Value
<i>allow-rstp-migration enable</i>	Enables a switch in MSTP mode to operate in RSTP mode and send RSTP BPDUs on a specific port. The default is disabled.
<i>cost <1-200000000></i>	Configures the path cost for a port. Valid values are 1 to 200000000
<i>edge-port <false true></i>	Enables or disables the port as an edge port.
<i>force-port-state enable</i>	Enables STP.
<i>hello-time <100-1000></i>	Configures the hello-time for a port.
<i>msti <1-63></i>	Configures the port MSTP MSTI.
<i>p2p {auto force-false force-true}</i>	Enables or disables point-to-point for a port.
<i>{slot/port[/sub-port]}</i>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Variable	Value
<code>priority <0-240></code>	Configures priority for the port.
<code>protocol-migration {false true}</code>	<p>If you chose true, the option initiates protocol migration for a port. If you chose false, the option terminates protocol migration for a port.</p> <p>An MSTP-compatible port transmits and receives only RSTP BPDUs. If an MSTP port receives an STP BPDU, it becomes an STP port. User intervention is required to change this port back to MSTP mode. This process is called Port Protocol Migration.</p> <p>You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.</p>

Configure Ethernet MSTP MSTI

Use the following procedure to configure the Ethernet MSTP MSTI parameters on a port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure Ethernet MSTP MSTI:

```
spanning-tree mstp msti <1-63> [cost <1-2000000000>] [force-port-state  
enable] [port {slot/port[-slot/port] [,...][slot/all][all]}] [priority  
<0-240>]
```

Example

Configure Ethernet MSTP MSTI:

```
Switch(config-if)#spanning-tree mstp msti 62 priority 32
```

Variable Definitions

Use the data in the following table to use the **spanning-tree mstp msti <1-63>** command.

Variable	Value
<1-63>	Specifies the instance ID.
<i>cost</i> <1-2000000000>	Configures the path cost for the port
<i>force-port-state enable</i>	Enables MSTI learning for the port.
<i>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</i>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<i>priority</i> <0-240>	Configures the priority for the port. Enter the priority value (0-240) as increments of 16.

Configure MSTP Restricted Role

About This Task

Perform the following procedure to ensure ports do not accept superior BPDUs or claims from other non-root bridges trying to become the root bridge.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[, ...] [slot/all] [all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure MSTP Restricted Role (Root Guard):

```
spanning-tree mstp restricted-role enable
```

Display MSTP Configurations

Perform the following procedure to display the MSTP-related bridge-level VLAN and region information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the MSTP configurations:

```
show spanning-tree mstp config
```

Example

Display the MSTP configurations:

```
Switch:1>show spanning-tree mstp config
```

```
=====
                        MSTP Configurations
=====
Mstp Module Status      : Enabled
Number of Msti Supported : 64
Cist Bridge Priority     : 32768 (0x8000)
Stp Version             : Mstp Mode
Cist Bridge Max Age     : 20 seconds
Cist Bridge Forward Delay : 15 seconds
Tx Hold Count           : 3
PathCost Default Type   : 32-bit
Max Hop Count           : 2000
Msti Config Id Selector : 0
Msti Region Name        : 00:15:e8:9e:10:01
Msti Region Version     : 0
Msti Config Digest      : b2:96:8d:23:9d:73:39:e4:4f:bd:94:c2:14:d4:8d:09
=====
```

Display MSTP Status

Perform the following procedure to display the MSTP-related status information known by the selected bridge.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the MSTP status:

```
show spanning-tree mstp status
```

Example

Display the MSTP status:

```
Switch:1>show spanning-tree mstp status
```

```
=====
                        MSTP Status
=====
Bridge Address          : 00:15:e8:9e:10:01
Cist Root               : 80:00:00:15:e8:9e:10:01
Cist Regional Root      : 80:00:00:15:e8:9e:10:01
Cist Root Port          : cpp
Cist Root Cost           : 0
Cist Regional Root Cost : 0
Cist Instance Vlan Mapped : 1-9,11-12,14-100,102-1024
Cist Instance Vlan Mapped2k : 1025-2048
Cist Instance Vlan Mapped3k : 2049-3072
Cist Instance Vlan Mapped4k : 3073-3999,4001-4094
=====
```

Cist Max Age	: 20 seconds
Cist Forward Delay	: 15 seconds

Display MSTP Port Information

Perform the following procedure to display the MSTP, CIST port, and MSTI port information maintained by every port of the common spanning tree.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Display the MSTP port information:
`show spanning-tree mstp port role [{slot/port[-slot/port]][,...][slot/all][all]}`



Note
If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

Display the MSTP port information:

```
Switch:1>show spanning-tree mstp port role 1/3
```

CIST Port Roles and States				
Port-Index	Port-Role	Port-State	PortSTPStatus	PortOperStatus
1/3	Disabled	Discarding	Enabled	Disabled

NEW Display MSTP Port Configuration Information

About This Task

Perform the following procedure to display MSTP port confirmation information in your network.

Procedure

- 1. Enter GigabitEthernet Interface Configuration mode:
`interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}`



Note
If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- 2. Display MSTP port configuration information in your network:
`show spanning-tree mstp port config {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}`

Example

Display MSTP port configuration information:

```
Switch:1>#show spanning-tree mstp port config 1/7
=====
MSTP Cist Port Information
(Port Priority Vector)
=====
Port Number           : 1/7
Cist Port Priority     : 128 (0x80)
Cist Port force-port-state : disable
Cist Port cost        : 200000000
Cist Port Designated Root : 80:00:40:88:2f:dc:94:00
Cist Port Designated Cost : 0
Cist Port Designated Bridge : 80:00:40:88:2f:dc:94:00
Cist Port Designated Port : 80:c6
Cist Port Regional Root : 80:00:40:88:2f:dc:94:00
Cist Port Regional PathCost : 0
Cist Port Protocol Migration : False
Cist Port Admin Edge Status : False
Cist Port Oper Edge Status : False
Cist Port Admin P2P Status : Auto
Cist Port Oper P2P Status : False

Cist Port Hello Time : 2
Cist Port Oper Proto-Version : Mstp
Cist Port Rstp Migration : Disabled
Cist Port Restricted Role : Disabled
Cist Port Restricted TCN : Disabled
```

Display MSTP MSTI Information

Display MSTP MSTI information to ensure the feature is configured correctly for your network.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display MSTI information:

```
show spanning-tree mstp msti [config <1-63>] [port config {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} | role {slot/port[-slot/port][,...][slot/all][all]} | statistics {slot/port[-slot/port][,...][slot/all][all]}
```



Note

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

Display MSTI information:

```
Switch:1>#show spanning-tree mstp msti config 62
=====
MSTP Instance Status
=====
```

```

Instance Id                : 62
Msti Bridge Regional Root  : 80:00:00:15:e8:9e:10:01
Msti Bridge Priority       : 32768 (0x8000)
Msti Root Cost            : 0
Msti Root Port            : cpp
Msti Instance Vlan Mapped : 
Msti Instance Vlan Mapped2k : 
Msti Instance Vlan Mapped3k : 
Msti Instance Vlan Mapped4k : 4000

Switch:1>#show spanning-tree mstp msti port statistics 1/1
=====
                        MSTP Instance-specific Per-Port Statistics
=====
Port Number                : 1/1
Instance Id                : 1
Msti Port Fwd Transitions  : 0
Msti Port Received BPDUs   : 0
Msti Port Transmitted BPDUs : 0
Msti Port Invalid BPDUs Rcvd : 0

```

Variable Definitions

Use the data in the following table to use the **show spanning-tree mstp msti** command.

Variable	Value
<i>config</i> [<i><1-63></i>]	Shows the configuration for one or all MSTP instance IDs.
<i>port</i>	Shows the configuration, role, or statistics information of a MSTP port. <ul style="list-style-type: none"> <code>config {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]</code> <code>role {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]</code> <code>statistics {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]</code>

Display MSTP Statistics

Display MSTP MSTI information to ensure the feature is configured correctly for your network.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display MSTP statistics:

```
show spanning-tree mstp statistics
```

Example

Display MSTP statistics:


```
Switch:1>show spanning-tree mstp statistics
```

```
=====
                        MSTP Bridge Statistics
=====
Mstp UP Count           : 1
Mstp Down Count         : 0
Region Config Change Count : 4
Time Since Topology Change : 0 seconds
Topology Change Count   : 0
New Root Bridge Count   : 1
```

Configure tc-receive-alarm-threshold

About This Task

You can establish a threshold rate for receiving TC/TCN packets by configuring values for count and interval. If the threshold rate is exceeded, you receive a warning message.

Use the following procedure to configure values for **tc-receive-alarm-threshold count** and **tc-receive-alarm-threshold interval**.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure a value for count:

```
spanning-tree tc-receive-alarm-threshold count <1-1000>
```
3. Configure a value for interval:

```
spanning-tree tc-receive-alarm-threshold interval <1-15>
```

Example

Configure count and interval values:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#spanning-tree tc-receive-alarm-threshold count 3
Switch:1(config)#spanning-tree tc-receive-alarm-threshold interval 2
```

Display the alarm threshold configuration:

```
Switch:1#show spanning-tree tc-receive-alarm-threshold
```

```
=====
                        Spanning Tree TC/TCN Rx Alarm Threshold Configuration
=====
Interval           : 2 minutes
Count              : 3
```

Variable definitions

Use the data in the following table to use the **tc-receive-alarm-threshold** command.

Variable	Value
count <1-1000>	Specifies the number of packets used to establish the threshold rate. The default is 2.
interval <1-15>	Specifies the time interval (in minutes) used to establish the threshold rate. The default is 1.

Display the tc-receive-alarm-threshold Configuration

About This Task

You can configure a threshold rate for receiving TC/TCN packets. If the threshold is exceeded, you receive a warning message.

Use the following procedure to display the tc-receive-alarm-threshold configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the tc-receive-alarm-threshold configuration:
`show spanning-tree tc-receive-alarm-threshold`

Example

```
Switch:1(config)#show spanning-tree tc-receive-alarm-threshold
```

```
=====
Spanning Tree TC/TCN Rx Alarm Threshold Configuration
=====
Interval          : 1 minutes
Count             : 2
```

Spanning Tree Configuration Using EDM

This section describes how to create, manage, and monitor spanning tree groups (STG). It also describes how to configure the Rapid Spanning Tree Protocol (RSTP) and the Multiple Spanning Tree Protocol (MSTP) using Enterprise Device Manager (EDM).



Important

The switch supports up to 64 STGs in a device, however, SPBM uses STG 63 and MSTI 62 for internal use. STG 63 or MSTI 62 cannot be used by other VLANs or MSTIs.

Configure the Spanning Tree Mode

About This Task

Configure the Spanning Tree mode to change the mode to MSTP or RSTP mode.



Important

After you change the mode, save the configuration, and then restart the system for the changes to take effect.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **Globals**.
3. Select the required spanning tree mode.
4. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
SpanningTreeAdminMode	Configures the spanning tree mode as either RSTP or MSTP. The default is MSTP.
SpanningTreeOperMode	Specifies the current mode of the spanning tree.

Configure BPDU Guard

About This Task

Configure BPDU Guard to block the root selection process or to prevent BPDU flooding from unknown devices.

To configure multiple ports simultaneously, select more than one port in the Device Physical View tab. The system displays the **BPDU Guard** tab as a table-based tab.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. Select **BpduGuardAdminEnabled** to enable BPDU Guard for the port.
6. (Optional) Type a value in **BpduGuardTimeout** to configure the timer for port-state recovery.
7. Select **Apply**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Index	Displays the index of the port, written in the slot/port[/sub-port] format.
Name	Configures the name of the port.
Descr	Displays the description of the port. A textual string containing information about the interface.
Type	Displays the type of connector plugged in the port.
Mtu	Displays the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram which can be sent or received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.
PhysAddress	Displays the physical address of the port. The address of the interface at the protocol layer immediately below the network layer in the protocol stack. For interfaces which do not have such an address, for example, a serial line, this object should contain an octet string of zero length.
VendorDescr	Displays the vendor of the connector plugged in the port.
DisplayFormat	Identifies the slot and port numbers (slot/port). If the port is channelized, the format also includes the sub-port in the format slot/port/sub-port
AdminStatus	Configures the port as enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
OperStatus	Displays the current status of the port. The status includes enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
ShutdownReason	Indicates the reason for a port state change.
LastChange	Displays the timestamp of the last change.
LinkTrap	Enable or disable link trapping.

Name	Description
AutoNegotiate	Enables or disables Auto-Negotiation for this port. The default Auto-Negotiation behavior depends on the switch model and transceiver type.
AutoNegAd	Specifies the port speed and duplex abilities to advertise during link negotiation. Supported speeds and duplex modes vary, depending on your hardware. The abilities specified in this object are only used when auto-negotiation is enabled on the port. If all bits in this object are disabled, and auto-negotiation is enabled on the port, then the physical link process on the port is disabled (if hardware supports this ability). Any change to this configuration restarts the auto-negotiation process, which has the same effect as physically unplugging and reattaching the cable attached to the port. If you select default , all capabilities supported by the hardware are advertised.
AdminDuplex	Configures the administrative duplex setting for the port.
OperDuplex	Indicates the operational duplex setting for the port.
AdminSpeed	Configures the administrative speed for the port.
OperSpeed	Indicates the operational speed for the port.
QoSLevel	Selects the Quality of Service (QoS) level for this port. The default is level1.
DiffServ	Enables the Differentiated Service feature for this port. The default is disabled.
Layer3Trust	Configures if the system should trust Layer 3 packets coming from access links or core links only. The default is core.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled (selected) or disabled (cleared) on the port. The default is disabled (clear).
MltId	Shows the MLT ID associated with this port. The default is 0.
Locked	Shows if the port is locked. The default is unlocked.

Name	Description
UnknownMacDiscard	Discards packets that have an unknown source MAC address, and prevents other ports from sending packets with that same MAC address as the destination MAC address. The default is disabled.
DirectBroadcastEnable	<p>Specifies that an Isolated Routing Port (IRP) can forward directed broadcast traffic. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcast on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling this function protects a host from possible denial of service (DoS) attacks.</p> <p>With the feature enabled, the Control Processor (CP) does not receive a copy of the directed broadcast. As a result, the system does not respond to a subnet broadcast ping sent from a remote subnet.</p> <p>The default is disabled.</p>
OperRouting	Shows the routing status of the port.
HighSecureEnable	Enables or disables the high secure feature for this port.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
FlexUniEnable	Enables Flex UNI on the port. The default is disabled.
EgressRateLimitState	Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the port. The default is disabled.
EgressRateLimit	<p>Specifies the egress rate limit in Kbps. You cannot configure the egress shaper rate to exceed the port capability.</p> <p>If you configure this value to 0, shaping is disabled on the port.</p>
TxFlowControl	<p>Configures if the port sends pause frames. By default, an interface does not send pause frames.</p> <p>You must also enable the flow control feature globally before an interface can send pause frames.</p>
TxFlowControlOperState	Shows the operational state of flow control.

Name	Description
BpduGuardTimerCount	Shows the time, starting at 0, since the port became disabled. When the BpduGuardTimerCount reaches the BpduGuardTimeout value, the port is enabled. Displays in 1/100 seconds.
BpduGuardTimeout	Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.
BpduGuardAdminEnabled	Enables BPDU Guard on the port. The default is disabled.
ForwardErrorCorrection	Configures one of the following options for Forward Error Correction (FEC) on the port: <ul style="list-style-type: none"> • CL 91 • CL 108 • CL 74 • disable • auto The disable option disables this configuration on the port.
ForwardErrorCorrectionApplicability	Displays whether FEC is applicable on the interface.
OperAutoNegotiate	Shows the operational state of Auto-Negotiation.
OperForwardErrorCorrection	Shows the negotiated operational FEC clause. If the value is off, the port supports FEC and is up but not configured for FEC. If the value is notApplicable, the port does not support FEC. If the value is unknown, the port supports FEC but is down.
Action	Performs one of the following actions on the port <ul style="list-style-type: none"> • none - none of the following actions • flushMacFdb - flush the MAC forwarding table • flushArp - flush the ARP table • flushIp - flush the IP route table • flushAll - flush all tables • triggerRipUpdate — manually triggers a RIP update The default is none.
Result	Displays the result of the selected action. The default is none.

Name	Description
AutoSense	Enables or disables Auto-sense on the specific port. The default value is disabled for existing configurations but enabled for new Zero Touch Fabric Configuration deployments.
AutoSenseKeepAutoConfig	Retains the Auto-sense configuration if you disable Auto-sense on the port. The dynamic configuration becomes a manual configuration and is visible in the show running-config output.
CustomAutoNegAdOrigin	Specifies the origin of Custom Auto Negotiation Advertisements (CANA) configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
BpduGuardOrigin	Specifies the origin of BPDU Guard configuration on the port. The supported values are: <ul style="list-style-type: none"> • config - Set by the user. • radius - Set by the Remote Authentication Dial-In User Service (RADIUS) attribute.
AutoSenseState	Displays the Auto-sense port state.
LinkDebounce	Specifies the extended debounce timer on the port. The value 0 milliseconds disables debounce time. The default value is 1000. Note: This field does not apply to Auto-sense ports. To configure Link Debounce on Auto-sense ports, see Configure Link Debounce for Auto-sense Ports on page 65.
AutoSenseDataIsid	Specifies the Auto-sense data I-SID per port.
AutoSenseWaitInterval	Specifies the Auto-sense Wait Interval for the port. The default is 35.
NodeAliasEnable	Enables or disables Node Alias on this specific port. The default is enabled on Auto-sense ports.

Name	Description
NodeAliasClear	Clears Node Alias on this specific port.
IpfixEnable	Enables or disables IPFIX on the port. The default is true, which means enabled. If you enable IPFIX globally, the switch learns flows on both UNI ports and NNIs automatically, because IPFIX is enabled on those interfaces by default. You must explicitly disable IPFIX on those interfaces to stop flow learning. If you disable IPFIX globally, no flow learning occurs, regardless of the UNI port or NNI configuration.

Configure RSTP Global Parameters

Perform this procedure to configure the RSTP global parameters.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **RSTP**.
3. Configure the parameters as required.
4. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
PathCostDefault	Specifies the version of the spanning tree default path costs that are used by this bridge. A value of 8021d1998 indicates the use of the 16-bit default path costs from IEEE Std. 802.1d-1998. A value of stp8021t2001 indicates the use of the 32-bit default path costs from IEEE Std. 802.1t.
TxHoldCount	Specifies the value used by the port transmit state machine to limit the maximum transmission rate. The default is 3.
Version	Specifies the version of STP that the bridge currently runs. The value stpCompatible indicates that the Spanning Tree Protocol as specified in IEEE 802.1d is in use; rstp indicates that the Rapid Spanning Tree Protocol as specified in IEEE 802.1w is in use.
EnableStp	Indicates whether the spanning tree protocol is active in this STG. The default is enabled.

Name	Description
Priority	Specifies the RSTP bridge priority.
BridgeMaxAge	Specifies the value that all bridges use for MaxAge while this bridge acts as the root.
BridgeHelloTime	The value that all bridges use for HelloTime while this bridge acts as the root.
BridgeForwardDelay	Specifies the value that all bridges use for forward delay while this bridge acts as the root.
DesignatedRoot	Specifies the unique bridge identifier of the bridge recorded as the root in the configuration BPDUs transmitted by the designated bridge for the segment to which the port is attached.
RootCost	Specifies the cost of the path to the root from this bridge.
RootPort	Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge.
MaxAge	Specifies the maximum age of Spanning Tree Protocol information in hundredths of a second learned from the network on any port before the port is discarded.
HelloTime	Specifies the amount of time in hundredths of a second between the transmission of configuration bridge PDUs by this node on any port while it is the root of the spanning tree (or trying to become the root).
ForwardDelay	Specifies a time value, measured in hundredths of a second, controls how fast a port changes its spanning state after moving towards the forwarding state. The value determines how long the port stays in each of the listening and learning states, which precede the forwarding state. This value is also used after a topology change is detected, and is underway, to age all dynamic entries in the forwarding database.
RstpUpCount	Specifies the number of times the RSTP module is enabled. A trap is generated on the occurrence of this event.
RstpDownCount	Specifies the number of times the RSTP module is disabled. A trap is generated on the occurrence of this event.
NewRootIdCount	Specifies the number of times this bridge detects a root identifier change. A trap is generated on the occurrence of this event.

Name	Description
TimeSinceTopology Change	Specifies the time (in hundredths of a second) since the TcWhile Timer for any port in this bridge was nonzero for Common Spanning Tree.
TopChanges	Specifies the number of times that there was at least one nonzero TcWhile Timer on this bridge for Common Spanning Tree.

Configure RSTP Ports

About This Task

Configure RSTP to reduce the recovery time after a network breakdown.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **RSTP**.
3. Select the **RSTP Ports** tab.
4. Use the fields in the **RSTP Ports** tab to configure the RSTP ports.
5. Select **Apply**.

RSTP Ports Field Descriptions

Use the data in the following table to use the **RSTP Ports** tab.

Name	Description
Port	Specifies a unique value, greater than zero, indicating the port number.
Priority	Specifies the value of the priority field.
PathCost	Specifies the contribution of this port to the path cost of paths towards the root that includes this port.
ProtocolMigration	Specifies a port to transmit RSTP BPDUs if operating in RSTP mode. Any other operation on this object has no effect, and RSTP mode returns false if read.
AdminEdgePort	Specifies the administrative value of the Edge Port parameter. A value of true indicates that this port is an edge-port, and a value of false indicates that this port is a nonedge-port.
OperEdgePort	Specifies the operational value of the Edge Port parameter. The object is initialized to the value of AdminEdgePort and is configured to false on reception of a BPDU.

Name	Description
AdminPointToPoint	Specifies the administrative point-to-point status of the LAN segment attached to this port. A value of forceTrue indicates that this port is treated as if it is connected to a point-to-point link. A value of forceFalse indicates that this port is treated as having a shared media connection. A value of auto indicates that this port is considered to have a point-to-point link if it is an aggregator and all of its members are aggregatable, or if the MAC entity is configured for full-duplex operation, either through autonegotiation or by management means.
OperPointToPoint	Specifies the operational point-to-point status of the LAN segment attached to this port. It indicates whether a port is considered to have a point-to-point connection or not. The value is determined by management or by autodetection as described in the AdminPointToPoint object.
OperVersion	Indicates if the port is in MSTP mode, RSTP mode or STP-compatible mode. MSTP mode transmits MST BDUs, RSTP mode transmits RST BPDUs and STP-compatible transmits Config/TCN BPDUs.

View RSTP Port Status

About This Task

View the RSTP port status to ensure proper functioning of RSTP.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **RSTP**.
3. Select the **RSTP Status** tab.

RSTP Status *Field Descriptions*

Use the data in the following table to use the **RSTP Status** tab.

Name	Description
Port	Specifies a unique value, greater than zero, indicating the port number.
State	Specifies the current state of the port as defined by application of the Spanning Tree Protocol. This state controls what action a port takes on reception of a frame.
Role	Indicates the current port role assumed by this port.
OperVersion	Indicates whether the port is operationally in the RSTP- or STP-compatible mode; that is, whether the port transmits RSTP BPDUs or Config/TCN BPDUs.
EffectivePortState	Specifies the effective operational state of the port. This object is configured to true if the port is operationally up in the Interface Manager, and if Force Port State for this port and the specified port state is enabled. Otherwise, this object is configured to false.

View RSTP Status Statistics**About This Task**

You can view status statistics for Rapid Spanning Tree Protocol (RSTP).

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **RSTP**.
3. On the **RSTP Status** tab, select a port, and then select **Graph**.

RSTP Stats *Field Descriptions*

Use the data in the following table to use the **RSTP Stats** tab.

Name	Description
RxRstBpduCount	Specifies the number of RSTP BPDUs this port received.
RxConfigBpduCount	Specifies the number of configuration BPDUs this port received.
RxTcnBpduCount	Specifies the number of TCN BPDUs this port received.
TxRstBpduCount	Specifies the number of RSTP BPDUs this port transmitted.

Name	Description
TxConfigBpduCount	Specifies the number of Config BPDUs this port transmitted.
TxTcnBpduCount	Specifies the number of TCN BPDUs this port transmitted.
InvalidRstBpduRxCount	Specifies the number of invalid RSTP BPDUs this port received. A trap is generated on the occurrence of this event.
InvalidConfigBpduRx Count	Specifies the number of invalid configuration BPDUs this port received. A trap is generated on the occurrence of this event.
InvalidTcnBpduRxCount	Specifies the number of invalid TCN BPDUs this port received. A trap is generated on the occurrence of this event.
ProtocolMigrationCount	Specifies the number of times this port migrated from one STP protocol version to another. The relevant protocols are STP-Compatible and RSTP. A trap is generated on the occurrence of this event.

View Port Spanning Tree Statistics

About This Task

View port spanning tree statistics to manage network performance.

Procedure

1. In the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Graph**.
3. Select **Port**.
4. Select the **Spanning Tree** tab.

Spanning Tree Field Descriptions

The following table describes parameters on the **Spanning Tree** tab.

Name	Description
InConfigBpdus	The number of Config BPDUs received.
InTcnBpdus	The number of Topology Change Notifications BPDUs received.
InBadBpdus	The number of unknown or malformed BPDUs received.
OutConfigBpdus	The number of Config BPDUs transmitted.
OutTcnBpdus	The number of Topology Change Notifications BPDUs transmitted.

Configure MSTP Global Parameters

Before You Begin

- The system must be in MSTP mode.

About This Task

Configure the global MSTP parameters to determine how MSTP operates for the system. Interface-level parameters override global settings.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **MSTP**.
3. Select the **Globals** tab.
4. Configure MSTP as required.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
PathCostDefaultType	Specifies the version of the spanning tree default path costs to be used by this bridge. A value of 8021d1998 denotes the use of the 16-bit default path costs from IEEE 802.1d-1998. A value of stp8021t2001 denotes the use of the 32-bit default path costs from IEEE 802.1t.
TxHoldCount	Specifies the value used by the port transmit state to limit the maximum transmission rate. The default is 3.
MaxHopCount	Assigns the MSTP bridge maximum hop count. The range is 100 to 4000 one hundredths of a second. The original MIB erroneously designated the value in hundredths of a second, when it should have been in hops. The replacement MIB kept the range at 100-4000 to remain backwards compatible. To convert this value to hops, divide by 100 so 100-4000 equals 1-40 hops.
NoOfInstancesSupported	Indicates the maximum number of spanning tree instances supported.
MstpUpCount	The number of times the MSTP module is enabled. A trap is generated on the occurrence of this event.
MstpDownCount	The number of times the MSTP module is disabled. A trap is generated on the occurrence of this event.

Name	Description
ForceProtocolVersion	<p>Specifies the version of Spanning Tree Protocol that the bridge currently runs. stpCompatible indicates that the Spanning Tree Protocol as specified in IEEE 802.1d is in use; rstp indicates that the Rapid Spanning Tree Protocol as specified in IEEE 802.1w is in use; and mstp indicates that the multiple spanning tree protocol as specified in IEEE 802.1s is in use.</p> <p>Although STP and MSTP are variations of the same spanning tree protocol, they communicate information differently. A switch in MSTI mode cannot recognize the spanning tree groups running on a chassis configured with STP. MSTP spanning tree groups are not the same as STP spanning tree groups. Using a switch in MSTP mode with a chassis in STP mode can create a loop in the network.</p> <p>The default is MSTP.</p>
BrgAddress	<p>Specifies the MAC address used by this bridge if it must be referred to in a unique fashion. This should be the numerically smallest MAC address of all ports that belong to this bridge. If concatenated with MstCistBridgePriority or MstBridgePriority, a unique bridge identifier is formed, which is used in the STP.</p>
Root	<p>Specifies the bridge identifier of the root of the common spanning tree as determined by the STP by this node. This value is used as the CIST root identifier parameter in all configuration bridge PDUs originated by this node.</p>
RegionalRoot	<p>Specifies the bridge identifier of the root of the multiple spanning tree region as determined by the STP as executed of this node. This value is used as the common and internal spanning tree (CIST) regional root identifier parameter in all configuration bridge PDUs originated by this node.</p>
RootCost	<p>Specifies the cost of the path to the CIST root from this bridge.</p>
RegionalRootCost	<p>Specifies the cost of the path to the CIST regional root from this bridge.</p>
RootPort	<p>Specifies the port number of the port which offers the lowest path cost from this bridge to the CIST root bridge.</p>

Name	Description
BridgePriority	Specifies the value of the writable portion of the bridge identifier comprising the first two octets. The values you enter for bridge priority must be in steps of 4096. The default is 32768.
BridgeMaxAge	Specifies the value that all bridges use for MaxAge while this bridge acts as the root. The granularity of this timer is specified as 1 second. An agent can return a bad value error if you attempt to configure a value which is not a whole number of seconds. The default is 2000.
BridgeForwardDelay	Specifies the value that all bridges use for forward delay if this bridge acts as the root. Note that 802.1d specifies that the range for this parameter is related to the value of BridgeMaxAge. The granularity of this timer is specified as 1 second. An agent can return a bad value error if you attempt to configure a value which is not a whole number of seconds. The default is 1500.
HoldTime	Determines the interval length in hundredths of a second during which no more than two configuration bridge PDUs can be transmitted by this node.
MaxAge	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the value that this bridge currently uses.
ForwardDelay	Specifies the time value, measured in units of hundredths of a second, that controls how fast a port changes its spanning state after moving towards the forwarding state. This value determines how long the port stays in a particular state before moving to the next state.
TimeSinceTopology Change	Specifies the time (in hundredths of a second) since the TcWhile Timer for any port in this bridge was nonzero for Common Spanning Tree.
TopChanges	Specifies the number of times that there was at least one nonzero TcWhile Timer on this bridge for Common Spanning Tree.
NewRootBridgeCount	Specifies the number of times this bridge detects a root bridge change for Common Spanning Tree. A trap is generated on the occurrence of this event.

Name	Description
RegionName	Specifies the name for the region configuration. By default, the region name is equal to the bridge MAC Address.
RegionVersion	Specifies the version of the MST region.
ConfigIdSel	Specifies the configuration identifier format selector used by the bridge. This has a fixed value of 0 to indicate RegionName. RegionVersions are specified as in the standard.
ConfigDigest	Specifies the configured MD5 digest value for this region, which must be 16 octets long.
RegionConfigChange Count	Specifies the number of times a region configuration identifier change is detected. A trap is generated on the occurrence of this event.

Configure CIST Ports for MSTP

About This Task

Configure Common and Internal Spanning Tree (CIST) ports to configure ports for MSTP.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **MSTP**.
3. Select the **CIST Port** tab.



Note

The MSTP, CIST Port tab contains information for each port that is common to all bridge and spanning tree instances.

4. Use the fields in the **CIST Port** box to configure the MSTP CIST port.
5. Select **Apply**.

CIST Port Field Descriptions

Use the data in the following table to use the **CIST Port** tab.

Name	Description
Port	Specifies the port number of the port for which this entry contains spanning tree information.
PathCost	Specifies the contribution of this port to the path cost of paths towards the CIST root that includes this port.

Name	Description
Priority	<p>Specifies the four most significant bits of the port identifier of the spanning tree instance which are modified by setting the CistPortPriority value. The values that are configured for port priority must be in steps of 16.</p> <p>Although port priority values can range from 0 to 255, only the following values are used: 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, 240.</p> <p>The default is 128.</p>
DesignatedRoot	Specifies the unique bridge identifier of the bridge recorded as the CIST root in the configuration BPDUs transmitted.
DesignatedCost	Specifies the path cost of the designated port of the segment that connects to this port.
DesignatedBridge	Specifies the unique bridge identifier of the bridge which that port considers to be the designated bridge for the ports segment.
DesignatedPort	Specifies the port identifier of the port on the designated bridge for this port segment.
RegionalRoot	Specifies the unique bridge identifier of the bridge recorded as the CIST regional root identifier in the configuration BPDUs transmitted.
RegionalPathCost	Specifies the contribution of this port to the path cost of paths towards the CIST regional root that include this port.
ProtocolMigration	<p>Specifies the protocol migration state of this port. If you chose true, the option initiates protocol migration for a port. If you chose false, the option terminates protocol migration for a port.</p> <p>An MSTP-compatible port transmits and receives only RSTP BPDUs. If an MSTP port receives an STP BPDU, it becomes an STP port. User intervention is required to change this port back to MSTP mode. This process is called Port Protocol Migration.</p> <p>You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.</p>

Name	Description
AdminEdgeStatus	Specifies the administrative value of the Edge Port parameter. A value of true indicates that this port is an edge-port, and a value of false indicates that this port is a nonedge-port.
OperEdgeStatus	Specifies the operational value of the Edge Port parameter. The object is initialized to the value of AdminEdgeStatus and is configured to false on reception of a BPDU.
AdminP2P	Specifies the administrative point-to-point status of the LAN segment attached to this port. A value of forceTrue indicates that this port is treated as if it connects to a point-to-point link. A value of forceFalse indicates that this port is treated as having a shared media connection. A value of auto indicates that this port is considered to have a point-to-point link if it is an aggregator and all of its members are aggregatable, or if the MAC entity is configured for full-duplex operation, either through autonegotiation or by management means.
OperP2P	Specifies the operational point-to-point status of the LAN segment attached to this port. It indicates whether a port is considered to have a point-to-point connection or not. The value is determined by management or by autodetection as described in the AdminP2P object.
HelloTime	Specifies the amount of time in hundredths of a second between the transmission of configuration bridge PDUs by this node on this port.
OperVersion	Specifies whether the port is operationally in the MSTP mode, the RSTP mode, or the STP-compatible mode; that is, whether the port transmits MST BPDUs, RST BPDUs, or Config/TCN BPDUs. Although STP and MSTP are variations of the same spanning tree protocol, they communicate information differently. A switch in MSTI mode cannot recognize the spanning tree groups running on a chassis configured with STP. MSTP spanning tree groups are not the same as STP spanning tree groups. Using a switch in MSTP mode with another chassis in STP mode can create a loop in the network.

Name	Description
EffectivePortState	Specifies the effective operational state of the port for CIST. This is true only if the port is operationally up at the interface and protocol levels for CIST. This is configured to false for all other conditions.
State	Specifies the current state of the port as defined by the common spanning tree protocol. It can be disabled, discarding, learning, or forwarding.
ForcePortState	Specifies the current state of the port. You can change the port to either Disabled or Enabled for the base spanning tree instance.
SelectedPortRole	Specifies the selected port role of the port for this spanning tree instance.
CurrentPortRole	Specifies the current port role of the port for this spanning tree instance.
RSTPMigration	Specifies whether the port operates in RSTP mode to send RSTP BPDUs. You can change the port to either Enabled or Disabled. The default is disabled.
RestrictedRole	Specifies whether the port accepts superior BPDUs or claims from other non-root bridges trying to become the root bridge. The default is disabled.
RestrictedTcn	Specifies whether the port propagates Topology Change Notifications (TCN) to other ports. The default is disabled.

Configure MSTI Bridges for MSTP

About This Task

Perform this procedure to configure multiple spanning tree instance (MSTI) bridges for MSTP.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **MSTP**.

3. Select the **MSTI Bridges** tab.

**Note**

The system generates MSTI bridge instances after you create a VLAN in MSTP mode.

4. Use the fields in the **MSTI Bridges** box to configure the MSTP bridge.
5. Select **Apply**.

MSTI Bridges *Field Descriptions*

Use the data in the following table to use the **MSTI Bridges** tab.

Name	Description
Instance	Specifies the spanning tree instance to which this information belongs.
RegionalRoot	Specifies the MSTI regional root identifier value for the instance. This value is used as the MSTI regional root identifier parameter in all configuration bridge PDUs originated by this node.
Priority	Specifies the writable portion of the MSTI bridge identifier comprising the first two octets. The values that are configured for bridge priority must be in steps of 4096. The default is 32768.
RootCost	Specifies the cost of the path to the MSTI regional root as seen by this bridge.
RootPort	Specifies the port number of the port that offers the lowest path cost from this bridge to the MSTI region root bridge.
TimeSinceTopologyChange	Specifies the time (in hundredths of a second) since the TcWhile Timer for any port in this bridge was nonzero for this spanning tree instance.
TopChanges	Specifies the number of times that there was at least one nonzero TcWhile Timer on this bridge for this spanning tree instance.
NewRootCount	Specifies the number of times this bridge detects a root bridge change for this spanning tree instance. A trap is generated on the occurrence of this event.
InstanceUpCount	Specifies the number of times a new spanning tree instance is created. A trap is generated on the occurrence of this event.
InstanceDownCount	Specifies the number of times a spanning tree instance is deleted. A trap is generated on the occurrence of this event.

Configure MSTI Ports for MSTP

About This Task

Perform the following procedure to configure MSTI ports for MSTP.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **MSTP**.
3. Select the **MSTI Port** tab.
4. Use the fields in the **MSTI Port** box to configure the MSTP.
5. Select **Apply**.

MSTI Port Field Descriptions

Use the data in the following procedure to use the **MSTI Port** tab.

Name	Description
Port	Specifies the port number of the port for which this entry contains spanning tree information.
Instance	Specifies the spanning tree instance to which the information belongs.
PathCost	Specifies the contribution of this port to the path cost of paths towards the MSTI root that includes this port.
Priority	Specifies the four most significant bits of the port identifier for a given spanning tree instance can be modified independently for each spanning tree instance supported by the bridge. The values configured for port priority must be in steps of 16. The default is 128.
DesignatedRoot	Specifies the unique bridge identifier of the bridge recorded as the MSTI regional root in the configuration BPDUs transmitted.
DesignatedBridge	Specifies the unique bridge identifier of the bridge that this port considers to be the designated bridge for the port segment.
DesignatedPort	Specifies the port identifier of the port on the designated bridge for this port segment.
State	Specifies the current state of the port, as defined by the MSTP. A port which is in forwarding state in one instance can be in discarding (blocking) state in another instance.

Name	Description
ForcePortState	Specifies the current state of the port, that is changed to either disabled or enabled for the specific spanning tree instance.
DesignatedCost	Specifies the path cost of the designated port of the segment connected to this port.
CurrentPortRole	Specifies the current port role of the port for this spanning tree instance.
EffectivePortState	Specifies the effective operational state of the port for a specific instance. This is configured to true if the port is operationally up at the interface and protocol levels for the specific instance. This is configured to false at all other times.

View VLAN and Spanning Tree CIST Statistics

About This Task

View CIST port statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **MSTP**.
3. Select the **CIST Port** tab.
4. Select a port, and then select **Graph**.

CIST Field Descriptions

The following table describes parameters on the **CIST** tab.

Name	Descriptions
ForwardTransitions	Specifies the number of times this port has transitioned to the forwarding state.
RxMstBpduCount	Specifies the number of MSTP BPDUs received on this port.
RxRstBpduCount	Specifies the number of RSTP BPDUs received on this port.
RxConfigBpduCount	Specifies the number of configuration BPDUs received on this port.
RxTcnBpduCount	Specifies the number of TCN BPDUs received on this port.
TxMstBpduCount	Specifies the number of MSTP BPDUs transmitted from this port.

Name	Descriptions
TxRstBpduCount	Specifies the number of RSTP BPDUs transmitted from this port.
TxConfigBpduCount	Specifies the number of configuration BPDUs transmitted from this port.
TxTcnBpduCount	Specifies the number of TCN BPDUs transmitted from this port.
InvalidMstBpduRxCount	Specifies the number of Invalid MSTP BPDUs received on this port.
InvalidRstBpduRxCount	Specifies the number of Invalid RSTP BPDUs received on this port.
InvalidConfigBpduRxCount	Specifies the number of invalid configuration BPDUs received on this port.
InvalidTcnBpduRxCount	Specifies the number of invalid TCN BPDUs received on this port. The number of times this port has migrated from one STP protocol version to another. The relevant protocols are STP-compatible and RSTP/MSTP. A trap is generated on the occurrence of this event.
ProtocolMigrationCount	Specifies the number of times this port has migrated from one STP protocol version to another. The relevant protocols are STP-compatible and RSTP. A trap is generated on the occurrence of this event.

View VLAN and Spanning Tree MSTI Statistics

About This Task

View multiple spanning tree instance (MSTI) port statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > VLAN > Spanning Tree**.
2. Select **MSTP**.
3. Select the **MSTI Port** tab.
4. Select a port, and then select **Graph**.

MSTI Field Descriptions

The following table describes parameters on the **MSTI** tab.

Name	Description
ForwardTransitions	Specifies the number of times this port has transitioned to the forwarding state for this specific instance.
ReceivedBPDUs	Specifies the number of BPDUs received by this port for this spanning tree instance.

Name	Description
TransmittedBPDUs	Specifies the number of BPDUs transmitted on this port for this spanning tree instance.
InvalidBPDUsRcvd	Specifies the number of invalid BPDUs received on this port for this spanning tree instance.



SPB-PIM Gateway configuration

[SPB-PIM Gateway Fundamentals](#) on page 3558

[Multicast Source Discovery Protocol Configuration](#) on page 3571

[Controller Configuration](#) on page 3609

[Gateway Configuration](#) on page 3620

[SPB-PIM Gateway Interface Configuration](#) on page 3626

[SPB-PIM Gateway Interface Deployment Scenarios](#) on page 3642

Table 307: SPB-PIM Gateway product support

Feature	Product	Release introduced
SPB-PIM Gateway controller node	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Not Supported
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1.5 Not supported on VSP4900-48P
	VSP 7400 Series	VOSS 8.0
SPB-PIM Gateway interface	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Not Supported
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

SPB-PIM Gateway configuration includes controller and gateway nodes, and Multicast Source Discovery Protocol (MSDP) configuration.

The following topics provide necessary concepts and procedures to configure SPB-PIM Gateway.

SPB-PIM Gateway Fundamentals

This section provides conceptual content to help you configure and customize SPB-PIM Gateway (SPB-PIM GW) on the switch.

IP Multicast over Fabric Connect in Protocol Independent Multicast Networks

IP Multicast over Fabric Connect provides simplicity in provisioning and deploying IP multicast bridging and routing. Also, due to the fact that only one control plane protocol (IS-IS) exists, convergence times in the event of a network failure, are typically sub second.

IP Multicast over Fabric Connect

IP Multicast over Fabric Connect introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. IP Multicast over Fabric Connect uses these extensions, along with the Internet Group Management Protocol (IGMP) Snooping and Querier functions at the edge of the SPBM cloud, to create sub-trees of the VSN SPB for each multicast group to transport IP multicast data.

With IP Multicast over Fabric Connect, the switch supports the following:

- Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network (Layer 2 VSN with IP Multicast over Fabric Connect). Example application: Multicast in data centers.
- IP multicast routing support for IP Shortcuts using SPBM in the core and IGMP on the access (IP Shortcuts with IP Multicast over Fabric Connect). Example applications: Video surveillance, TV/Video/Ticker/Image distribution, VX-LAN.
- Layer 3 Virtual Services Network with VRF based routing support for IP Multicast over Fabric Connect in the core and IGMP on the access (Layer 3 VSN with IP Multicast over Fabric Connect). Example applications: Video surveillance, TV/Video/Ticker/Image Distribution, VXLAN, Multi-tenant IP multicast.



Important

Sources must be IGMP enabled to support discovery functions specific to the multicast applications in use.

IP Multicast over Fabric Connect Restrictions

- IP Multicast over Fabric Connect cannot connect to an IP Multicast router outside the SPB network.
- You can only deploy IP Multicast over Fabric Connect in environments where there are no multicast routers between the edge of the SPB network and the IP Multicast hosts that connect to the network.
- An existing network which is Protocol Independent Multicast (PIM) based cannot participate in the SPB network either by connecting to SPB originated streams or by injecting PIM network streams into the SPB network.
- In certain environments it is not possible to deploy an SPB network all the way to the point where the SPB network directly connects to an IGMP edge.

You encounter these restrictions during the following typical deployment scenarios:

- **Scenario 1:** You deployed IP Multicast using PIM and want to expand the network by deploying SPB for the new portion of the network. You want multicast applications to work across the old and new portion of the network.
- **Scenario 2:** Multicast traffic is exchanged between independent network operators at the boundary between their networks. PIM is the multicast routing protocol. A network operator wants to upgrade or replace the existing network to an SPB network. The inter-domain multicast traffic exchanges with other networks should not be disrupted.

The following figure shows the traditional Multicast over Fabric Connect environment with no PIM routers.




Figure 225: IP Multicast over Fabric Connect streams

In the previous figure, sources and receivers on the edges of the SPB network are IGMP hosts and sources of multicast data. Hence, the traditional Multicast over Fabric Connect host-to-host deployment works.

The following figure shows the traditional Multicast over Fabric Connect environment with PIM routers.




Figure 226: IP Multicast over Fabric Connect Streams

In the previous figure, sources and receivers on the edges of the SPB network, which are IGMP or source hosts can communicate over the SPB network. Sources and receivers connected to PIM routers cannot participate in the SPB network.

SPB-PIM Gateway

Multicast over Fabric Connect cannot connect to a PIM router that is external to the SPB network. When a receiver joins the SPB network for a specific group, the receiver must receive multicast streams in the neighboring multicast domains (PIM network). Similarly, a receiver in the neighboring multicast domain (PIM network) must receive multicast streams from sources in the SPB network. SPB-PIM Gateway (SPB-PIM GW) provides multicast inter-domain communication between an SPB network and a PIM network. SPB-PIM GW accomplishes this inter-domain communication across a special gateway VLAN. The gateway VLAN communicates with the PIM network through a subset of the full protocol messaging required for RFC 4601 compliance of a PIM interface, and translates the PIM network requirements into SPB language and vice versa.

SPB-PIM GW provides the following functionality:

- One or more SPB domains can share streams with one or more PIM domains.

- SPB-PIM GW can connect two independent SPB domains. The independent SPB domains connected by SPB-PIM GW share a subset of multicast streams without a PIM network in between.

SPB-PIM GW is supported in the GRT and in VRFs.

Multicast over Fabric Connect with SPB-PIM GW

In a Multicast over Fabric Connect environment with SPB-PIM GW, the SPB network connects sources and receivers from one or more PIM networks. The multicast traffic is then delivered across the domain boundaries through a path that transports the multicast traffic.

Multicast over Fabric Connect with SPB-PIM GW functionality consists of SPB nodes, which act as SPB-PIM Controller nodes and SPB-PIM Gateway nodes. The SPB Controller uses the Multicast Source Discovery Protocol (MSDP) to discover foreign sources.

The following figure shows the Multicast over Fabric Connect environment with SPB-PIM GW.

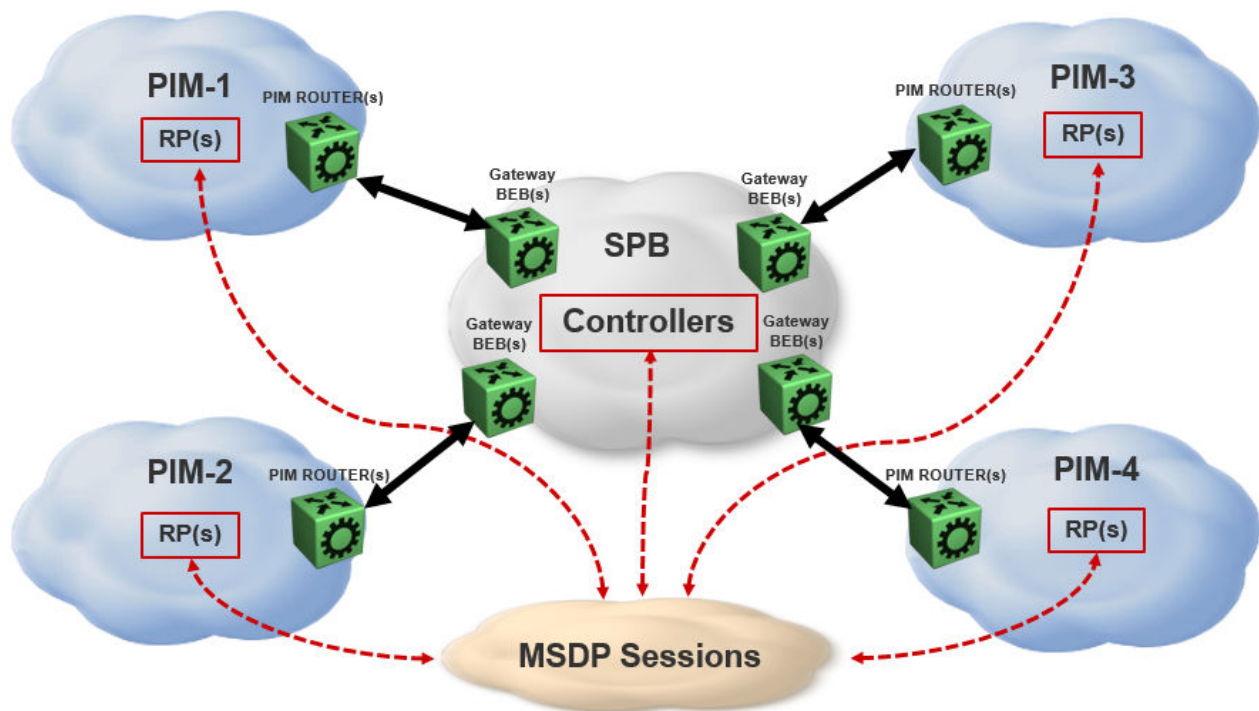


Figure 227: Multicast over Fabric Connect with SPB-PIM GW

SPB-PIM GW Components

SPB-PIM GW has two functional components:

- SPB-PIM Gateway Controller Nodes (Controller), which are used for multicast source discovery.
- SPB-PIM Gateway Nodes (Gateway), on which the SPB-PIM Gateway interfaces reside.



Note

The Controller and Gateway can reside in a single node.

The following figure shows the SPB-PIM GW components.

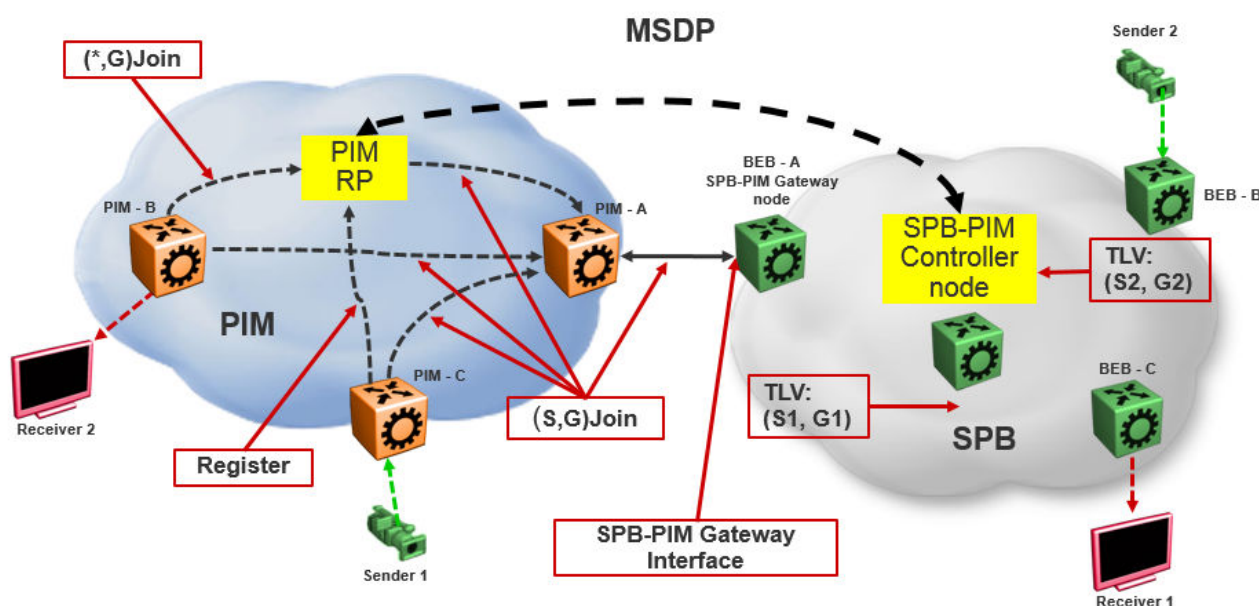


Figure 228: SPB-PIM GW components

SPB-PIM Gateway Controller Node

SPB-PIM Gateway Controller Node (Controller) shares stream information between the local SPB domain and a foreign domain. The foreign domain is the PIM network Rendezvous point (RP) or another SPB domain Controller.

The following list outlines Controller functionality:

- The Controller discovers PIM sources for a specific multicast group and distributes them to the Gateways.

**Note**

PIM source discovery is either through MSDP or static configuration of foreign streams at the Controller.

- The Controller advertises local SPB originated streams through MSDP to another PIM domain or another SPB domain.
- The Controller references the Unicast IP route table to determine which Gateway has the best route to the PIM source. The Controller then assigns the stream to the selected Gateway.

The Controller Node has the following components:

- Source Discovery (MSDP and static configuration)
- Gateway Selection Controller

Source Discovery (MSDP and Static Configuration)

MSDP resides in the Controller BEB or Controller BCB and PIM network RP that want to advertise multicast source information between domains.

You can implement SPB-PIM GW under the following scenarios:

- The multicast source resides in the Protocol Independent Sparse Module (PIM-SM) domain. The multicast source must be discovered by MSDP residing on the Gateway Controller in the SPB domain.
- The multicast source resides in the SPB domain. The multicast source must be advertised to the neighboring PIM domain through MSDP peers.

For more information on MSDP, see [MSDP Overview](#) on page 3567.

**Note**

You can also configure multicast sources statically on the Controller. Static configuration is useful for SSM multicast group range streams in the foreign domain, which are not advertised by MSDP. Static configuration is also useful for when two SPB domains are connected through a PIM Gateway, and want to only advertise a subset of streams to each other, without enabling MSDP.

Gateway Selection Controller

The Gateway Selection Controller resides in the Controller BEB or Controller BCB node in the SPB network. The Gateway Selection Controller receives source information from MSDP or through static configuration. The source information consists of the following components:

- Sender IP address (S)
- Group IP address (G)
- VRF ID of the stream
- RP of the source (optional)

Gateway Selection Controller finds the best BEB (Gateway Node) in the SPB network through which the sender sends traffic to group G. The Gateway Selection Controller performs the following tasks:

- The Gateway Selection Controller uses Layer 3 reachability information to reach S, which is retrieved from the ISIS IP Shortcuts (IPSC) database.
- The VSN identifier (I-SID) is determined by using the VRF ID provided by MSDP.
- The Gateway Selection Controller uses the VRF ID and searches the IP shortcut database to determine which Gateway is closest to S.

**Note**

If multiple BEBs have a route to S, the BEB with the lowest Layer 3 metric is selected as the Gateway.

**Note**

If a Gateway link fails or the cost of the route changes, the selection process identifies the link failure as a route change and selects another best Gateway BEB.

The selected Gateway BEB for a stream must satisfy the following criteria:

- The selected Gateway BEB for a stream must announce a route to the source of the foreign stream through ISIS.

**Note**

Among all the routes to the source of the foreign stream announced by different BEBs through ISIS, the route announced by the selected Gateway has the longest prefix match and has the lowest external route metric.

- If multiple BEBs meet the Gateway selection criteria, a deterministic hash function of system ID, source IP address, and group IP address is used. The deterministic hash function is computed for each of the BEBs that meet the Gateway selection criteria. The BEB that generates the lowest hash value is selected as the Gateway for the stream.

The result of the Gateway selection process is saved in the Gateway assignment table. The Gateway assignment table consists of VSN identifier or I-SID, S, G, and the selected Gateway BEB. Having only one selected Gateway BEB ensures that traffic from source S is drawn into the SPB network by only one BEB, the selected Gateway BEB. The selection Controller then distributes the Gateway assignment table information to all the Gateway nodes.

SPB-PIM Gateway Node

The SPB-PIM Gateway Node (Gateway) has the following components:

- Gateway Selection Agent
- SPB-PIM Gateway interface

Gateway Selection Agent

The Gateway Selection Agent (Agent) resides in the Gateway BEB Node in the SPB network. The Gateway BEB has connections into the foreign network over SPB-PIM

Gateway Interfaces. The Agent receives foreign network source information from the Controller BEB or the Controller BCB Node. The source information consists of the following components:

- Sender IP address (S)
- Group IP address (G)
- VSN identifier (I-SID)
- Gateway assigned to the stream

The Agent interacts with SPB-PIM Gateway interface and creates the multicast path. The Agent receives foreign source information from the Controller and creates a foreign source address (SA) cache after validating the reachability to S. The Agent interacts with the SPB-PIM Gateway interface to validate that the next-hop ip address toward the source is a valid PIM adjacency. The foreign SA cache includes the following components:

- Source IP address (information received from the Gateway Controller)
- Group IP address (information received from the Gateway Controller)
- Ingress port (The port through which S is accessible)
- Upstream IP address (The next-hop IP address, which is also the PIM neighbor across the SPB-PIM Gateway interface which is used to reach S as indicated by the unicast routing entry)
- Ingress VLAN ID

If multiple next-hops are available, then the first valid PIM neighbor next-hop is used for the upstream.

**Note**

If the Agent receives the same source information from multiple Controllers, then the Agent takes action only for the information received from the preferred Controller. The Controller with the lowest system ID is the preferred Controller.

SPB-PIM Gateway Interface

The SPB-PIM Gateway interface provides inter-domain multicast services. The SPB-PIM Gateway interface connects senders and receivers of multicast streams across a PIM Domain and a SPB network boundary over a Gateway interface. The SPB-PIM Gateway interface provides the following functionality:

- PIM HELLO exchanges
- Issuing Joins and Leaves
- Process received Joins and Leaves
- Implements the Gateway assignment table by acting as the Ingress BEB for streams for which the SPB-PIM Gateway interface is the selected Gateway
- Enforces the Gateway assignment table and does not forward streams for which the SPB-PIM Gateway interface is not the selected Gateway

- Forwards local and remote SPB streams to satisfy stream requests from neighboring multicast domains
- SPB-PIM Gateway Interfaces supports both SM and SSM multicast group range joins and prunes. *G joins are only supported in SM group range.

The PIM Gateway interface resides in the SPB-PIM Gateway Node (Gateway). The SPB-PIM Gateway interface connects to a PIM router in a PIM network or to another Gateway BEB in an SPB network. Local hosts (IGMP member hosts and multicast data source hosts) are not supported on SPB-PIM Gateway interfaces, only PIM Routers or another SPB BEB with SPB-PIM Gateway interface configured. Multicast data from local source hosts and IGMP reports from local hosts are dropped. An SPB Node must be configured as a SPB-PIM Gateway Node if the SPB Node is connected to a foreign PIM network or a foreign SPB network. A single Gateway Node can have multiple SPB-PIM Gateway interfaces. The SPB-PIM Gateway interface can be a VLAN or a router port, can reside on an MLT and is fully virtualized. The SPB-PIM Gateway interface is a translation mechanism between the PIM protocol and SPB TLVs.

**Note**

- Only PIM protocol messages are communicated over the SPB-PIM Gateway interface
- Only SPB TLVs are communicated over Fabric Connect over SPB
- The SPB-PIM Gateway interface is the only component that handles the translation mechanism

The SPB-PIM Gateway interface communicates with the PIM router through the standard PIM protocol messaging HELLO, JOIN, and PRUNE. The SPB-PIM Gateway interface then forms a normal PIM adjacency with the PIM router or another SPB Gateway Node. The SPB-PIM Gateway Interface processes received SG joins and prunes, *G joins and prunes, and SG-RPT joins and prunes. The SPB-PIM Gateway interface transmits SG joins and prunes, but never *G joins. The SPB-PIM Gateway Interface does not have RP capabilities, and therefore has no need for group-to-RP mapping configurations. A *G JOIN received on a SPB-PIM Gateway Interface is accepted if the destination IP is the IP address of the interface or of a neighbor on the interface if the neighbor is learned on another port in the interface. However, the RP address within the *G JOIN message is ignored by the SPB-PIM Gateway Interface.

MSDP Overview

MSDP enables advertisement of multicast source information between different PIM-SM domains. This function of MSDP in SPB-PIM GW topologies is to advertise multicast source information between SPB domains and PIM domains. MSDP routers in a PIM-SM or SPB domain have a peering relationship with MSDP peers in another domain. The peering relationship is a TCP connection in which the control information is exchanged. The TCP connection between peers uses the underlying unicast routing system.

Source Active Messages

In a PIM domain, MSDP enabled routers are RPs. MSDP routers form adjacencies through TCP port 639 to share multicast source information. This functionality is similar to the Border Gateway Protocol (BGP). When an MSDP router receives multicast source information, the routers use reachability information to perform Reverse Path Forwarding (RPF) checks. The reachability information is exchanged through BGP or any other unicast routing protocol.

When a RP router learns of a new (S,G), the RP router saves the (S,G) information and the RP address in the MSDP Source Active (SA) local cache. The RP router learns the new (S,G) through a directly connected source or PIM register message. The RP router then sends an SA update message which contains (S,G,RP) information to the MSDP peers. The MSDP peers broadcast the SA to RPs in their local domains and to their MSDP peers in other PIM-SM domains.

**Note**

A PIM domain is a set of routers in a single Autonomous System (AS), which uses the same RP for any given multicast group.

When an SPB-PIM Gateway Controller in the SPB domain learns of a new (S,G) in its own domain, the Controller saves the (S,G) information in the local SA cache. The Controller learns the new (S,G) through a directly connected source or Intermediate-System-to-Intermediate-System (IS-IS). The controller sends an SA update message to the MSDP peers in the PIM domain. SA uses the CLIP address configured on the controller as the RP address.

**Note**

Configure CLIP before you enable MSDP. Peer connections use the CLIP address as the local address.

Reverse Path Forwarding Check

When an MSDP peer receives the SA from a peer, the MSDP performs an RPF check. The RPF check ensures that the SA received from the MSDP peer is the closest to the originating RP. An RPF check prevents SA loops.

**Note**

This RPF check is different from the multicast routing RPF check.

If the RPF checks pass, then the receiving MSDP enabled router saves the SA information in the SA foreign cache and makes it available to the local domain. Each MSDP peer floods the SA information away from the originating RP. The flooding process is called peer RPF flooding.

SA Redistribution and Filtering

Redistribution and filtering is used to control SA flooding. The MSDP redistribute policy is applied on the MSDP node that originates the SAs to control which SAs are advertised on all MSDP peers. An SA filter is applied to a specific MSDP peer in the

inbound direction or outbound directions or both inbound and outbound direction on any MSDP router. Filtering is multicast group based.

When configuring MSDP redistribution, use prefix lists to create the route policies. When a route policy is created it must match the group prefix with the name of the prefix list created for the group address. If deny action is set for the lists in the route-policy, the policy blocks the matching groups from all the sources. If permit action is set for the lists in the route-policy, the policy accepts the matching groups from all the sources. MSDP redistribution does not refer to the redistribution of SPB domain sources to MSDP. MSDP redistribution refers to SAs which needs to be redistributed to other MSDP peers.

**Note**

MSDP redistribution is applied globally to all MSDP peers. SA filtering is used to filter SAs on a peer-to-peer level.

MSDP and SPB-PIM GW

The SPB-PIM GW functional component for MSDP resides in the SPB-PIM Gateway Controller node (Controller).

Overview

Controllers from an SPB network discover sources through MSDP sessions with RPs from a PIM network. Once the SA packed is received at the MSDP module of the Controller, the IP routing table is examined to determine which peer is the next hop towards the originating RP of the SA message. Once the SA RPF test passes, the SA packet is saved in the foreign cache and passed to the Controller.

The Controller nodes also distribute the sources from an SPB domain to a PIM domain through an MSDP session with RPs in the PIM network. Similar to RP, when a Controller in the SPB domain learns of a new (S,G), through a directly connected source or ISIS, the Controller saves the new (S,G) in its local SA cache. The Controller then transmits an SA update message for this source to its MSDP peers in the PIM domain. The Controller that sends the SA to the MSDP is viewed as the RP (circuitless IP interface is used).

MSDP as Ppart of SPB-PIM GW

MSDP does not work with the traditional PIM implementation. MSDP communicates only with the Controller and should be configured as an IP endpoint.

MSDP configuration considerations:

- A circuitless IP interface (CLIP) is used in the context of global router or VRF, hence at least one CLIP in each VRF should be configured.
- MSDP should use a single CLIP address as the source for establishing all MSDP connections in the same VRF.
- For SPB sources, this CLIP is used as the RP in all SA messages advertised.

- MSDP source IP address should be one of the CLIP interfaces pre-configured on the global router or VRF.
- The originator-id should be configured before enabling MSDP.

**Note**

MSDP transmits encapsulated multicast data packets inside forwarded MSDP messages. If the received SA is an encapsulated SA, then the switch parses the TTL value of the encapsulated data and compares it against the configured value. If the configured value is less than or equal to the parsed value, then the switch forwards the encapsulated data along with the SA, otherwise the switch forwards the SA alone by stripping the encapsulate data. By default, MSDP forwards encapsulated data along with the SA message. MSDP does not forward the encapsulated data to the local receivers.

When MSDP generates SA messages for SPB sources, the local cache miss data cannot be encapsulated into the SA messages that are sent to the peers.

The switch supports forwarding SA messages with encapsulated data from sources to MSDP peers but not from MSDP peers to the receivers.

For more information on MSDP configuration, see [Multicast Source Discovery Protocol Configuration](#) on page 3571.

Full Mesh Group

MSDP mesh groups are full mesh of MSDP peers and is a subset of MSDP speakers. MSDP mesh groups are used for SA flooding which is similar to the BGP route reflector concept. MSDP floods the SA to all the members of the mesh group when:

- The MSDP peers are fully meshed
- The MSDP enabled router learns a new SA from a non-member of its mesh group
- The SA passes the RPF check

The receiving routers accept the SA and forwards it only to any non-mesh Group MSDP peers.

The SPB-PIM Gateway is deployed in two models:

- Model 1: All multicast networks have peering agreements with one another. The full mesh MSDP is setup.
- Model 2 : An inter-domain multicast provider exists. All multicast networks setup MSDP peering with the provider.

The controllers relay SA messages between individual networks.

**Note**

Consider the following when you deploy SPB-PIM Gateway:

- Use mesh group of MSDP peers (PIM RP's and SPB-PIM Gateway Controller nodes) to avoid flooding and RPF failure.

**Note**

Since MSDP uses CLIP interface in its peering relation, the MSDP peer may not fall in any of the RFC rules and the MSDP SA messages will be rejected.

- Controllers from the same SPB network must not have MSDP sessions with each other, regardless of whether mesh groups are used or not.
- When using mesh groups, all Controllers within one SPB domain should peer with the same set of RPs and Controllers in adjacent domains, ie, one Controller should not peer with an RP that the other Controllers do not peer with.

Multicast Source Discovery Protocol Configuration

This section provides procedures to configure Multicast Source Discovery Protocol (MSDP) using the Command Line Interface (CLI) and Enterprise Device Manager (EDM).

Basic MSDP Configuration Using CLI

This section provides procedures to configure Basic MSDP using the command line interface (CLI).

Configure the MSDP Originator ID

About This Task

Configure the originator ID to set the Rendezvous Point (RP) address inside the Source Active (SA) message. The RP address must be a pre-configured CLIP interface on the global router or a VRF. The RP address is also the local IP address in all peer relations.

**Note**

To delete the originator ID, you must first disable MSDP.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the MSDP originator ID:

```
ip msdp originator-id {A.B.C.D}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp originator-id 2.0.2.2
```

Variable Definitions

The following table defines parameters for the **ip msdp originator-id** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP source IP address. The IP address must be one of the CLIP interfaces configured on the global router or a VRF.

Configure MSDP on a VRF

About This Task

Create an MSDP instance on a user defined VRF to allow further configuration to take place. This command does not exist in the Global Configuration mode because the MSDP instance for a default VRF is created by default.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the MSDP instance:

```
ip msdp
```

Enable MSDP Globally

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Enable MSDP globally on the switch:

```
ip msdp enable
```

Create an MSDP Peer

About This Task

Perform the following procedure to establish a peer relationship between the local MSDP enabled router and a peer in another domain.



Important

Do not enable more than 20 active peers.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional:

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create an MSDP peer:

```
ip msdp peer {A.B.C.D}
```
3. Enable an MSDP peer:

```
ip msdp peer {A.B.C.D} enable
```



Note

MSDP peer is disabled by default.

4. (Optional) Specify the remote autonomous system (AS) number of the MSDP peer:

```
ip msdp peer {A.B.C.D} remote-as WORD<0-11>
```

Example

```
Switch:1>enable  
Switch:1#configure terminal
```

```
Switch:1(config)#ip msdp peer 21.0.0.2
Switch:1(config)#ip msdp peer 21.0.0.2 enable
Switch:1(config)#ip msdp peer remote-as 1
```

Variable Definitions

The following table defines parameters for the **ip msdp peer** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP peer IP address.
WORD<0-11>	Specifies the AS number of the MSDP peer, 0-65535 (2-Byte AS) 0-4294967295 (4-Byte AS).

MSDP Peer Configuration Using CLI

This section provides procedures to configure MSDP peer using the command line interface (CLI).

Configure a Peer Description

About This Task

Configure a peer description to add descriptive text to an MSDP peer for easy identification of a peer.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the peer description:

```
ip msdp description {A.B.C.D} WORD<1-255>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp peer 21.0.0.2 primary
```

Variable Definitions

The following table defines parameters for the **ip msdp description** command.

Variable	Value
<code>{A.B.C.D}</code>	Specifies the MSDP peer IP address.
<code>WORD<1-255></code>	Specifies a descriptive text to a MSDP peer in the range of 1-255 characters. To include spaces in the peer description, enclose the text string in quotation marks.

Secure Control Messages

About This Task

Configure Message Digest (MD) 5 authentication to secure control messages on the TCP connection between MSDP peers.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Enable MD5 authentication:

```
ip msdp md5-authentication {A.B.C.D} [enable]
```

3. Specify the case sensitive password for MD5 authentication:

```
ip msdp password peer {A.B.C.D} WORD<1-80>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp md5-authentication 21.0.0.2 enable
Switch:1(config)#ip msdp password peer 21.0.0.2 helloworld
```

Variable Definitions

The following table defines parameters for the **ip msdp** command.

Variable	Value
<code>{A.B.C.D}</code>	Specifies the MSDP peer IP address.
<code>WORD<1-80></code>	Specifies the MD5 authentication password.

Configure the MSDP Peer SA Limit

About This Task

Configure the SA limit to limit the number of SA messages from an MSDP peer that the router saves in the SA cache. The default value is 6,144 messages.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the SA limit:

```
ip msdp sa-limit {A.B.C.D} <0-6144>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp sa-limit 21.0.0.2 6100
```

Variable Definitions

The following table defines parameters for the `ip msdp sa-limit` command.

Variable	Value
<code>{A.B.C.D}</code>	Specifies the MSDP peer IP address.
<code><0-6144></code>	Specifies the maximum number of SA messages to keep in SA cache.

Configure the Time-To-Live Threshold

About This Task

Configure the time-to-live (TTL) threshold to limit which multicast data packets the router encapsulated in SA Messages forwarded to an MSDP peer. The TTL limits the number of hops a packet can take before the router drops the packet. The router sends out SA Messages with encapsulated data only if TTL equals or exceeds the value you

configure. If the TTL is lower than the value you configure, the router drops the data packet and forwards the SA Message without the encapsulated data.



Note

MSDP transmits encapsulated multicast data packets inside forwarded MSDP messages. If the received SA is an encapsulated SA, then the switch parses the TTL value of the encapsulated data and compares it against the configured value. If the configured value is less than or equal to the parsed value, then the switch forwards the encapsulated data along with the SA, otherwise the switch forwards the SA alone by stripping the encapsulate data. By default, MSDP forwards encapsulated data along with the SA message. MSDP does not forward the encapsulated data to the local receivers.

When MSDP generates SA messages for SPB sources, the local cache miss data cannot be encapsulated into the SA messages that are sent to the peers.

The switch supports forwarding SA messages with encapsulated data from sources to MSDP peers but not from MSDP peers to the receivers.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the MSDP peer TTL threshold:

```
ip msdp ttl-threshold {A.B.C.D} <1-255>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp ttl-threshold 21.0.0.2 10
```

Variable Definitions

The following table defines parameters for the **ip msdp ttl-threshold** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP peer IP address.
<1-255>	Specifies the TTL value. Default value is 1.

*Configure the MSDP Peer Keep Alive Messages***About This Task**

Configure keep alive messages to adjust the interval in seconds at which an MSDP peer sends keep alive messages (default is 60 seconds) and the interval at which the MSDP peer waits for keep alive messages from other peers before it declares them down (default is 75 seconds).

**Note**

In a peer relationship, the keep alive interval configured on one peer must be at least 1 second less than the hold time configured on the other side of the peer relationship. This is not applicable when the hold time is set to 0 seconds.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`

**Note**

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the MSDP peer keep alive interval:

```
ip msdp keepalive {A.B.C.D} <0-21845> <0-65535>
```

Example

```
Switch:1>enable
Switch:1#configure termina
Switch:1(config)#ip msdp keepalive 21.0.0.2 70 71
```

Variable Definitions

The following table defines parameters for the **ip msdp keepalive** command.

Variable	Value
<code>{A.B.C.D}</code>	Specifies the MSDP peer IP address.
<code><0-21845></code>	Specifies the keep alive interval in seconds. The default is 60 seconds.
<code><0-65535></code>	Specifies the hold time interval in seconds. The default is 75 seconds. Note: 0 seconds means the peer never expires. Values 1 and 2 are not allowed.

Configure the MSDP Peer Connect-Retry Period

About This Task

Configure the connect-retry period to specify the amount of time, in seconds, between connection attempts for peering sessions. The default is 30 seconds.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the MSDP peer connect-retry period:

```
ip msdp connect-retry {A.B.C.D} <1-65535>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp connect-retry 21.0.0.2 40
```

Variable Definitions

The following table defines parameters for the **ip msdp connect-retry** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP peer IP address.
<1-65535>	Specifies the connect-retry interval in seconds. The default is 30 seconds.

Clear the Peer Connection

About This Task

Clear the peer connection to clear the TCP connection to the specified MSDP peer and reset all MSDP message counters.



Note

This procedure does not clear the SA cache entries the router learns from the peer.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear the peer connection:

```
clear ip msdp peer {A.B.C.D} [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#clear ip msdp peer 21.0.0.2
```

Variable Definitions

The following table defines parameters for the **clear ip msdp peer** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP peer IP address.
vrf WORD<0-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

Delete an MSDP Peer

Procedure

1. Enter Privileged EXEC mode:
enable
2. Clear the peer connection:
no ip msdp peer {A.B.C.D}

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#no ip msdp peer 21.0.0.2
```

MSDP Message Control Using CLI

Filter PIM Routes

About This Task

Filter SPB routes to filter which (S,G,RP) entries sent out to all MSDP peers. This procedure applies only to the rendezvous point (RP) that originates the MSDP SA messages and not the intermediate MSDP peers that forward the received SA messages.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the MSDP filter:

```
ip msdp redistribute
```

3. Create the route policy name:

```
ip msdp redistribute route-policy WORD<1-64>
```

4. Apply the redistribution filters:

```
ip msdp apply redistribute
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp redistribute
Switch:1(config)#ip msdp redistribute route-policy helloworld
Switch:1(config)#ip msdp apply redistribute
```

Variable Definitions

The following table defines parameters for the **clear ip redistribute** command.

Variable	Value
<i>WORD<1-64></i>	Specifies the route policy name.

Filter SA Messages

About This Task

Filter SA messages to determine which SA messages to accept from a peer and which SA messages to send to a peer. By default, no inbound or outbound filter exists.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create the inbound filter:

```
ip msdp sa-filter in {A.B.C.D}
```

3. Create the inbound filter route policy name:

```
ip msdp sa-filter in {A.B.C.D} route-policy WORD<1-64>
```

4. Create the outbound filter:

```
ip msdp sa-filter out {A.B.C.D}
```

5. Create the outbound filter route policy name:

```
ip msdp sa-filter out {A.B.C.D} route-policy WORD<1-64>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp sa-filter in 21.0.0.2 route-policy helloworld
```

Variable Definitions

The following table defines parameters for the **ip msdp sa-filter** command.

Variable	Value
<code>{A.B.C.D}</code>	Specifies the MSDP peer IP address.
<code>route-policy WORD<1-64></code>	Specifies the route policy name for an inbound or outbound filter.

Configure MSDP Mesh Groups

About This Task

Configure mesh groups to reduce SA flooding. A mesh group does not forward SA messages to other group members. The originator, which is also a mesh group member, forwards SA messages to all group members. Create MSDP mesh groups from a group of meshed MSDP speakers from a domain.



Note

The MSDP router does not belong to any mesh group by default.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

Optional: `router vrf WORD<1-16>`



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure the MSDP mesh group:

```
ip msdp mesh-group WORD<1-64> {A.B.C.D}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip msdp mesh-group helloworld 21.0.0.2
```

Variable Definitions

The following table defines parameters for the **ip msdp mesh-group** command.

Variable	Value
<i>WORD<1-64></i>	Specifies the mesh group name.
<i>{A.B.C.D}</i>	Specifies the MSDP peer IP address.

Clear the MSDP SA Cache

About This Task

Clear the SA cache to clear the SA entries the router learns from all peers or a specific peer.



Note

This procedure clears the foreign cache. This procedure does not clear the local cache.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear the SA cache for all peers:

```
clear ip msdp sa-cache [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

3. Clear the SA cache for a specific peer:

```
clear ip msdp sa-cache peer {A.B.C.D} [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

4. Clear the SA cache for a specific group range, source range, and RP.

```
clear ip msdp sa-cache [source prefix/len] [group prefix/len] [rp
{A.B.C.D}] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#clear ip msdp sa-cache peer 21.0.0.2
```

Variable Definitions

The following table defines parameters for the **clear ip msdp sa-cache** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP peer IP address.
vrf WORD<0-16>	Specifies the VRF names.
vrfids WORD<0-512>	Specifies the VRF ID.

MSDP Verification Using CLI

This section provides procedures to verify MSDP using the command line interface (CLI).

Display the Peer Information

About This Task

Use the following procedure to display the peer configuration and SA message information.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the peer information:
show ip msdp peer {A.B.C.D} [vrf WORD<0-16>] [vrfids WORD<0-512>]

Example

```
Switch:#enable
Switch:1#show ip msdp peer 2.2.2.2
=====
MSDP Peer - GlobalRouter
=====
MSDP Peer 2.2.2.2, AS 109Admin Status: Enabled
Operational Status: Enabled
Description:
Connection status:
FSM State: Established, Establish Count: 9,
Connection source: 2.2.2.17
Uptime (Downtime): 1d10h, Messages sent/received:
436765/429062
Connection and counters cleared 1w2d ago
SA Filtering:
Input (S,G) route-policy: none
```

```

Output (S,G) route-policy: none
SA In count: SA out Count:
SA-Requests:
Input filter: none
Sending SA-Requests to peer: disabled
SA Request In Count: SA Request out Count:
SA Response In Count: SA Response out Count:
Peer ttl threshold: 0
SAs learned from this peer: 32, SAs limit: 500
Peer RPF failure Count:
KeepAlive In Count:
KeepAlive out count:
Encapsulated Data packets In:
Encapsulated Data Packets out:
KeepAlive Timer:
Peer Hold timer:
Connection Retry timer:
Encapsulation type:
MD5 Authentication: Enabled, MD5 Password:
%d462277d77
Peer FSM Established Time:
Peer In Message Time:
Remote port: Local port:
Number of connection Attempts:
Discontinuity timeout:
Too Short MSDP message Rx count:
Bad MSDP message Rx count:

```

Variable Definitions

The following table defines parameters for the **show ip msdp peer** command.

Variable	Value
{A.B.C.D}	Specifies the MSDP peer IP address.
vrf WORD<0-16>	Displays configuration info for a particular VRF.
vrfids WORD<0-512>	Displays configuration info for a particular VRF ID.

Display the SA Cache

About This Task

Use the following procedure to display the (S,G) state learned from MSDP peers and the local (S,G) state. The local (S,G) is the SPB (S,G) sent to MSDP.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the SA cache:

```
show ip msdp sa-cache [local] [vrf WORD<0-16>] [vrfids WORD<0-512>]
[group {A.B.C.D}] [rp {A.B.C.D}] [source {A.B.C.D}]
```

Example

```
Switch:#enable
Switch:1#show ip msdp sa-cache local
=====
MSDP Foreign SA Cache - GlobalRouter
=====
MSDP Source-Active Foreign Cache - 8 entries
(2.10.1.100, 224.5.5.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.10.1.100, 224.5.6.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.10.1.100, 224.5.7.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.10.1.100, 224.5.8.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.11.2.100, 224.6.5.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.11.2.100, 224.6.6.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.11.2.100, 224.6.7.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35
(2.11.2.100, 224.6.8.0), RP 3.3.3.3, BGP/AS 1,
00:01:53/00:05:35

Switch:1#show ip msdp vrf msdpvrf
=====
MSDP Local SA Cache - VRF msdpVrf
=====
MSDP Source-Active Local Cache - 12 entries
(5.12.5.100, 224.7.5.0), RP 5.5.5.5
(5.12.5.100, 224.7.6.0), RP 5.5.5.5
(5.12.5.100, 224.7.7.0), RP 5.5.5.5
(5.12.5.100, 224.7.8.0), RP 5.5.5.5
(5.13.7.100, 224.8.5.0), RP 5.5.5.5
(5.13.7.100, 224.8.6.0), RP 5.5.5.5
(5.13.7.100, 224.8.7.0), RP 5.5.5.5
(5.13.7.100, 224.8.7.0), RP 5.5.5.5
(5.13.7.100, 224.8.7.0), RP 5.5.5.5
(7.14.8.100, 224.9.6.0), RP 5.5.5.5
(7.14.8.100, 224.9.7.0), RP 5.5.5.5
(7.14.8.100, 224.9.8.0), RP 5.5.5.5
```

Variable Definitions

The following table defines parameters for the **show ip msdp sa-cache** command.

Variable	Value
<i>group</i> {A.B.C.D}	Displays all SA cache entries that match the group IP address.
<i>local</i>	Displays the local SA cache.
<i>rp</i> {A.B.C.D}	Displays all SA cache entries that match the RP IP address.
<i>source</i> {A.B.C.D}	Displays all SA cache entries that match the source IP address.

Variable	Value
<code>vrf WORD<0-16></code>	Displays configuration information for a particular VRF.
<code>vrfs WORD<0-512></code>	Displays configuration information for a particular VRF ID.

Display the MSDP Count

About This Task

Use the following procedure to display the number of sources and groups sent and received.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the MSDP count:
show ip msdp count [`vrf WORD<0-16>`] [`vrfs WORD<0-512>`] [`<0-65535>`]

Example

```
Switch:#enable
Switch:1#show ip msdp count
=====
MSDP Count - GlobalRouter
=====
SA state per peer Counters, <peer>: <# SA learned>
192.135.250.116: 24
144.228.240.253: 3964
172.17.253.19: 10
172.17.170.110: 11
SA state per ASN Counters, <asn>: <# SA-count>
Total entries: 4009
?: 192, 9: 1, 14: 107, 17: 5
18: 4, 25: 23, 26: 39, 27: 2
32: 19, 38: 2, 52: 4, 57: 1
68: 4, 73: 12, 81: 19, 87: 9
```

Variable Definitions

The following table defines parameters for the **show ip msdp count** command.

Variable	Value
<code><0-65535></code>	Specifies the AS number.
<code>vrf WORD<0-16></code>	Displays configuration information for a particular VRF.
<code>vrfs WORD<0-512></code>	Displays configuration information for a particular VRF ID.

Display the MSDP Summary

About This Task

Use the following procedure to display the MSDP global and peer status.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the MSDP summary:
show ip msdp summary [*vrf WORD<0-16>*] [*vrfids WORD<0-512>*]

Example

```
Switch:#enable
Switch:1#show ip msdp summary
=====
MSDP Summary - GlobalRouter
=====
MSDP Status Summary
  MSDP Global Status: enabled
  cache status: enabled
  cache-lifetime: 390 seconds
  cache-count: 8
  Originator id: 5.5.5.5
  Redistribute: route-policy:
  SA Limit: 6144

MSDP Peer Status Summary

Peer Address AS State      Uptime/  Established SA
                Downtime Count      Count
4.5.35.3      1  Established 00:00:27 3          8
5.7.56.7      2  Established 00:00:31 2          0
```

Variable Definitions

The following table defines parameters for the **show ip msdp summary** command.

Variable	Value
<i>vrf WORD<0-16></i>	Displays configuration information for a particular VRF.
<i>vrfids WORD<0-512></i>	Displays configuration information for a particular VRF ID.

Display the RPF Peer Information

About This Task

Use the following procedure to display the MSDP peer information for a specific RP. The SA messages are received from the MSDP peer.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the RPF peer information:
show ip msdp rpf {*A.B.C.D*} [*vrf WORD<0-16>*] [*vrfids WORD<0-512>*]

Example

```
Switch:#enable
Switch:1#show ip msdp rpf 172.16.10.13
=====
MSDP RPF - GlobalRouter
=====
RPF peer information for (172.16.10.13)
RPF peer: (172.16.121.10)
RPF route/mask: 172.16.0.0/255.255.0.0
RPF rule: Peer is IGP next hop of best route
RPF type: unicast (ospf)
```

Variable Definitions

The following table defines parameters for the **show ip msdp rpf** command.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the RP IP address.
<i>vrf WORD<0-16></i>	Displays configuration information for a particular VRF.
<i>vrfids WORD<0-512></i>	Displays configuration information for a particular VRF ID.

*Display the MSDP Mesh Group Information***About This Task**

Use the following procedure to display the configured mesh groups.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the MSDP mesh group information:
show ip msdp mesh-group [*vrf WORD<0-16>*] [*vrfids WORD<0-512>*]
[*WORD<1-64>*]

Example

```
Switch:#enable
Switch:1#show ip msdp mesh-group
=====
MSDP Mesh Group - GlobalRouter
=====
NAME                ADDRESS
-----
test                1.1.1.1
-----
```

Variable Definitions

The following table defines parameters for the **show ip msdp mesh-group** command.

Variable	Value
<code>vrf WORD<0-16></code>	Displays configuration information for a particular VRF.
<code>vrfids WORD<0-512></code>	Displays configuration information for a particular VRF ID.
<code>WORD<1-64></code>	Specifies the mesh group name.

Display the SA Check Information

About This Task

Use the following procedure to display the peer information from which the router accepts SA originating from the RP. The following procedure also checks if the specified (S,G,RP) will be accepted from the peer.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the SA check information:
show ip msdp sa-check source {A.B.C.D} group {A.B.C.D} rp {A.B.C.D}
[peer {A.B.C.D}] [vrf WORD<0-16>] [vrfids WORD<0-512>]

Example

```
Switch:#enable
Switch:1#show ip msdp sa-check source 10.10.10.1 group 225.1.1.1 rp 172.16.10.13 peer
3.3.3.1

MSDP SA Check - GlobalRouter
=====

RPF peer information for (172.16.10.13)
RPF peer: (172.16.121.10)
RPF route/mask: 172.16.0.0/255.255.0.0
RPF rule: Peer is IGP next hop of best route
RPF type: unicast (ospf)
(10.10.10.1, 225.1.1.1, 172.16.10.13) - SA Accepted

Switch:1#show ip msdp sa-check source 5.5.5.1 group 225.1.1.1 rp 172.16.10.13 vrf msdpvrf
=====

MSDP SA Check- VRF msdpVrf
=====

RPF peer information for (172.16.10.13)
RPF peer: (172.16.121.10)
RPF route/mask: 172.16.0.0/255.255.0.0
RPF rule: Peer is IGP next hop of best route
RPF type: unicast (ospf)
(5.5.5.1, 225.1.1.1, 172.16.10.13) - SA Filtered by IN
filter route-policy abc

Switch:1#show ip msdp sa-check source 5.5.5.1 group 225.1.1.1 rp 59.59.59.1 peer 3.3.3.1
=====

MSDP SA Check - GlobalRouter
=====

(5.5.5.1, 225.1.1.1, 172.16.10.13) - SA not accepted due
to RPF peer mismatch
```

Variable Definitions

The following table defines parameters for the **show ip msdp sa-check** command.

Variable	Value
<i>group</i> {A.B.C.D}	Specifies the group IP address.
<i>peer</i> {A.B.C.D}	Specifies the MSDP peer IP address.
<i>rp</i> {A.B.C.D}	Specifies the RP IP address.
<i>source</i> {A.B.C.D}	Specifies the source IP address.
<i>vrf</i> WORD<0-16>	Displays configuration information for a particular VRF.
<i>vrfids</i> WORD<0-512>	Displays configuration information for a particular VRF ID.

Display All MSDP Information

About This Task

Use the following procedure to display all the MSDP information.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display all MSDP information:
show ip msdp show-all [file WORD<1-99>] [vrf WORD<0-16>] [vrfids WORD<0-512>]

Example

```
Switch:#enable
Switch:1#show ip msdp show-all

=====
MSDP Show-all - GlobalRouter
=====

# show ip msdp count
SA State per Peer Counters, Peer: # SA learned
 4.5.35.3: 8
 5.7.56.7: 0
AS Num : SA Count
 1: 8

# show ip msdp mesh-group
No Mesh Group exists

# show ip msdp peer

MSDP Peer 4.5.35.3, AS 1
Admin Status : enabled
Operational Status : enabled
Description:
Connection status:
  FSM State: Established, Established Count: 3,
Connection source: 4.5.35.5
Uptime (Downtime): 00:00:20 ago, Messages
```

```
sent/received: 10839/174
  Connection and counters cleared 00:00:27 ago
SA Filtering:
  Input (S,G) route-policy:
  Output (S,G) route-policy:
  SA In count: 8 SA out Count: 10836
SA-Requests:
  Input filter: none
  Sending SA-Requests to peer: disabled
  SA Request In Count: 0 SA Request out Count: 0
  SA Response In Count: 0 SA Response out Count: 0
Peer ttl threshold: 0
SAs learned from this peer: 8, SAs limit: 6144
Peer RPF failure Count: 0
KeepAlive In Count: 166
KeepAlive out count: 3
Encapsulated Data packets In: 8
Encapsulated Data Packets out: 6152
KeepAlive Timer: 60
Peer Hold timer: 75
Connection Retry timer: 30
Encapsulation type: 6
MD5 Authentication: enable
Md5 password: %d462277d77
Peer FSM Established Time: 01:20:57
Peer In Message Time: 01:21:00
Remote port: 49156 Local port: 639
Number of connection Attempts: 0
Discontinuity timeout:01:20:50
Too Short MSDP message Rx count: 0
Bad MSDP message Rx count: 0

MSDP Peer 5.7.56.7, AS 2
Admin Status : enabled
Operational Status : enabled
Description:
Connection status:
  FSM State: Established, Established Count: 2,
Connection source: 5.7.56.5
  Uptime (Downtime): 00:00:27 ago, Messages
sent/received: 4677/77
  Connection and counters cleared 00:00:30 ago
SA Filtering:
  Input (S,G) route-policy:
  Output (S,G) route-policy:
  SA In count: 0 SA out Count: 4675
SA-Requests:
  Input filter: none
  Sending SA-Requests to peer: disabled
  SA Request In Count: 0 SA Request out Count: 0
  SA Response In Count: 0 SA Response out Count: 0
Peer ttl threshold: 0
SAs learned from this peer: 0, SAs limit: 6144
Peer RPF failure Count: 0
KeepAlive In Count: 77
KeepAlive out count: 2
Encapsulated Data packets In: 0
Encapsulated Data Packets out: 8
KeepAlive Timer: 60
Peer Hold timer: 75
Connection Retry timer: 30
Encapsulation type: 6
MD5 Authentication: disable
Md5 password:
```

```

Peer FSM Established Time: 01:20:53
Peer In Message Time: 01:20:53
Remote port: 639 Local port: 49164
Number of connection Attempts: 3
Discontinuity timeout:01:20:50
Too Short MSDP message Rx count: 0
Bad MSDP message Rx count: 0

# show ip msdp sa-cache

MSDP Source-Active Foreign Cache - 8 entries
(2.10.1.100, 224.5.5.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07
(2.10.1.100, 224.5.6.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07
(2.10.1.100, 224.5.7.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07
(2.10.1.100, 224.5.8.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07
(2.11.2.100, 224.6.5.0), RP 3.3.3.3, BGP/AS 1,
00:00:23/00:06:06
(2.11.2.100, 224.6.6.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07
(2.11.2.100, 224.6.7.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07
(2.11.2.100, 224.6.8.0), RP 3.3.3.3, BGP/AS 1,
00:00:22/00:06:07

# show ip msdp summary
MSDP Status Summary
    MSDP Global Status: enabled
cache status: enabled
cache-lifetime: 390 seconds
cache-count: 8
Originator id: 5.5.5.5
Redistribute: route-policy:
SA Limit: 6144
MSDP Peer Status Summary

Peer Address AS State      Uptime/   Established SA
                        Downtime Count      Count
4.5.35.3      1  Established 00:00:27 3          8
5.7.56.7      2  Established 00:00:31 2          0

MSDP Source-Active Local Cache - 12 entries
(5.12.5.100, 224.7.5.0), RP 5.5.5.5
(5.12.5.100, 224.7.6.0), RP 5.5.5.5
(5.12.5.100, 224.7.7.0), RP 5.5.5.5
(5.12.5.100, 224.7.8.0), RP 5.5.5.5
(5.13.7.100, 224.8.5.0), RP 5.5.5.5
(5.13.7.100, 224.8.6.0), RP 5.5.5.5
(5.13.7.100, 224.8.7.0), RP 5.5.5.5
(5.13.7.100, 224.8.8.0), RP 5.5.5.5
(7.14.8.100, 224.9.5.0), RP 5.5.5.5
(7.14.8.100, 224.9.6.0), RP 5.5.5.5
(7.14.8.100, 224.9.7.0), RP 5.5.5.5

```

Variable Definitions

The following table defines parameters for the **show ip msdp show-all** command.

Variable	Value
<code>file WORD<1-99></code>	Specifies the file name to save the display output.
<code>vrf WORD<0-16></code>	Displays configuration information for a particular VRF.
<code>vrfids WORD<0-512></code>	Displays configuration information for a particular VRF ID.

Clear IPv4 MSDP Statistics

About This Task

Use the following procedure to clear all IPv4 Multicast Source Discovery Protocol (MSDP) statistics for all peers or a specific peer.

If you do not specify a VRF or VRF ID, the switch defaults to GRT.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Clear IPv4 MSDP statistics for all peers:
`clear ip msdp statistics [vrf WORD<0-16>] [vrfids WORD<0-512>]`
3. Clear IPv4 MSDP statistics for a specific peer:
`clear ip msdp statistics {A.B.C.D.} [vrf WORD<0-16>] [vrfids WORD<0-512>]`

Variable Definitions

Use the data in the following table to use the **clear ip msdp statistics** command.

Variable	Value
<code>vrf WORD<0-16></code>	Specifies a VRF instance by VRF name.
<code>vrfids WORD<0-512></code>	Specifies the ID of the VRF.
<code>{A.B.C.D.}</code>	Specifies the IPv4 MSDP address for a specific peer.

Basic MSDP Configuration Using EDM

This section provides procedures to configure basic MSDP using EDM. You can perform many tasks on the same EDM tab. For a quick look at field descriptions for that tab, see [Globals Field Descriptions](#) on page 3596. For additional tasks, see [MSDP Message Control using EDM](#) on page 3605

Configure the MSDP Originator ID

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted

(*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Configure the originator ID to set the RP address inside the SA message. The RP address must be a pre-configured CLIP interface on the global router or a VRF. The RP address is also the local IP address in all peer relations.



Note

Originator ID cannot be deleted if MSDP is enabled.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Globals** tab.
4. In the **RPAddress** box, type the IP address to use as the originator ID.
5. Select **Apply**.

Enable MSDP

Before You Begin

- You must configure the originator ID before you enable MSDP.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Enable or disable MSDP globally on the switch to permit further configuration to take place.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Globals** tab.
4. Select **Enabled** to enable MSDP.
5. Select **Apply**.

What to Do Next

[Create an MSDP Peer](#) on page 3596

Globals *Field Descriptions*

Use the data in the following table to use the **Globals** tab.

Name	Description
Enabled	Enables MSDP. The default is clear (disabled).
CacheLifetime	Configures the lifetime given to SA cache entries when created or refreshed.
NumSACacheEntries	Displays the total number of entries in the SA cache.
RPAddress	Specifies the IP address to use as the originator ID. If the address is not a system local address, the system rejects the configuration.
RouteMapName	Specifies the name of the optional route policy to create or modify. You do not need to create a route policy to use the redistribution filter. Note: To delete the route map name, clear the field and select Apply .
RedistributeFilterEnabled	Filters the (S,G,RP) entries provided by PIM to MSDP. The default is clear (disabled).
RedistruteFilterApply	Applies the changes made to the redistribute filter.
StatsClear	Clears MSDP statistics.

MSDP Peer Configuration using EDM

This section provides procedures to configure MSDP peer using EDM. You can perform many tasks on the same EDM tab. For a quick look at field descriptions for that tab, see [Peers Field Descriptions](#) on page 3601. For additional tasks, see [MSDP Message Control using EDM](#) on page 3605.

Create an MSDP Peer

Create an MSDP peer to establish a peer relationship between the local MSDP enabled router and a peer in another domain.

**Important**

Do not enable more than 20 active peers.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about

how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. Select **Insert**.
5. In **RemoteAddress**, type the IP address of the peer.
6. Select **Insert**.

Secure Control Messages

Configure Message Digest (MD) 5 authentication to secure control messages on the TCP connection between MSDP peers.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **Md5AuthEnabled**, and then select true.
5. In the row for the peer, double-click **Md5AuthPassword**, and then type a password.
6. Select **Apply**.

Configure the MSDP Peer SA Limit

Configure the SA limit to limit the number of SA messages from an MSDP peer. The router saves the SA messages in the local cache.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.

4. In the row for the peer, double-click **SALimit**, and then type a value.
5. Select **Apply**.

Configure a Peer Description

Configure a peer description to add a descriptive text to an MSDP peer, for easy identification of a peer.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **Description**, and then type a description for the peer.
5. Select **Apply**.

Configure the MSDP Peer Time-to-Live Threshold

Configure the time-to-live (TTL) threshold to limit which multicast data packets the router encapsulated in SA Messages forwarded to an MSDP peer. The TTL limits the number of hops a packet can take before the router drops the packet. The router sends out SA Messages with encapsulated data only if TTL equals or exceeds the value you configure. If the TTL is lower than the value you configure, the router drops the data packet and forwards the SA Message without the encapsulated data.



Note

MSDP transmits encapsulated multicast data packets inside forwarded MSDP messages. If the received SA is an encapsulated SA, then the switch parses the TTL value of the encapsulated data and compares it against the configured value. If the configured value is less than or equal to the parsed value, then the switch forwards the encapsulated data along with the SA, otherwise the switch forwards the SA alone by stripping the encapsulate data. By default, MSDP forwards encapsulated data along with the SA message. MSDP does not forward the encapsulated data to the local receivers.

When MSDP generates SA messages for SPB sources, the local cache miss data cannot be encapsulated into the SA messages that are sent to the peers. The switch supports forwarding SA messages with encapsulated data from sources to MSDP peers but not from MSDP peers to the receivers.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **Md5AuthEnabled**, and then select true.
5. In the row for the peer, double-click **DataTtl**, and then type a value.
6. Select **Apply**.

Configure the MSDP Peer Keepalive Messages

About This Task

Configure keepalive messages to adjust the interval in seconds at which an MSDP peer sends keep alive messages (default is 60 seconds) and the interval at which the MSDP peer waits for keep alive messages from other peers before it declares them down (default is 75 seconds).



Note

In a peer relationship, the keep alive interval configured on one peer must be at least 1 second less than the hold time configured on the other side of the peer relationship. This is not applicable when the hold time is set to 0 seconds.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **KeepAliveConfigured**, and then type the interval at which to send keepalive messages.
5. In the row for the peer, double-click **HoldTimeConfigured**, and then type the interval at which to wait for keepalive messages.
6. Select **Apply**.

Configure the MSDP Peer connect-retry Period

Configure the connect-retry period to specify the amount of time, in seconds, between connection attempts for peering sessions. The default is 30 seconds.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted

(*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **ConnectRetryInterval**, and then type the interval.
5. Select **Apply**.

Change the MSDP Peer Status

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Change the peer status to administratively enable or disable a configured peer. Disable the peer to stop the peering relationship.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **AdminEnabled**, and then select **true**.
5. Select **Apply**.

Clear the MSDP Peer Connection

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Clear the TCP connection to the specified MSDP peer, and reset all MSDP message counters.

The default is disabled (false).

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **ClearPeer**, and then select **true**.
5. Select **Apply**.

Delete an MSDP Peer

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. Select a row from the peer to delete.
5. Select **Delete**.
6. Select **Yes**.

View Peer Information

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.

Peers Field Descriptions

Use the data in the following table to use the **Peers** tab.

Name	Description
RemoteAddress	Shows the IP address of the remote MSDP peer.
State	Shows the state of the default peer. An MSDP node only accepts SA messages from an operational default peer. Only one default peer can be operational; the configured default peers provide redundancy.

Name	Description
AdminEnabled	Changes the peer status to administratively enable or disable a configured peer. The default value is disabled (false).
ClearPeer	Clears the TCP connection to the specified MSDP peer and resets all MSDP message counters. The default value is disabled (false).
ConnectRetryInterval	Specifies the connection retry period, in seconds, for this peer. The default value is 30 seconds.
HoldTimeConfigured	The default value is 75 seconds.
KeepAliveConfigured	Specifies the keepalive period, in seconds, configured for this MSDP speaker with this peer. If the value is 0 seconds, no periodic keepalive messages are sent to the peer after the MSDP connection is established. The default value is 60 seconds.
DataTtl	Specifies the time-to-live value, from 1–255. The default value is 1, which means that the router forwards all SA messages with encapsulated data.
InSAFilterEnabled	Activates the inbound SA filter for the peer.
InSAFilterRouteMapName	Specifies the name of the route map. If you configure the route map name, the filter accepts only the SA messages that meet the match criteria in the route map map-name with a permit keyword. If you do not configure the route map name, the system blocks all inbound SA messages from this peer.
OutSAFilterEnabled	Activates the outbound SA filter for the peer.
OutSAFilterRouteMapName	Specifies the name of the route map. If you configure the route map name, the filter sends only the SA messages that meet the match criteria in the route map map-name with a permit keyword. If you do not configure the route map name, the system blocks all outbound SA messages from this peer.
Description	Specifies the text description, up to 255 characters, for the peer.
SALimit	Specifies the maximum number of SA messages from an MSDP peer to keep in the SA cache. The valid values are from 0–6144; the default value is 6144.

Name	Description
Md5AuthEnabled	Activates MD5 authentication on the TCP connection between peers. The default is false.
Md5AuthPassword	Specifies a case-sensitive password, up to 80 characters, for MD5 authentication.
RemotePort	Shows the remote port for the TCP connection between the MSDP peers.
LocalPort	Shows the local port for the TCP connection between the MSDP peers.
OperEnabled	Shows the operational status of the peer.
RPFFailures	Shows the number of SA messages received from this peer that failed the Peer-RPF check. Discontinuities in the value of this counter can occur at reinitialization of the management system.
InSAs	Shows the number of MSDP SA messages received on this connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
OutSAs	Shows the number of MSDP SA messages sent on this connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
InSARRequests	Shows the number of MSDP SA-Request messages received on this connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
OutSARRequests	Shows the number of MSDP SA-Request messages sent on this connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
InSAResponses	Shows the number of MSDP SA-Response messages received on this connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
OutSAResponses	Shows the number of MSDP SA-Response messages sent on this connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.

Name	Description
InControlMessages	Shows the total number of MSDP messages received on this TCP connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
OutControlMessages	Shows the total number of MSDP messages transmitted on this TCP connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
InDataPackets	Shows the total number of encapsulated packets received on this TCP connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
OutDataPackets	Shows the total number of encapsulated packets transmitted on this TCP connection. Discontinuities in the value of this counter can occur at reinitialization of the management system.
FsmEstablishedTransitions	Shows the total number of times the BGP transitioned to the established state.
FsmEstablishedTime	Shows the time when the peer transitioned to the established state.
InMessageTime	Shows the time when the last MSDP message was received from the peer.
ConnectionAttempts	Shows the number of times the state machine has transitioned from inactive to connecting.
DiscontinuityTime	Shows the sysUpTime value (the time, in hundredths of a second, since the network management portion of the system last reinitialized) when one or more of the counters for this entry suffered a discontinuity. Discontinuities can occur at peer connection establishment. If no discontinuities occurred since the last reinitialization of the local management subsystem, the value is zero.
AsNumber	Specifies the autonomous system number of the MSDP peer. A peer can appear to be in another autonomous system (other than the one in which it really resides) if you use an MSDP peering session but do not use a Border Gateway Protocol peer session with that peer. If another autonomous system injects the prefix of the peer, the system displays the prefix as the autonomous system number of the peer.

Name	Description
TooShortMessages	Shows the number of short messages received from this peer.
InBadMessages	Shows the number of bad MSDP messages received from this peer.
InKeepAliveMessages	Shows the number of keepalive messages received from this peer.
OutKeepAliveMessages	Shows the number of keepalive messages transmitted to this peer.
SAsLearnedFromThisPeer	Shows the total number of SAs learned from this peer.
SAsAdvertisedToThisPeer	Shows the total number of SAs advertised from this peer.
UpOrDownTime	Shows the duration a peer has been up or down.
ConnAndStatsClearedTime	Shows the duration of connection and statistics cleared.

MSDP Message Control using EDM

This section provides procedures to configure MSDP message control using EDM.

Filter PIM Routes

Configure MSDP global filter for which the SA local cache are distributed to all MSDP peers. All SA local cache entries generate SA messages.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Globals** tab.
4. Select **RedistributeFilterEnabled**.
5. Type the name of the route policy in **RouteMapName**.
6. Select **RedistributeFilterApply** to apply the changes to the redistribute filter.

You do not need to apply the changes in the following situations:

- You create the redistribute filter without a route policy.
- You disable the redistribute filter.
- You remove a route policy from the redistribute filter.

7. Select **Apply**.

Filter SA Messages

Filter SA messages to determine which SA messages to accept from a peer and which SA messages to send to a peer. By default, no inbound or outbound filter exists.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Peers** tab.
4. In the row for the peer, double-click **InSAFilterEnabled**, and then select true.
5. In the row for the peer, double-click **InSAFilterRouteMapName**, and then type the route map name for the IN SA Filter of the peer.
6. In the row for the peer, double-click **OutSAFilterEnabled**, and then select true.
7. In the row for the peer, double-click **OutSAFilterRouteMapName**, and then type the route map name for the OUT SA Filter of the peer.
8. Select **Apply**.

Configure the MSDP Mesh Groups

Configure mesh groups to reduce SA flooding. A mesh group does not forward SA messages to other group members in the same mesh group. The originator, which is also a mesh group member, forwards SA messages to all group members. Create MSDP mesh groups from a group of meshed MSDP speakers from a domain. Do not create MSDP peerings between Controllers within the same SPB domain.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **Mesh Group** tab.
4. Select **Insert**.
5. In **Name**, type a name for the mesh group.

6. In **PeerAddress**, type the IP address of the peer to add the mesh group.
7. Select **Insert**.

Mesh Group Field Descriptions

Use the data in the following table to use the **Mesh Group** tab.

Name	Description
Name	Specifies the mesh group ID; the name of the mesh group from 1-64 characters.
PeerAddress	Specifies the IP address of the MSDP router that is the peer.

Clear the MSDP SA Cache

Clear the SA cache to clear the SA entries the router learns from all the peers or a specific peer.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **SA-Cache** tab.
4. Select **Clear SA-Cache**.

SA-Cache Field Descriptions

Use the data in the following table to use the **SA-Cache** tab.

Name	Description
GroupAddr	Shows the group IP address of the SA cache entry.
SourceAddr	Shows the source IP address of the SA cache entry.
OriginRP	Shows the RP address of the SA cache entry.
PeerLearnedFrom	Shows the peer from which this SA cache entry was accepted.
RPFPeer	Shows the peer from which an SA message corresponding to the cache entry is accepted.

Name	Description
InSAs	Discontinuities in the value of this counter can occur at reinitialization of the management system.
InDataPackets	Shows the number of MSDP encapsulated data packets received that are relevant to this cache entry.
UpTime	Shows the time since this entry was first placed in the SA cache.
ExpiryTime	Shows the time remaining before this entry expires from the SA cache.

MSDP Verification using EDM

This section provides procedures to verify MSDP using EDM. You can perform many tasks on the same EDM tab. For a quick look at field descriptions for that tab, see [SA-Cache-Records Field Descriptions](#) on page 3609.

View the Local SA Cache

View the local SA cache to display the (S, G) state the router learns from local Protocol Independent Multicast - Sparse Mode (PIM-SM) entries.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **SA-Cache-Records** tab.

View the Foreign SA Cache

View the foreign SA cache to display the (S, G) state the router learns from SA messages.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **MSDP**.
3. Select the **SA-Cache-Records** tab.

SA-Cache-Records Field Descriptions

Use the data in the following table to use the **SA-Cache-Records** tab.

Name	Description
TypeInformation	Shows the SA cache type. The SA cache type can be local or foreign cache.
GroupAddr	Shows the group IP address of the SA cache entry.
SourceAddr	Shows the source IP address of the SA cache entry.
OriginRP	Shows the RP address of the SA cache entry.
OriginatorAsNumber	Shows the AS number of the originator.
RouteType	Shows the type of route used for Reverse Path Forwarding checking. The value can be rip (1), ospf (2), static (3), bgp (4), isis(5) or none (6).

Controller Configuration

This section provides procedures to configure the Controller using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

The Controller functionality is configured at the global (switch-wide) level.

Controller Configuration using CLI

This section provides procedures to configure controller using the command line interface (CLI).

Enable the Controller Globally

Procedure

1. Enter IS-IS Router Configuration mode:
`enable`
`configure terminal`
`router isis`
2. Enable the Controller globally:
`spbm <1-100> multicast spb-pim-gw controller enable`

Variable Definitions

The following table defines parameters for the **spb** command.

Variable	Value
<1-100>	Specifies the IS-IS SPBM instance-id to create the SPBM instance.
<i>controller enable</i>	Enables the SPB-PIM Gateway Controller.

Display the Controller Admin Status

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the Controller Status:

```
show isis spbm
```

Example



Note

The SPB-PIM-GW column displays either Controller, Gateway, or Controller/Gateway if the Controller and or Gateway functionality is configured.

```
Switch:1>show isis spbm

=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY  NICK      LSDB      IP      IPV6      MULTICAST  SPB-PIM-GW
INSTANCE          VLAN      NAME      TRAP
-----
1          10,20      10       0.00.77   disable   enable   disable   enable     controller
=====

                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB  SMLT-VIRTUAL-EMAC  SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
INSTANCE
-----
1          primary      00:00:00:00:00:00              N/A                  N/A
-----

Total Num of SPBM instances: 1
=====
```

Display the Active Controller and Gateway Nodes

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the SPB-PIM Gateway active Controller and Gateway Nodes:

```
show ip spb-pim-gw node [controller | gateway] [spb-node-as-mac]
```


Example

Display all node lists:

```
Switch:1>show ip spb-pim-gw node
```

```
=====
                        Spb-pim-gw Active Controller/Gateway
=====

HOST-NAME      Role
-----
BEB3-4037      Gateway
BEB5-4011      Controller

Total Number of Nodes = 2/2
=====
```

Display Controller node lists only:

```
Switch:1>show ip spb-pim-gw node controller
```

```
=====
                        Spb-pim-gw Active Controller/Gateway
=====

HOST-NAME      Role
-----
BEB5-4011      Controller

Total Number of Nodes = 1/2
=====
```

Display Gateway node lists only:

```
Switch:1>show ip spb-pim-gw node gateway
```

```
=====
                        Spb-pim-gw Active Controller/Gateway
=====

HOST-NAME      Role
-----
BEB3-4037      Gateway

Total Number of Nodes = 1/2
=====
```

Display all node lists with MAC address:

```
Switch:1>show ip spb-pim-gw node spb-node-as-mac
```

```
=====
                        Spb-pim-gw Active Controller/Gateway
=====

HOST-NAME      Role      Mac Address
-----
BEB3-4037      Gateway    00:37:00:37:00:37
BEB5-4011      Controller 00:11:00:11:00:11
Total Number of Nodes = 2/2
=====
```

Display Controller node lists with MAC address:

```
Switch:1>show ip spb-pim-gw node controller spb-node-as-mac
```

```
=====
                               Spb-pim-gw Active Controller/Gateway
=====
```

HOST-NAME	Role	Mac Address
BEB5-4011	Controller	00:11:00:11:00:11

```
Total Number of Nodes = 1/2
=====
```

Display Gateway node lists with MAC address:

```
Switch:1>show ip spb-pim-gw node gateway spb-node-as-mac
```

```
=====
                               Spb-pim-gw Active Controller/Gateway
=====
```

HOST-NAME	Role	Mac Address
BEB3-4037	Gateway	00:37:00:37:00:37

```
Total Number of Nodes = 1/2
=====
```

Configure a Static Foreign Source on the Global Router

About This Task

Configure a static foreign source on the global router. Configuration is done at the Controller. Statically configure foreign sources, such as streams in a Source Specific Multicast (SSM) group range that are not advertised by the foreign network through MSDP. Non-SSM range group multicast address streams are advertised by MSDP and do not need to be statically configured.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Configure a static foreign source:


```
ip spb-pim-gw foreign-source {A.B.C.D} group {A.B.C.D}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip spb-pim-gw foreign-source 10.0.0.1 group 240.0.0.1
```

Variable Definitions

The following table defines parameters for the **ip spb-pim-gw** command.

Variable	Value
<i>foreign-source</i> {A.B.C.D}	Specifies the multicast foreign source IP address.
<i>group</i> {A.B.C.D}	Specifies the group IP address.

Configure a Static Foreign Source on a VRF

About This Task

Configure a static foreign source on a VRF; configuration is done at the Controller.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Configure a static foreign source:

```
ip spb-pim-gw foreign-source {A.B.C.D} group {A.B.C.D}
```

Example

In the following example, vrf-10 is configured with vrf id 10.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrf-10
Switch:1(router-vrf)#ip spb-pim-gw foreign-source 10.0.0.1 group 240.0.0.1
```

Variable Definitions

The following table defines parameters for the **ip spb-pim-gw** command.

Variable	Value
<i>foreign-source</i> {A.B.C.D}	Specifies the multicast foreign source IP address.
<i>group</i> {A.B.C.D}	Specifies the group IP address.

Display Foreign Sources

About This Task

Use the following procedure to display the foreign sources learned from MSDP or statically configured at the Controller.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the foreign source information:

```
show ip spb-pim-gw foreign-source [all] [controller | gateway] [vrf
WORD<0-16>] [vrfs WORD<0-512>] [source {A.B.C.D}] [group {A.B.C.D}]
[static | msdp] [spb-node-as-mac]
```

Example



Note

The command **show ip spb-pim-gw**, which specifies the parameter *controller*, the OWNER column displays the RP address of the MSDP peer from which the foreign source was learned. If the *gateway* parameter is specified, then the OWNER column displays MSDP rather than the actual RP address.

```
Switch:1>show ip spb-pim-gw foreign-source controller
=====
SPB-PIM-GW Controller Foreign Source
=====
SOURCE    GROUP    SPB-PIM-GW    VRF    OWNER
-----
10.0.0.1  240.0.0.1  beeb-1        GlobalRouter  47.17.0.1
10.0.0.2  240.0.0.2  beeb-1        GlobalRouter  static
10.0.0.3  240.0.0.3  -             GlobalRouter  47.17.0.2
=====
```

Display the foreign sources from a specific VRF:

```
Switch:1>show ip spb-pim-gw foreign-source controller vrf green
=====
SPB-PIM-GW Controller Foreign Source
=====
SOURCE    GROUP    SPB-PIM-GW    VRF    OWNER
-----
10.0.0.1  240.0.0.1  beeb-1        green    47.17.0.1
10.0.0.2  240.0.0.2  beeb-1        green    static
10.0.0.3  240.0.0.3  -             green    47.17.0.2
=====
```

Display all the foreign sources at the Controller with the Gateway in the SPB-PIM-GW shown as a mac address rather than a nickname:

```
Switch:1>show ip spb-pim-gw foreign-source controller vrf green spb-node-as-mac
=====
SPB-PIM-GW Controller Foreign Source
=====
SOURCE    GROUP    SPB-PIM-GW    VRF    OWNER
-----
10.0.0.1  240.0.0.1  00:0b:eb:00:00:a1  green    47.17.0.1
10.0.0.2  240.0.0.2  00:0b:eb:00:00:a1  green    static
10.0.0.3  240.0.0.3  -             green    47.17.0.2
=====
```

Variable Definitions

The following table defines parameters for the **show ip spb-pim-gw foreign source** command.

Variable	Value
<i>all</i>	Displays information for all the VRF IDs from the Controller and Gateway foreign source database.
<i>controller</i>	Displays information from the Controller foreign source database. Only displays information on nodes configured as Controller.
<i>gateway</i>	Displays information from the Gateway foreign source database. Only displays information on nodes configured as Gateway.
<i>vrf WORD<0-16></i>	Displays information from the Controller foreign source database for a specific VRF name.
<i>vrfids WORD<0-512></i>	Displays information from the Controller foreign source database for a range of VRF IDs. Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.
<i>source {A.B.C.D}</i>	Displays information for the specific source IP address from the Controller foreign source database.
<i>group {A.B.C.D}</i>	Displays information for the specific multicast group IP address from the Controller foreign source database.
<i>static</i>	Displays information from the Controller foreign source database that is configured statically.
<i>msdp</i>	Displays information from the Controller foreign source database that is learned through MSDP.
<i>spb-node-as-mac</i>	Displays the MAC address for the assigned SPB-PIM Gateway.

Display Multicast over Fabric Connect Sources

About This Task

Use the following procedure to display all the SPB Multicast over Fabric Connect sources distributed to MSDP. This procedure is only valid on a Controller node.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the SPB source information:

```
show ip spb-pim-gw spbmc-source [vrf WORD<1-32>] [vrfrids WORD<0-512>]
[source {A.B.C.D}] [group {A.B.C.D}] [originator WORD<1-32>] [spb-
node-as-mac]
```

Example

```
Switch:1>show ip spb-pim-gw spbmc-source
```

```
=====
```

```
SPB-PIM-GW SPB Source
```

```
=====
```

SOURCE	GROUP	VRF	ORIGINATOR
10.0.0.1	240.0.0.1	GlobalRouter	bcb-1
10.0.0.2	240.0.0.2	GlobalRouter	bcb-2
10.0.0.3	240.0.0.3	GlobalRouter	bcb-2

```
=====
```

Display the SPB Multicast over Fabric Connect from a specific VRF:

```
Switch:1>show ip spb-pim-gw spbmc-source vrf green
```

```
=====
```

```
SPB-PIM-GW Foreign Source
```

```
=====
```

SOURCE	GROUP	SPB-PIM-GW	VRF	OWNER	CONTROLLER
10.0.0.1	240.0.0.1	beb-1	green	47.17.0.1	bcb-2
10.0.0.2	240.0.0.2	beb-1	green	static	bcb-2
10.0.0.3	240.0.0.3	-	green	47.17.0.2	bcb-2

```
=====
```

Display all the SPB Multicast over Fabric Connect sources advertised to MSDP with the originator value shown as a MAC address rather than a host name:

```
Switch:1>show ip spb-pim-gw spbmc-source vrf green spb-node-as-mac
```

```
=====
```

```
SPB-PIM-GW SPB Source
```

```
=====
```

SOURCE	GROUP	VRF	ORIGINATOR
10.0.0.1	240.0.0.1	green	00:0b:cb:00:00:c2
10.0.0.2	240.0.0.2	green	00:0b:cb:00:00:c2
10.0.0.3	240.0.0.3	green	00:0b:cb:00:00:c2

```
=====
```

Variable Definitions

The following table defines parameters for the **show ip spb-pim-gw spbmc-source** command.

Variable	Value
<code>vrf WORD<1-32></code>	Displays SPB originated sources for a specific VRF.
<code>vrfids WORD<0-512></code>	Displays SPB originated sources for a range of VRF IDs. Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.
<code>source {A.B.C.D}</code>	Displays information for a specific source IP address from SPB originated sources database.
<code>group {A.B.C.D}</code>	Displays information for a specific multicast group IP address from SPB originated sources database.
<code>originator WORD<0-32></code>	Displays information for a specific originator host name from SPB originated sources database.
<code>spb-node-as-mac</code>	Displays the originator of SPB originated sources as a MAC address rather than a nickname.

Controller Configuration using EDM

This section provides procedures to configure controller using EDM.

Enable the Controller

About This Task

Enable the Controller globally.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **SPBM** tab.
4. In the row for the ID, double-click **McastSpbPimGwControllerEnable**, and then select **true**.
5. Select **Apply**.

Display the Controller and Gateway Admin Status

About This Task

Use the following procedure to display the admin status of the Controller and Gateway.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **SPBM** tab.

Display Active Controller and Gateway Nodes

About This Task

Use the following procedure to display the active Controllers and Gateways in the SPBM domain.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Node** tab.

Node Field Descriptions

Use the data in the following table to use the **Node** tab.

Name	Description
MacAddress	Shows the MAC address of the active node.
HostName	Shows the host name of the active node.
RoleType	Shows the role of the active node: either gateway, controller, or both.

Configure a Static Foreign Source

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Configure a static foreign source on the global router or VRF. Configuration is done at the Controller.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Controller-Foreign-Source** tab.
4. Select **Insert**.

5. In **SourceAddress**, type the multicast foreign source IP address.
6. In **GroupAddress**, type the group IP address.
7. Select **Insert**.

Controller-Foreign-Source Field Descriptions

Use the data in the following table to use the **Controller-Foreign-Source** tab.

Name	Description
SourceAddress	Specifies the source IP address from a foreign multicast domain.
GroupAddress	Specifies the multicast group IP address associated with the foreign source.
GatewaySysId	Displays the system ID of the node selected as the Gateway for this foreign source. GatewaySysId field will have a valid value if the Gateway is assigned to a source. If the Gateway is not assigned to a source the value is 0.
GatewayHostName	Displays the host name of the node selected as the Gateway for this foreign source. GatewayHostName field will have valid values if the Gateway is assigned to a source. If the Gateway is not assigned to a source the value is NULL.
Type	Displays the owner type for this source.
Owner	Displays the IP address of the MSDP peer if the foreign source is MSDP.

Display Foreign Sources for the Controller

About This Task

Use the following procedure to display the foreign sources learned from MSDP or statically configured at the Controller.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Controller-Foreign-Source** tab.

Controller-Foreign-Source Field Descriptions

Use the data in the following table to use the **Controller-Foreign-Source** tab.

Name	Description
SourceAddress	Specifies the source IP address from a foreign multicast domain.
GroupAddress	Specifies the multicast group IP address associated with the foreign source.

Name	Description
GatewaySysId	Displays the system ID of the node selected as the Gateway for this foreign source. GatewaySysId field will have a valid value if the Gateway is assigned to a source. If the Gateway is not assigned to a source the value is 0.
GatewayHostName	Displays the host name of the node selected as the Gateway for this foreign source. GatewayHostName field will have valid values if the Gateway is assigned to a source. If the Gateway is not assigned to a source the value is NULL.
Type	Displays the owner type for this source.
Owner	Displays the IP address of the MSDP peer if the foreign source is MSDP.

Display Multicast over Fabric Connect Sources

About This Task

Use the following procedure to display all the SPB Multicast over Fabric Connect sources distributed to MSDP.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **spbm-source** tab.

Spbmc-Source Field Descriptions

Use the data in the following table to use the **Spbmc-Source** tab.

Name	Description
SourceAddress	Displays the source IP address from SPBM multicast domain.
GroupAddress	Displays the multicast group IP address associated with the SPBM source.
OriginatorSysId	Displays the system ID of the node from which the source originates.
OriginatorHostName	Displays the host name of the node from which the source originates.

Gateway Configuration

This section provides procedures to configure the Gateway using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

The Gateway functionality is configured at the global (switch-wide) level. SPB-PIM Gateway Interfaces are configured at the interface level. For more information on SPB-PIM Gateway Interfaces configuration, see [SPB-PIM Gateway Interface Configuration](#) on page 3626.

Gateway Configuration using CLI

This section provides procedures to configure gateway using the command line interface (CLI).

Enable the Gateway

About This Task

Enable the Gateway at the global (switch-wide) level.

Procedure

- 1. Enter IS-IS Router Configuration mode:
enable

configure terminal

router isis
- 2. Enable the Gateway globally:
spbm <1-100> multicast spb-pim-gw gateway enable

Variable Definitions

The following table defines parameters for the **spb** command.

Variable	Value
<1-100>	Specifies the IS-IS SPBM instance-id to create the SPBM instance.
gateway enable	Enables the SPB-PIM Gateway.

Display the Gateway Admin Status

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Display the Gateway admin status:
show isis spbm

Example

```
Switch:1>show isis spbm

=====
                        ISIS SPBM Info
=====
SPBM      B-VID    PRIMARY  NICK    LSDB    IP    IPV6    MULTICAST  SPB-PIM-GW
INSTANCE  VLAN          NAME     TRAP
=====
```

```
1          10,20      10          0.00.77  disable  enable   disable  enable   Gateway

=====
                        ISIS SPBM SMLT Info
=====

SPBM      SMLT-SPLIT-BEB  SMLT-VIRTUAL-BMAC    SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
INSTANCE
-----
1          primary      00:00:00:00:00:00          N/A                  N/A

-----
Total Num of SPBM instances: 1
-----
```

If the controller and gateway are both enabled on the node

```
Switch:1>show isis spbm

=====
                        ISIS SPBM Info
=====

SPBM      B-VID      PRIMARY  NICK      LSDB      IP        IPV6      MULTICAST  SPB-PIM-GW
INSTANCE  VLAN      NAME     TRAP
-----
1          10,20      10       0.00.77  disable  enable   disable  enable
controller
          /gateway

=====
                        ISIS SPBM SMLT Info
=====

SPBM      SMLT-SPLIT-BEB  SMLT-VIRTUAL-BMAC    SMLT-PEER-SYSTEM-ID  SMLT-PEER-NICK-NAME  SMLT-PEER-HOST-NAME
INSTANCE
-----
1          primary      00:00:00:00:00:00          N/A                  N/A

-----
Total Num of SPBM instances: 1
-----
```

Display Foreign Sources Information

About This Task

Use the following procedure to display the Gateway foreign sources database. If executed on a Gateway node, it displays the foreign sources assigned to the Gateway by the Controller. Foreign sources are originally learned from MSDP or statically configured on the Controller.

Procedure

- 1. To enter User EXEC mode, log on to the switch.

2. Display the foreign sources information:

```
show ip spb-pim-gw foreign-source [all] [controller | gateway] [vrf
WORD<0-16>] [vrfsids WORD<0-512>] [source {A.B.C.D}] [group {A.B.C.D}]
[from-controller <0x00:0x00:0x00:0x00:0x00:0x00> | preferred][static |
msdp] [spb-node-as-mac]
```

Example

```
Switch:1>show ip spb-pim-gw foreign-source gateway
=====
SPB-PIM-GW Gateway Foreign Source
=====
```

SOURCE	GROUP	SPB-PIM-GW	VRF	PORT	VLAN	CONTROLLER	OWNER
10.0.0.1	240.0.0.1	beb-2	GlobalRouter	1/1	200	bcb-2	msdp
10.0.0.2	240.0.0.2	beb-2	GlobalRouter	1/1	200	bcb-2	static
10.0.0.3	240.0.0.3	beb-2	GlobalRouter	-	-	bcb-2	msdp

```
=====
```



Note

The SPB-PIM-GW column displays the node that is selected as the Gateway for the particular source or group stream. The OWNER column displays either msdp or static depending on how the source was originally learned at the assigning Controller. The PORT and VLAN columns represent the port or VLAN toward the source.

Display the foreign sources from a specific VRF:

```
Switch:1>show ip spb-pim-gw foreign-source gateway vrf green
=====
SPB-PIM-GW Gateway Foreign Source
=====
```

SOURCE	GROUP	SPB-PIM-GW	VRF	PORT	VLAN	CONTROLLER	OWNER
10.0.0.1	240.0.0.1	beb-2	green	1/1	200	bcb-2	msdp
10.0.0.2	240.0.0.2	beb-2	green	1/1	200	bcb-2	static
10.0.0.3	240.0.0.3	beb-2	green	-	-	bcb-2	msdp

```
=====
```

Display all the foreign sources available at the Gateway with SPB-PIM-GW and Controller as mac:

```
Switch:1>show ip spb-pim-gw foreign-source gateway vrf green spb-node-as-mac
=====
SPB-PIM-GW Gateway Foreign Source
=====
```

SOURCE	GROUP	SPB-PIM-GW	VRF	PORT	VLAN	CONTROLLER	OWNER
10.0.0.1	240.0.0.1	00:0b:eb:00:00:a2	green	1/1	200	00:0b:cb:00:00:c2	msdp
10.0.0.2	240.0.0.2	00:0b:eb:00:00:a2	green	1/1	200	00:0b:cb:00:00:c2	static
10.0.0.3	240.0.0.3	00:0b:eb:00:00:a2	green	-	-	00:0b:cb:00:00:c2	msdp

```
=====
```

Display all the foreign sources available at the Gateway which are statically configured at the Controller:

```
Switch:1>show ip spb-pim-gw foreign-source gateway vrf green static
```

```
=====
SPB-PIM-GW Gateway Foreign Source
=====
```

SOURCE	GROUP	SPB-PIM-GW	VRF	PORT	VLAN	CONTROLLER	OWNER
10.0.0.2	240.0.0.2	beb-2	green	1/1	200	bcb-2	static

```
=====
```

Variable Definitions

The following table defines parameters for the **show ip spb-pim-gw foreign source** command.

Variable	Value
<i>all</i>	Displays information for all the VRF IDs from the Controller and Gateway foreign source database.
<i>controller</i>	Displays information from the Controller foreign source database. Only displays information on nodes configured as Controller.
<i>gateway</i>	Displays information from the Gateway foreign source database. Only displays information on nodes configured as Gateway.
<i>vrf WORD<0-16></i>	Displays information from the Gateway foreign source database for a specific VRF name.
<i>vrfids WORD<0-512></i>	Displays information from the Gateway foreign source database for a range of VRF IDs. Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.
<i>source {A.B.C.D}</i>	Displays information for the specific source IP address from the Gateway foreign source database.
<i>group {A.B.C.D}</i>	Displays information for the specific multicast group IP address from the Gateway foreign source database.
<i>from-controller</i> <i>0x00:0x00:0x00:0x00:0x00:0x00</i>	Displays information filtering on a specific Controllers assignments, where the Controller is specified as a mac address.

Variable	Value
<i>from-controller preferred</i>	Displays information from Gateway source database filtering on a preferred Controller or chosen by the Gateway.
<i>static</i>	Displays information from the Gateway foreign source database that is configured statically at the assigning Controller.
<i>msdp</i>	Displays information from the Gateway foreign source database that is learned through MSDP.
<i>spb-node-as-mac</i>	Displays the MAC address for the assigned PIM-GW.

Gateway Configuration using EDM

This section provides procedures to configure gateway using EDM.

Enable the Gateway Globally

About This Task

Use this procedure to enable the Gateway at the global (switch-wide) level.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **SPBM**.
3. Select the **SPBM** tab.
4. In the row for the ID, double-click **McastSpbPimGwGatewayEnable**, and then select **true**.
5. Select **Apply**

Display Foreign Sources for the Gateway

About This Task

Use the following procedure to display the Gateway foreign source database. Foreign sources are originally learned from MSDP or statically configured on the Controller.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Gateway-Foreign-Source** tab.

Gateway-Foreign-Source Field Descriptions

Use the data in the following table to use the **Gateway-Foreign-Source** tab.

Name	Description
SourceAddress	Displays the foreign source IP address.
GroupAddress	Displays the multicast group IP address associated with the foreign source.
ControllerSysId	Displays the system ID of the controller node that sends this foreign source.
ControllerHostName	Displays the host name of the controller node that sends this foreign source.
GatewaySysId	Displays the system ID of the node selected as the gateway for this foreign source.
GatewayHostName	Displays the host name of the node selected as the gateway for this foreign source.
InVid	Displays the VLAN ID of the SPB-PIM Gateway interface through which the source of this source is reachable.
InPort	Displays the physical interface through which the source of this source is reachable.
OwnerType	Displays if the owner is MSDP or static.

SPB-PIM Gateway Interface Configuration

This section provides procedures to configure the SPB-PIM Gateway interface using the Command Line Interface (CLI) and Enterprise Device Manager (EDM).

The SPB-PIM Gateway interface is either a VLAN or a Brouter port interface. An SPB-PIM Gateway interface is configured separately from the global (switch-wide) Gateway functionality. The global Gateway configuration does not affect the administrative or the operational state of the SPB-PIM Gateway interfaces which function independently. However, the Gateway node functionality works in conjunction with the Gateway Interface functionality. Configure SPB-PIM Gateway on an interface that connects to a router in a foreign PIM or SPB multicast domain.

SPB-PIM Gateway Interface Configuration using CLI

This section provides procedures to configure SPB-PIM gateway interface using the command line interface (CLI).

*Enable SPB-PIM Gateway on a VLAN***Procedure**

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable SPB-PIM Gateway on a VLAN:

```
ip spb-pim-gw enable
```

**Note**

SPB-PIM Gateway cannot be enabled under the following circumstances:

- If the IP interface does not exist on the VLAN. An IP Address must first be configured on the VLAN.
- If the **spbm_config_mode** boot flag is set to false.
- If the VLAN is configured with a circuitless IP.
- If the interface is a management VLAN.
- If **ip igmp snooping** is enabled.
- If the spb-multicast is enabled on the VLAN.
- If the VLAN has SMLT ports.
- If the VLAN has an i-sid configured.
- If the VLAN is a vIST VLAN.

*Enable SPB-PIM Gateway on a Brouter Port Interface***Procedure**

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable SPB-PIM Gateway on a brouter port:

```
ip spb-pim-gw enable
```

**Note**

SPB-PIM Gateway cannot be enabled under the following circumstances:

- If the IP interface is not configured using the **brouter** command
- If the IP interface does not exist on the brouter port
- If the **spbm_config_mode** boot flag is set to false
- If **ip igmp snooping** is enabled
- If the spb-multicast is enabled on the brouter port
- If the brouter port is part of an SMLT or vIST Vlan
- If the brouter port has an i-sid configured

Configure the SPB-PIM Gateway VLAN Optional Parameters

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Configure the SPB-PIM Gateway VLAN HELLO interval:

```
ip spb-pim-gw hello-interval <0-18724>
```

3. Configure the SPB-PIM Gateway VLAN JOIN PRUNE interval:

```
ip spb-pim-gw ip join-prune-interval <1-18724>
```

Variable Definitions

The following table defines parameters for the **ip spb-pim-gw** command.

Variable	Value
<i>hello-interval</i> <0-18724>	Specifies the HELLO interval in seconds. The default value is 30 seconds.
<i>join-prune-interval</i> <1-18724>	Specifies the JOIN PRUNE interval in seconds. The default value is 60 seconds.

*Configure the SPB-PIM Gateway Brouter Port Optional Parameters***Procedure**

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the SPB-PIM Gateway brouter port HELLO interval:

```
ip spb-pim-gw hello-interval <0-18724>
```

3. Configure the SPB-PIM Gateway brouter port JOIN PRUNE interval:

```
ip spb-pim-gw join-prune-interval <1-18724>
```

Variable Definitions

The following table defines parameters for the **ip spb-pim-gw** command.

Variable	Value
<i>hello-interval</i> <0-18724>	Specifies the HELLO interval in seconds. The default value is 30 seconds.
<i>join-prune-interval</i> <1-18724>	Specifies the JOIN PRUNE interval in seconds. The default value is 60 seconds.

*Display the SPB-PIM Gateway Interface Default Values***About This Task**

Use the following procedure to display the default values used for the SPB-PIM Gateway interface HELLO and JOIN PRUNE intervals unless specifically configured on the individual interface.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the SPB-PIM Gateway interface default values:

```
show ip spb-pim-gw
```

Example

```
Switch:1>show ip spb-pim-gw
```

```
=====
                        Spb-pim-gw General Group
=====
Hello Interval           : 30
Join-Prune Interval      : 60
```

Display the SPB-PIM Gateway Brouter Port Information

About This Task

Use the following procedure to display the SPB-PIM Gateway brouter port information. This procedure displays the administrative (configured) state of the interface as well as the operational state, and the HELLO and JOIN PRUNE intervals. An interface can be administratively ENABLED but operationally DISABLED if, for example, mvpn is not enabled on the VRF, or `spbm <spbm-instance> multicast enable` is not configured.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display the SPB-PIM Gateway interface information:

```
show ip spb-pim-gw interface [gigabitethernet{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]] [vrf WORD<1-16>] [vrfrids WORD<0-512>]
```

Example

Switch:1#show ip spb-pim-gw interface gigabitethernet 1/2

Port Ip Spb-pim-gw				
ORT-NUM	OPSTATE	ADMINSTATE	HELLOINT	JPINT
1/2	Disabled	Enabled	30	60

Variable Definitions

The following table defines parameters for the `show ip spb-pim-gw interface gigabit` command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>vrf WORD<1-16></code>	Displays SPB-PIM Gateway interface information for a specific VRF.
<code>vrfids WORD<0-512></code>	<p>Displays SPB-PIM Gateway interface information for a range of VRF IDs.</p> <p>Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.</p>

Display the SPB-PIM Gateway VLAN Information

Use the following procedure to display the SPB-PIM Gateway VLAN interface information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the SPB-PIM Gateway VLAN interface information:

```
show ip spb-pim-gw interface [vlan]
```

Example

```
Switch:1>show ip spb-pim-gw interface
=====
Spb-pim-gw Interface - GlobalRouter
=====
```

IF	ADDR	MASK	JPINT	HELLOINT	OPSTATE	ADMINSTATE
Vlan50	50.1.1.1	255.255.255.0	60	30	Disabled	Enabled
Vlan123	123.1.1.1	255.255.255.0	60	30	Disabled	Disabled
Vlan142	142.1.1.1	255.255.255.0	60	30	Enabled	Enabled
Vlan400	100.1.1.2	255.255.255.0	60	30	Disabled	Disabled

```
Total spb-pim-gw Interfaces Displayed 4/4
```

Variable Definitions

The following table defines parameters for the **show ip spb-pim-gw interface** command.

Variable	Value
vlan-id	The VLAN ID of an interface to display.

Display the SPB-PIM Gateway Neighbor Information

Use the following procedure to display the SPB-PIM Gateway interfaces neighbor information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the SPB-PIM Gateway neighbor information:

```
show ip spb-pim-gw neighbor [vrf WORD<1-16>] [vrfrids WORD<0-512>]
```

Example

```
Switch:1>show ip spb-pim-gw neighbor

=====
                        Spb-pim-gw Neighbor - GlobalRouter
=====
INTERFACE ADDRESS      UPTIME                EXPIRE
Vlan26      26.1.1.10          0 day(s), 00:11:36    0 day(s), 00:01:28

Total SPB-PIM-GW Neighbors Displayed = 1/1
=====
```

Variable Definitions

The following table defines parameters for the **show ip spb-pim-gw neighbor** command.

Variable	Value
vrf WORD<1-16>	Displays the SPB-PIM Gateway interface neighbor information for a specific VRF name.
vrfrids WORD<0-512>	Displays the SPB-PIM Gateway interface neighbor information for a range of VRF IDs. Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.

*Display the SPB-PIM Gateway Multicast Routes***About This Task**

Use the following procedure to display the SPB-PIM Gateway multicast routes. This procedure displays upstream (toward the foreign source) information and downstream (receiver) information on the SPB-PIM Gateway interfaces. This command does not display the following information:

- Upstream information for streams ingressing on spb-multicast interfaces
- Upstream information for streams ingressing from a remote SPB node
- Receivers in spb-multicast interfaces

Use the **show isis spbm ip-multicast-route** command to display information on all multicast streams and the multicast streams ingress interfaces and egress interfaces.

Use the **show ip spb-pim-gw mroute** command to display information only on SPB-PIM Gateway interfaces.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the SPB-PIM Gateway multicast routes:

```
show ip spb-pim-rw mroute [source {A.B.C.D}] [group {A.B.C.D}] [vrf
WORD<1-16>] [vrfids WORD<0-512>]
```

Example**Note**

The **show ip spb-pim-gw mroute** command displays active upstream information and active downstream information per (*,g) and (s,g) per port, which includes *G Join or Prune pending state, SG join or prune pending state, and SG Rpt PRUNE or PRUNE pending state.. There can be upstream information for a specific SG and no downstream, and vice-versa, as the information in the command **show ip spb-pim-gw mroute** reflects only information known on the SPB-PIM Gateway interfaces. For example, there might be downstream receivers on a SPB-PIM Gateway Interface for a particular stream which is ingressing on an spb-multicast interface; thus, the upstream information will not be displayed using this command.

```
Switch:1>show ip spb-pim-gw mroute
=====
                Spb-pim-gw Active PIM Multicast Route - GlobalRouter
=====

Src: 0.0.0.0          Grp: 225.1.1.1
Flags: WC
Joined Ports:
Vlan    Ports          Join Timer
----    -
Vlan30  1/3              155
-----
```

```

Src: 123.1.1.101   Grp: 225.1.1.1   Upstream: 50.1.1.1   Incoming Port: Vlan50-1/5

Flags: SG
SG Joined Ports:
Vlan    Ports          Join Timer
----    -
Vlan30  1/3                184

SG Prune Pending Ports:
Vlan    Ports          Prune Pending Timer
----    -
Vlan40  1/4                180

SG Rpt Pruned Ports:
Vlan    Ports          RPT Prune Timer
----    -
Vlan90  1/9                156

SG Rpt Prune Pending Ports:
Vlan    Ports          RPT Prune Pending Timer
----    -
Vlan60  1/6                164

-----

Total Num of Entries Displayed 2/2
Flags Legend:
WC=(*,Grp) entry, SG=(Src,Grp) entry

```

Variable Definitions

The following table defines parameters for the **show ip spb-pim-gw mroute** command.

Variable	Value
<i>vrf</i> <i>WORD</i> <1-16>	Displays the SPB-PIM Gateway mroute information for a specific VRF name.
<i>vrfids</i> <i>WORD</i> <0-512>	Displays the SPB-PIM Gateway interface mroute information for a range of VRF IDs. Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.
<i>group</i> { <i>A.B.C.D</i> }	Displays mroute information specific to a group IP address.
<i>source</i> { <i>A.B.C.D</i> }	Displays mroute information specific to a source IP address.

Display the IP mroute Routes

About This Task

Use the following procedure to display multicast routes ingressing on either SPB-PIM Gateway interfaces or SPB multicast interfaces.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the multicast routes:

```
show ip mroute mroute
```

Example



Note

The UPSTREAM_NBR field is populated only if the stream was learned across a spb-pim-gw interface, and the upstream neighbor is the PIM neighbor IP address toward the source. The PROT field equals spb-pim-gw when the stream's source is learned on a VLAN configured for protocol spb-pim-gw. If the stream's source is learned on a VLAN configured for spb-multicast, the PROT field equals spb.

```
Switch:1>show ip mroute route
```

```
=====
                        Mroute Route - GlobalRouter
=====
GROUP                SOURCE                SRCMASK                UPSTREAM_NBR    IF    EXPIR    PROT
-----
233.252.0.1          0.0.0.0                0.0.0.0                0.0.0.0        V3    30      spb-access
233.252.0.1          192.0.2.102            255.255.255.0          0.0.0.0        -     0       spb-network
233.252.0.2          0.0.0.0                0.0.0.0                0.0.0.0        V2    30      pimsm
225.1.1.1            198.51.100.99          255.255.255.0          0.0.0.0        V3    173     spb-pim-gw

Total 4
```

Variable Definitions

The following table defines parameters for the **show ip mroute route** command.

Variable	Value
<code>vrf WORD<1-32></code>	Displays the multicast mroute information for a specific VRF name.
<code>vrfids WORD<0-512></code>	Displays the multicast mroute information for a range of VRF IDs. Note: Enter a single VRF ID or multiple VRF IDs separated by ',' or enter a range of VRF IDs 'x-y'.

SPB-PIM Gateway Interface Configuration using EDM

This section provides procedures to configure SPB-PIM gateway interface using EDM.

*Enable SPB-PIM Gateway on a VLAN***Procedure**

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. In the row for the VLANs, double-click **SpbPimGatewayMulticast**, and then select **enable**.
5. Select **Apply**.

*Enable SPB-PIM Gateway on a Brouter Port Interface***About This Task**

Enable SPB-PIM Gateway on a Brouter port.

Procedure

1. In the Device Physical View tab, select the port you need to configure.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **SPB-PIM-GW** tab.
5. Select **Enable**.
6. Select **Apply**.

SPB-PIM-GW Field Descriptions

Use the data in the following table to use the **SPB-PIM-GW** tab.

Name	Description
Enable	Enables SPB-PIM Gateway on the interface. The default is disabled.
OperState	Displays the current operational state of this SPB-PIM Gateway interface.
Address	Displays the primary IP address of this router on this SPB-PIM Gateway interface.
AddressMask	Displays the primary IP address mask of this router on this SPB-PIM Gateway interface.
HelloInterval	Configures the PIM HELLO transmission interval.
JoinPruneInterval	Configures the PIM JOIN PRUNE transmission interval.

*Configure SPB-PIM Gateway VLAN Optional Parameters***About This Task**

Configure the SPB-PIM Gateway interface parameters on a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a row for the VLAN, and then select **IP**.
5. Select the **SPB-PIM-GW** tab.
6. In **HelloInterval**, type the hello transmission interval.
7. In **JoinPruneInterval**, type the join prune transmission interval.
8. Select **Apply**.

SPB-PIM-GW Field Descriptions

Use the data in the following table to use the **SPB-PIM-GW** tab.

Name	Description
OperState	Displays the current operational state of this SPB-PIM Gateway interface.
Address	Displays the primary IP address of this router on this SPB-PIM Gateway interface.
AddressMask	Displays the primary IP address mask of this router on this SPB-PIM Gateway interface.
HelloInterval	Configures the PIM HELLO transmission interval. This SPB-PIM Gateway VLAN level interval setting overrides the inherited global SPB-PIM Gateway HELLO interval setting. Setting the HELLO Interval to 0 causes the neighbors to never expire its neighborship with this local SPB-PIM Gateway interface.
JoinPruneInterval	Configures the PIM JOIN PRUNE transmission interval. This SPB-PIM Gateway VLAN level interval setting overrides the inherited global level JOIN PRUNE transmission interval setting.

Configure the SPB-PIM Gateway Optional Parameters

About This Task

Perform this procedure to configure the optional parameters on existing SPB-PIM Gateway interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Interfaces** tab.
4. In the row for the interfaces, double-click **HelloInterval**, and then type the hello interval in seconds.

5. In the row for the interfaces, double-click **JoinPruneInterval**, and then type the join prune interval in seconds.
6. Select **Apply**.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	Shows the interface index.
OperState	Shows the operational state of the interface.
AddressType	Shows the address type of the interface.
Address	Shows the address assigned to the interface.
AddressMask	Shows the address mask associated with the interface address.
HelloInterval	Configures the PIM HELLO transmission interval. The default is 30 seconds.
JoinPruneInterval	Configures the PIM JOIN PRUNE transmission interval. The default is 60 seconds.

Display the SPB-PIM Gateway Interface Default Values

About This Task

Use the following procedure to display the default values used for the SPB-PIM Gateway interface HELLO and JOIN PRUNE intervals unless specifically configured on the individual interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Globals** tab.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
HelloInterval	Displays the PIM HELLO transmission interval.
JoinPruneInterval	Displays the PIM JOIN PRUNE transmission interval.

*Display the SPB-PIM Gateway Router Port Information***About This Task**

Use the following procedure to display the SPB-PIM Gateway interface information.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Interfaces** tab.

Interfaces Field Descriptions

Use the data in the following table to use the **Interfaces** tab.

Name	Description
IfIndex	Shows the interface index.
OperState	Shows the operational state of the interface.
AddressType	Shows the address type of the interface.
Address	Shows the address assigned to the interface.
AddressMask	Shows the address mask associated with the interface address.
HelloInterval	Configures the PIM HELLO transmission interval. The default is 30 seconds.
JoinPruneInterval	Configures the PIM JOIN PRUNE transmission interval. The default is 60 seconds.

*Display the SPB-PIM Gateway VLAN Information***About This Task**

Use the following procedure to display the SPB-PIM Gateway VLAN information.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a row for the VLAN, and then select **IP**.
5. Select the **SPB-PIM-GW** tab.

SPB-PIM-GW Field Descriptions

Use the data in the following table to use the **SPB-PIM-GW** tab.

Name	Description
OperState	Displays the current operational state of this SPB-PIM Gateway interface.
Address	Displays the primary IP address of this router on this SPB-PIM Gateway interface.
AddressMask	Displays the primary IP address mask of this router on this SPB-PIM Gateway interface.
HelloInterval	Configures the PIM HELLO transmission interval. This SPB-PIM Gateway VLAN level interval setting overrides the inherited global SPB-PIM Gateway HELLO interval setting. Setting the HELLO Interval to 0 causes the neighbors to never expire its neighborship with this local SPB-PIM Gateway interface.
JoinPruneInterval	Configures the PIM JOIN PRUNE transmission interval. This SPB-PIM Gateway VLAN level interval setting overrides the inherited global level JOIN PRUNE transmission interval setting.

*Display the SPB-PIM Gateway Neighbor Information***About This Task**

Use the following procedure to display the SPB-PIM Gateway neighbor information

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **SPB-PIM-GW**.
3. Select the **Neighbors** tab.

Neighbors Field Descriptions

Use the data in the following table to use the **Neighbors** tab.

Name	Description
IfIndex	Specifies the IfIndex for the interface which is used to reach this SPB-PIM Gateway neighbor.
AddressType	Specifies the address type of this SPB-PIM Gateway neighbor.
Address	Specifies the primary IP address of this router on this SPB-PIM Gateway neighbor.

Name	Description
UpTime	Specifies the time since this SPB-PIM Gateway neighbor last became a neighbor of the local router.
ExpiryTime	Specifies the minimum time remaining before this SPB-PIM Gateway neighbor times out.

Display the IP mroute Routes

About This Task

Use the following procedure to display multicast routes ingressing on either SPB-PIM Gateway interfaces or SPB multicast interfaces.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **Multicast**.
3. Select the **Routes** tab.

Routes Field Descriptions

Use the data in the following table to use the **Routes** tab.

Name	Description
Group	Displays the IP multicast group address for this entry that contains multicast routing information.
Source	Displays the network address that, when combined with the corresponding route SourceMask value, identifies the source that contains multicast routing information.
SourceMask	Displays the network mask that, when combined with the corresponding route Source value, identifies the multicast source.
UpstreamNeighbor	Shows the address of the upstream neighbor from which the IP datagrams from these sources are received. The address is 0.0.0.0 if the network is local.
Interface	Displays the interface, slot and portnumber, or VLAN ID where IP datagrams sent by these multicast sources to this multicast address are received.

Name	Description
ExpiryTime	Displays the amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: <ul style="list-style-type: none"> • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb (12) • spbpimgw(13)

SPB-PIM Gateway Interface Deployment Scenarios

This section provides information about different deployment scenarios for SPB-PIM gateway interface.

SPB-PIM Gateway Base Case Deployment Scenario

There are several different customer topology scenarios for the SPB-PIM Gateway (SPB-PIM GW) feature deployment . One of these scenarios, shown in [Figure 229](#) on page 3643, is fully described here, along with configuration details, and display information.

The deployment scenario described here has two domains:

- SPB domain
- PIM domain

The PIM network has 5 PIM routers:

- RP
- PIM-A1
- PIM-A2
- PIM-B
- PIM-C

The RP router is the PIM-SM rendezvous point. PIM router PIM-B has receiver host R2 attached to it and source S1 is connected to PIM router PIM-C.

The SPB domain has the following components:

- SPB-PIM Gateway Controller node
- Two Gateway nodes, BEB-A1 and BEB-A2
- BEB-A1 is connected to the PIM network through a SPB-PIM Gateway interface to the PIM router PIM-A1

- BEB-A2 is connected to the PIM network through a SPB-PIM Gateway interface to the PIM router PIM-A2

**Note**

PIM-A1 and PIM-A2 routers attached to the BEBs SPB-PIM Gateway interface have standard PIM configured on their side of the interface.

- The SPB cloud has a BEB-B to which the source S2 is connected
- The SPB cloud has a BEB-C to which a receiver R1 is connected

**Note**

You can place the controller anywhere in the SPB cloud. The controller can be in the boundary or the core. Anywhere there is a connection into the PIM network from the SPB network, there must be a Gateway node(s) and Gateway interface(s).




Figure 229: SPB-PIM Gateway base case configuration

The SPB-PIM Gateway nodes BEB-A1 and BEB-A2 from the SPB cloud are connected to the PIM network. The connection is established by SPB-PIM Gateway interface connections to PIM routers PIM-A1 and PIM-A2. A MSDP connection is established between the RP from the PIM domain and the SPB-PIM Gateway Controller from the SPB domain. A MSDP connection is established to exchange the multicast source information between the RP and the SPB Controller. Unicast routing or reachability is setup before establishing the MSDP connection between the RP and the Gateway controller. The Unicast setup is not shown in the previous figure.

*SPB-PIM Gateway Base Case Configuration Example***Before You Begin**

- The Shortest Path Bridging (SPB) infrastructure must be configured and setup in the SPB domain (not shown in this example)
- Protocol Independent Multicast (PIM) infrastructure must be configured and setup in the PIM domain (not shown in this example)
- Unicast routing table must be setup in the PIM domain to ensure reachability of the Multicast Source Discovery Protocol (MSDP) peer and source S2 from the SPB network
- Unicast routing table must be setup in the SPB domain to ensure reachability of the MSDP peer and source S1 from the PIM network

Example

Node: Gateway controller

Configure ISIS and SPBM:

```
Switch:1#configure terminal
Switch:1(config)#spbm
Switch:1(config)#router isis
Switch:1(config-isis)#system-id 0026.0026.0026
Switch:1(config-isis)#manual-area 01.0202.0303.04
Switch:1(config-isis)#spbm 1
Switch:1(config-isis)#spbm 1 nick-name 0.00.26
Switch:1(config-isis)#spbm 1 b-vid 10,20 primary 10
Switch:1(config-isis)#vlan create 10 name bvlan1 type spbm-bvlan
Switch:1(config)#vlan create 20 name bvlan2 type spbm-bvlan
Switch:1(config)#router isis enable
```

NNI Configuration:

```
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#isis
Switch:1(config-if)#isis spbm 1
Switch:1(config-if)#isis enable
Switch:1(config-if)#no shutdown
```

- Enable IPSC to setup Unicast route table. This setup ensures reachability of Rendezvous Point (RP) in the PIM network
- Create a loopback interface 2.0.2.2 to enable IPSC and MSDP originator-id
- Enable Multicast over Fabric Connect
- Configure static route and direct route redistribution
- Apply the static route and direct route redistribution

**Note**

The following configuration uses a static route to reach the RP.

```
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 2.0.2.2/32
Switch:1(config)#router isis
Switch:1(config-isis)#ip-source-address 2.0.2.2
Switch:1(config-isis)#spbm 1 ip enable
Switch:1(config-isis)#spbm 1 multicast enable
Switch:1(config-isis)#
```

```
Switch:1(config-isis)#redistribute static
Switch:1(config-isis)#redistribute static enable
Switch:1(config-isis)#redistribute direct
Switch:1(config-isis)#redistribute direct enable
Switch:1(config-isis)#
Switch:1(config-isis)#end
Switch:1#isis apply redistribute static
Switch:1#isis apply redistribute direct
```

MSDP Configuration:

Create an instance for the MSDP session. This IP interface is used for establishing an MSDP session with an RP in the PIM network. The source IP address used for the MSDP session must not be the newly created IP interface. The originator-id specifically configured for MSDP is used as the source IP address to establish the MSDP session. The originator-id is also used by the RP in the source active (SA) messages sent to the MSDP peers. The CLIP configured earlier is used as the originator-id.

```
Switch:1#configure terminal
Switch:1(config)#vlan create 2100 type port-mstprstp 0
Switch:1(config)#vlan members add 2100 1/3 portmember
Switch:1(config)#interface vlan 2100
Switch:1(config-if)#ip address 21.0.0.1/24
Switch:1(config-if)#exit

Switch:1(config)#ip msdp originator-id 2.0.2.2
Switch:1(config)#ip msdp enable
Switch:1(config)#ip msdp peer 21.0.0.2
Switch:1(config)#ip msdp peer 21.0.0.2 enable
```

SPB-PIM Gateway Controller configuration:

Enable SPB-PIM Gateway Controller

```
Switch:1(config)#router isis
Switch:1(config-isis)#spbm 1 multicast spb-pim-gw controller enable
```

Node: BEB-A1 (SPB-PIM Gateway)

Configure ISIS and SPBM:

```
Switch:1#configure terminal
Switch:1(config)#spbm
Switch:1(config)#
Switch:1(config)#router isis
Switch:1(config-isis)#system-id 0015.0015.0015
Switch:1(config-isis)#manual-area 01.0202.0303.04
Switch:1(config-isis)#spbm 1
Switch:1(config-isis)#spbm 1 nick-name 0.00.15
Switch:1(config-isis)#spbm 1 b-vid 10,20 primary 10
Switch:1(config-isis)#
Switch:1(config-isis)#vlan create 10 name bvlan1 type spbm-bvlan
Switch:1(config)#vlan create 20 name bvlan2 type spbm-bvlan
Switch:1(config)#router isis enable
```

NNI Configuration:

```
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#isis
Switch:1(config-if)#isis spbm 1
Switch:1(config-if)#isis enable
Switch:1(config-if)#no shutdown
```

- Enable IP Shortcuts (IPSC) to setup Unicast route table. This setup ensures reachability of sources in the PIM network
- Create a loopback interface 1.0.1.1 to enable IPSC
- Enable Multicast over Fabric Connect
- Configure static route and direct route redistribution
- Apply the static route and direct route redistribution

**Note**

The following configuration uses a static route to reach sources in the PIM network.

```
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 1.0.1.1/32
Switch:1(config-if)#router isis
Switch:1(config-isis)#ip-source-address 1.0.1.1
Switch:1(config-isis)#spbm 1 ip enable
Switch:1(config-isis)#spbm 1 multicast enable
Switch:1(config-isis)#
Switch:1(config-if)#router isis
Switch:1(config-isis)#redistribute static
Switch:1(config-isis)#redistribute static enable
Switch:1(config-isis)#redistribute direct
Switch:1(config-isis)#redistribute direct enable
Switch:1(config-isis)#
Switch:1(config-isis)#end
Switch:1#isis apply redistribute static
Switch:1#isis apply redistribute direct
```

SPB-PIM Gateway interface configuration:

- Create an IP interface
- Enable SPB-PIM Gateway on the IP interface

**Note**

For the following sample configuration, the SPB-PIM Gateway interface is on VLAN 2000 with IP address 20.0.0.1

```
Switch:1#configure terminal
Switch:1(config)#vlan create 2000 type port-mstprstp 0
Switch:1(config)#vlan members add 2000 1/1 portmember
Switch:1(config)#interface vlan 2000
Switch:1(config-if)#ip address 20.0.0.1/24
Switch:1(config-if)#ip spb-pim-gw enable
```

Enable SPB-PIM Gateway node functionality:

```
Switch:1(config-if)#router isis
Switch:1(config-isis)#spbm 1 multicast spb-pim-gw gateway enable
```

Similar configuration is done for the SPB-PIM Gateway node BEB-A2.

PIM Sparse Mode (PIM-SM) is enabled at the PIM routers PIM-A1 and PIM-A2 on the interfaces connecting SPB-PIM Gateway nodes BEB-A1 and BEB-A2. The SPB-PIM Gateway nodes BEB-A1 and BEB-A2 see the PIM routers PIM-A1 and PIM-A2 as PIM neighbors. The PIM routers PIM-A1 and PIM-A2 see the SPB-PIM Gateway nodes as PIM neighbors. The SPB-PIM Gateway nodes BEB-A1 and BEB-A2 have IP reachability to the PIM source S1 with PIM neighbors as the next hop.

The route to reach source S1 is distributed to the Gateway controller through IPSC. The Gateway controller uses this route information to select only one of the Gateways to which the source S1 will be assigned for a specific group. The Gateway node is the only node that can draw the source S1 stream into the SPB network on behalf of SPB receivers, by sending an SG Join across a Gateway Interface to the nexthop toward the source S1. This ensures that the data is not duplicated from multiple ingress interfaces from the PIM network. Other Gateway nodes that are not assigned as the Gateway to the source S1 will not establish multicast path.

When S1 from the PIM network sends traffic to G1, RP from the PIM network sends MSDP SA message for (S1,G1) to the Gateway Controller. If the Gateway Controller selects BEB-A1 as the Gateway for the foreign source S1 and group G1, the Gateway Controller assigns BEB-A1 to (S1,G1). The Gateway Controller then sends the Gateway assignment information to all the nodes. When BEB-A1 receives the assignment information, it sees that it is assigned as the Gateway to (S1,G1). The BEB-A1 then checks if the next hop to reach S1 is a valid PIM neighbor. If the next hop is a valid PIM neighbor, the BEB-A1 interacts with Multicast over Fabric Connect and advertises a sender TLV for the (S1,G1) into the SPB cloud. The BEB-A2 also receives the Gateway assignment information but silently saves the received Gateway assignment information since it is not the selected Gateway. If the interested receiver R1 is found at BEB-C, as part of Multicast over Fabric Connect processing, BEB-C sends receiver TLV for the group G1 to the advertising node, BEB-A1. Upon receiving this receiver TLV, BEB-A1 establishes the multicast stream through its Gateway interface which is upstream towards the PIM neighbor, by sending out a PIM SG Join message toward the source S1. This causes PIM-A1 node to forward multicast data from S1 to BEB-A1.

When the local source S2 (local to the SPB network) at BEB-B in the SPB network sends traffic to group G1, BEB-B advertises a sender TLV for (S1,G1). The controller sees this sender TLV and sends an MSDP SA message for (S2,G1) to the RP in the PIM network.

**Note**

The controller does not send SA messages for (S1,G1) to the PIM network since S1 is a foreign source.

When PIM network receiver R2 is interested in group G1, the PIM router PIM-B sends PIM a (*,G) Join message to the RP. RP in turn sends (S2,G1) Join towards the source S2. The unicast IP reachability to source S2 which is setup in the RP is used for sending (S2,G1) joins hop-by-hop towards the source. From the RP point of view, the next hop to reach the source S2 is one of the PIM routers PIM-A1 or PIM-A2 (depending on the unicast route table next hop address). For this example, the next hop is PIM-A1. The RP sends the (S2,G1) Join message towards the PIM-A1. PIM-A1 then sends an SG Join

to BEB-A1. Upon receiving the Join for (S2,G1) from the PIM network, BEB-A1 sends a receiver TLV into the SPB network to the S2 advertising router BEB-B. When the BEB-B receives the receiver TLV, the BEB-B establishes the multicast stream from source S2 toward the receiver.

Source Specific Multicast

PIM-SSM does not use a Rendezvous Point to centralize the receivers and sources. A PIM-SSM router which has a receiver for a group multicast address in the SSM address range joins directly to a source for that group by sending an SG join toward the source, not a *G join toward the RP for the group. Because of this, MSDP is not required in order for the PIM Network to learn of an SSM range stream in the SPB network. However, in order for the SPB network to know where a PIM SSM source resides, it must statically configure S1,SSM-G1 at the controllers. In this way, a Gateway can be chosen for the stream, even in the absence of MSDP.

The following figure shows a non-MSDP SSM environment where stream PIM network source S1, for an SSM group, must be statically configured at the SPB Controllers:

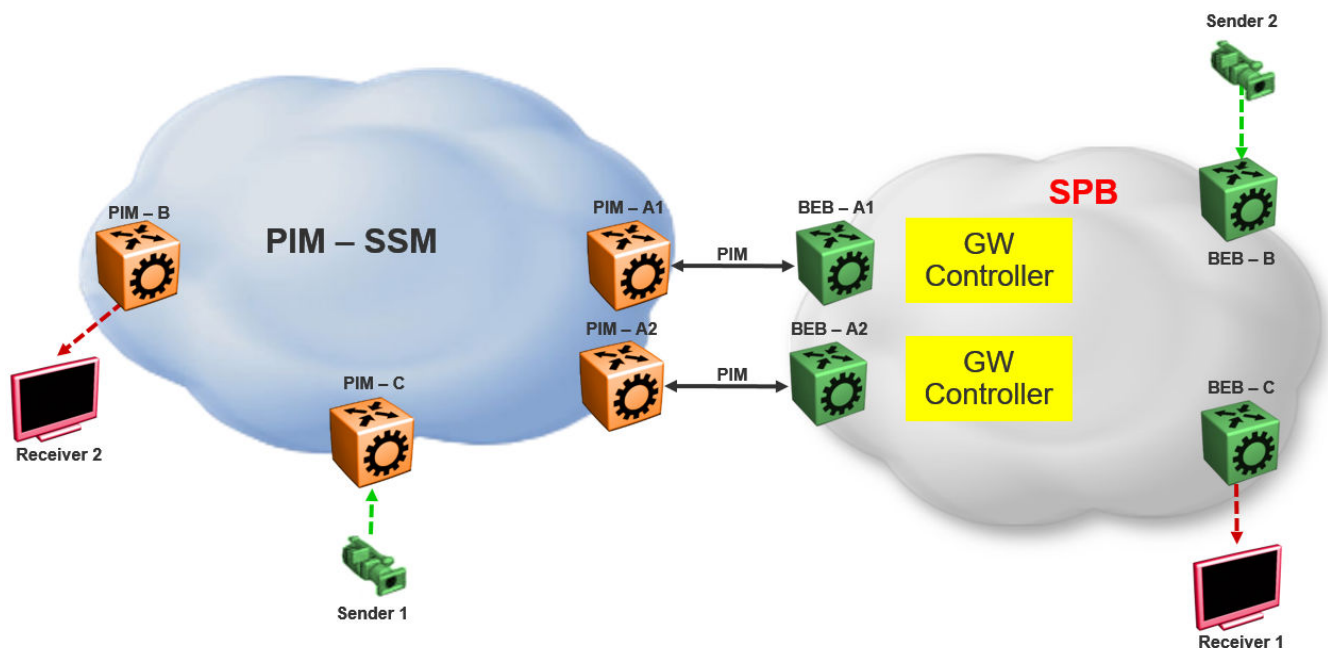


Figure 230: Static configuration of SSM groups in Controllers

Peer Mesh Group

The following figure shows the Peer Mesh Group configuration:



Note

Controllers within a single SPB network must never peer with each other, regardless of whether mesh groups exist. In addition, both Controllers must have the same peerings configured with other networks RPs.




Figure 231: Peer Mesh Group

MSDP Peer Mesh Group Configuration Example

Example

Configure MSDP:

If the MSDP speakers are fully meshed, the speakers can be configured into a mesh group in order to prevent excessive SA forwarding and RPF checks. To configure a mesh group, specify a name along with the MSDP peer. For example, on router RP1, configure a mesh group which includes RP2, RP3, and both Gateway controllers. On RP2, configure the same mesh group with members RP1, RP3, and both Gateway controllers. On RP3, configure the same mesh group with members RP1, RP2, and both Gateway controllers. On each Gateway controller, configure the same mesh group with members RP1, RP2, and RP3, but never with another Gateway controller in the same SPB domain.

```
Switch:1(config)#ip msdp originator-id 2.0.2.2
Switch:1(config)#ip msdp enable
Switch:1(config)#ip msdp peer 21.0.0.2 enable
Switch:1(config)#ip msdp mesh-group mgTest 21.0.0.2
```

Multi Domain

The following figure shows a multi domain scenario, where two PIM domains and one SPB domain share multicast streams:




Figure 232: Multi domain configuration

SPB Domain Interconnect

The following figure shows the SPB domain interconnect configuration. In this scenario, two SPB domains are connected by PIM Gateway interfaces, and there is no traditional PIM Network involved. The Controllers from each SPB domain form MSDP adjacencies the Controllers in the other domain (but not within the same domain) other to share their multicast sources. The SPB-PIM Gateway nodes see the other SPB-PIM Gateway nodes as PIM neighbors on the SPB-PIM Gateway interfaces.

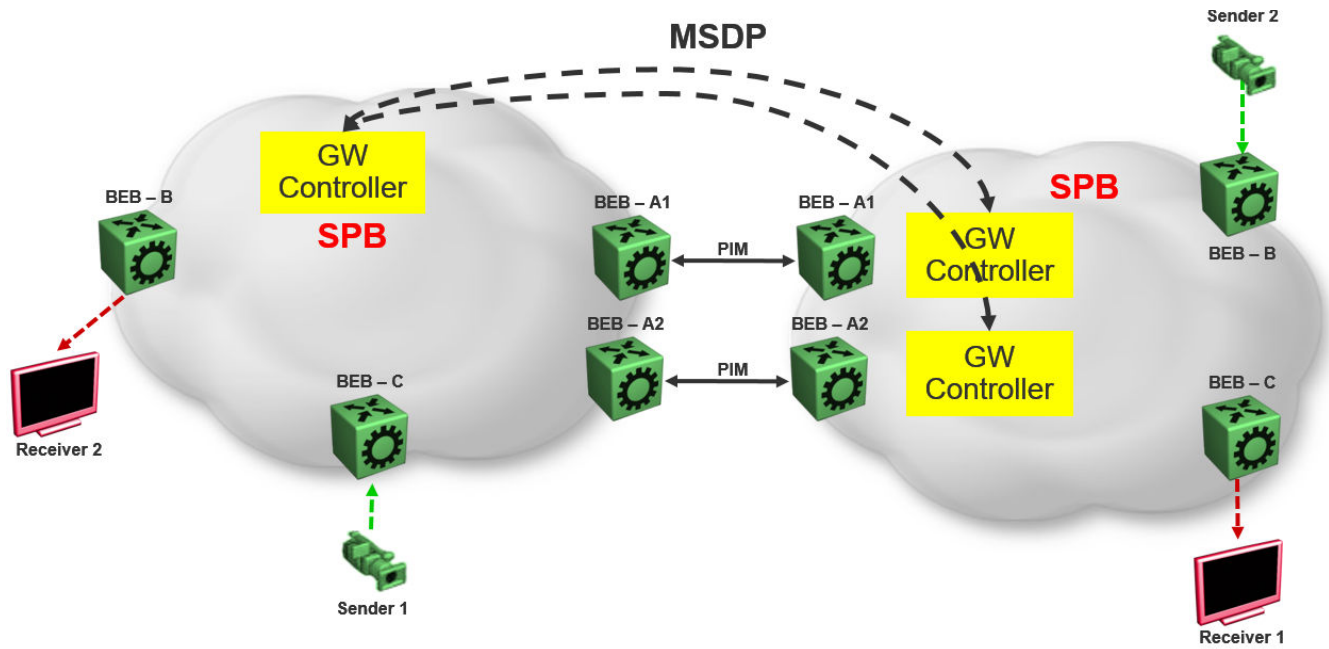


Figure 233: SPB domain interconnect



TACACS+

[TACACS+ Fundamentals](#) on page 3653

[TACACS+ configuration using CLI](#) on page 3664

[TACACS+ configuration using EDM](#) on page 3673

[TACACS+ Configuration Example](#) on page 3678

Table 308: TACACS+ product support

Feature	Product	Release introduced
TACACS+	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
TACACS+ secure communication using IPSec for IPv4	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

TACACS+ Fundamentals

The switch supports the TACACS+ client. TACACS+ is a remote authentication protocol that provides centralized validation of users who attempt to gain access to a router or Network Access Server (NAS).

The TACACS+ feature is a client and server-based protocol that allows the switch to accept a user name and password and send a query to a TACACS+ authentication server, sometimes called a TACACS+ daemon. The TACACS+ server allows access or denies access based on the response by the client.

The TACACS+ feature facilitates the following services:

- Login authentication and authorization for CLI access through Secure Shell (SSH), Telnet, or serial port.
- Login authentication for web access through EDM.
- Command authorization for CLI through SSH, Telnet, or serial port.
- Accounting of CLI through SSH, Telnet, and serial port.

The following figure displays the basic layout of the switch and the TACACS+ server for a remote user connection.




Figure 234: Switch and TACACS+ server

The TACACS+ feature uses Transmission Control Protocol (TCP) for its transport to ensure reliable delivery of packets. TACACS+ provides security by encrypting all traffic between the switch, which acts as the Network Access Server, and the TACACS+ server.

TACACS+ is a newer version of TACACS and provides separate authentication, authorization, and accounting (AAA) services. TACACS+ does not support earlier versions of TACACS.

TACACS+ Operation

The switch acts as an NAS to provide a connection to a single user, to a network, subnetwork or interconnected networks. The switch acts as a gateway to guard access to the TACACS+ server and network. Encryption relies on a secret key that is known to the client and the TACACS+ server.

Similar to the Remote Access Dial-In User Services (RADIUS) protocol, TACACS+ provides the ability to centrally manage the users who want to access a remote device. TACACS+ provides management of remote and local users who try to access a device through:

- Secure Shell (SSHv2)
- Telnet
- serial port
- web management

A TACACS+ daemon, which typically runs on a UNIX or Windows NT workstation, maintains the TACACS+ authentication, authorization, and accounting services.

You configure users in the TACACS+ server. If you enable authentication, authorization, and accounting services, the following occurs:

- During the logon process, the TACACS+ client initiates the TACACS+ authentication session with the TACACS+ server.
- After successful authentication the TACACS+ client initiates the TACACS+ authorization session with the TACACS+ server. This is transparent to the user. The switch receives the user access level after a successful TACACS+ authorization. The TACACS+ server authorizes every command the user issues if TACACS+ command authorization is enabled for that user access level.
- After successful authorization, if you enable TACACS+ accounting, the TACACS+ client sends accounting information to the TACACS+ server.

A TACACS+ session establishes with the server in one of two ways:

- Multi-connection mode (also known as per-session): For every authentication, authorization, and accounting (AAA) request the switch establishes a session with the TACACS+ server, and then after the request finishes, the session is torn down. Multi-connection mode is the default mode.
- Single-connection mode: The first AAA request establishes the session, which is only torn down if TACACS+ is disabled or due to inactivity.

TACACS+ Architecture

You can connect the TACACS+ server to the switch:

- In-band through one of the data ports.
- Out-of-band through the management port.

Connect the TACACS+ server through a local interface. Management PCs can reside on an out-of-band management Ethernet port, or on the corporate network. Place the TACACS+ server on the corporate network so you can route it to the switch.

Before you configure the switch, you must configure at least one TACACS+ server and a key.

The TACACS+ server and the switch must have the same:

- Encryption key
- Connection mode (single connection or per-session connection. Per-session connection is the same as multi-connection mode.)
- TCP port number

You can configure a secondary TACACS+ server for backup authentication. You specify the primary authentication server when you configure the switch.

Authentication, Authorization, and Accounting

A fundamental feature of TACACS+ is the separation of authentication, authorization, and accounting (AAA) services, which allows you to selectively implement one or more TACACS+ services.

TACACS+ Authentication

TACACS+ authentication provides control of authentication through login and password.

Authentication uses a database of users and passwords to determine:

- who a user is
- to grant or deny the user access to the NAS



Important

Prompts for log on and password occur prior to the authentication process. If TACACS+ fails because no valid servers exist, the device uses the user name and password from the local database. If TACACS+ or the local database returns an access denied packet, the authentication process stops. The device attempts no other authentication methods.

The following figure illustrates the authentication process after the user attempts to make a connection to the switch.




Figure 235: Authentication process

TACACS+ Authorization

The transition from TACACS+ authentication to the authorization phase is transparent to the user. After successful completion of the authentication session, an authorization

session starts with the authenticated user name. The authorization session provides access level functionality.

Authorization cannot occur without authentication.

Authorization:

- determines what a user can do
- allows administrators fine-grained control over the capabilities of users during sessions

The following figure illustrates the authorization process after the user attempts to make a connection to the switch..

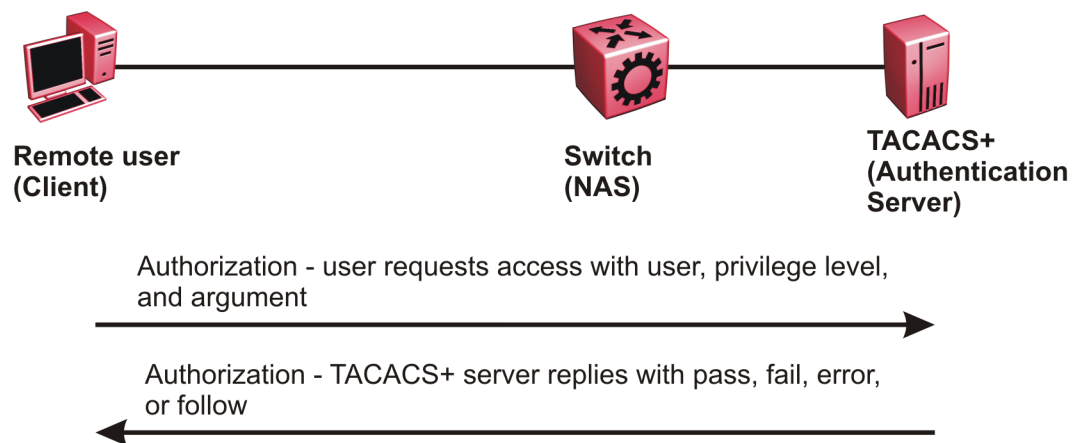


Figure 236: Authorization process

Authorization determines what a user can do. Authorization gives you the ability to limit network services to certain users and to limit the use of certain commands to certain users. The TACACS+ feature enhances the security by tightly policing the command execution for a particular user. After you enable command authorization, all commands, no matter the access level to which they belong, are sent to the TACACS+ server for authorization. Authorization cannot occur without first enabling authentication. You must configure command authorization globally and at individual access levels.

Two kinds of authorization requests exist:

1. Login authorization: Login authorization happens immediately after authentication and is transparent to the user. When the user logs on to the device, authorization provides the user access level. With log on, the device does not send a command to the TACACS+ server. You cannot configure login authorization.
2. Command authorization: When you configure command authorization for a particular level, all commands that you issue are sent to the TACACS+ server for authorization. The device can only issue the commands the TACACS+ server

authorizes. You need to configure command authorization globally and at individual access levels, which are visible to the users.

**Note**

You must verify that the switch can reach the TACACS+ server and that you configure TACACS+ properly before you enable command authorization.

If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow the user to issue any command that has privilege level command authorization enabled. In such a case, the user can only issue logout and exit commands.

If a user tries to log in and the TACACS+ server does not exist or is not reachable, then, as discussed before, a local database in the switch authenticates the user. The switch authorizes a locally authenticated user and a locally authenticated user is not eligible for TACACS+ command authorization.

After the switch requests authorization, the logon credentials are sent to the TACACS+ daemon for authorization. If logon authorization fails, the user receives a permission denied message.

If TACACS+ logon authorization succeeds, the switch uses information from the user profile, which exists in the local user database or on the TACACS+ server, to configure the session for the user.

After you enable TACACS+ command authorization all commands are visible to all users; however, the user can only issue those commands that the TACACS+ server configuration allows.

The switch cannot enforce command access level. The TACACS+ server returns an access level to the switch. The switch allows the user to access the switch according to the access level. The device grants the user access to a command only if the profile for the user allows the access level.

You preconfigure command authorization on the TACACS+ server. You specify a list of regular expressions that match command arguments, and you associate each command with an action to deny or permit.

All members in a group have the same authorization. If you place a user in a group, the daemon looks in the group for authorization parameters if it cannot find them in the user profile.

TACACS+ Accounting

TACACS+ accounting enables you to track the services users access and the amount of network resources users consume.

TACACS+ accounting allows you to track:

- what a user does
- when a user does certain actions

The accounting record includes the following information:

- User name
- Date
- Start/stop/elapsed time
- Access server IP address
- Reason

You can use accounting for an audit trail, to bill for connection time or resources used, or for network management. TACACS+ accounting provides information about user sessions using the following connection types: Telnet, SSH, and web-based management.

With separation of AAA, accounting can occur independently from authentication and authorization.

The following figure illustrates the accounting process after the user attempts to make a connection to the switch..

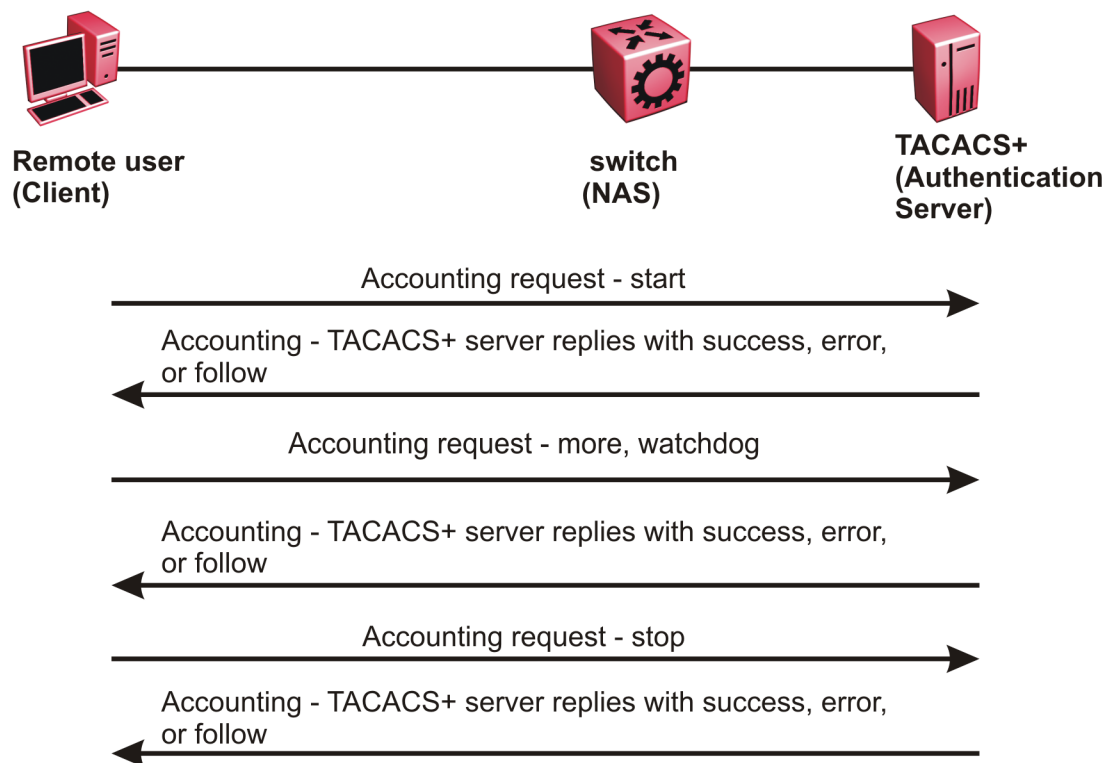


Figure 237: Accounting process

After you enable accounting, the switch reports user activity to the TACACS+ server in the form of accounting records. Each accounting record contains accounting attribute value (AV) pairs. AV pairs are strings of text in the form “attribute-value” sent between

the switch and a TACACS+ daemon as part of the TACACS+ protocol. The TACACS+ server stores the accounting records.

You cannot customize the set of events the switch monitors and logs with TACACS+ accounting. TACACS+ accounting logs the following events:

- User logon and logoff
- Logoff generated because of activity timeout
- Unauthorized command
- Telnet session closed (not logged off)

Privilege Level Changes at Runtime

You can change your privilege level at runtime with the **tacacs switch level** command.

You need to configure separate profiles in the TACACS+ server configuration file for the switch level. The switch supports only levels 1 to 6 and level 15. The switch uses the profile when you issue the command **tacacs switch level <1-15>**. As part of the profile, you specify a user name, level, and password. To preconfigure a dummy user for that level on the TACACS+ daemon, the format of the user name for the dummy user is **\$enab<n>\$**, where **<n>** is the privilege level to which you want to allow access.

The following is an example of a TACACS+ server profile, which you configure on the TACACS+ server:

```
user = $enab6$ {
member = level6
login = cleartext get-me-on-6
}
```

The following table maps user accounts to TACACS+ privilege level.

Switch access level	TACACS+ privilege level	Description
NONE	0	If the TACACS+ server returns an access level of 0, the user is denied access. You cannot log into the device if you have an access level of 0.
READ ONLY	1	Permits you to view only configuration and status information.
LAYER 1 READ WRITE	2	Permits you to view most of the switch configuration and status information and change physical port settings.

Switch access level	TACACS+ privilege level	Description
LAYER 2 READ WRITE	3	Permits you to view and change configuration and status information for Layer 2 (bridging and switching) functions.
LAYER 3 READ WRITE	4	Permits you to view and change configuration and status information for Layer 2 and Layer 3 (routing) functions.
READ WRITE	5	Permits you to view and change configuration and status information across the switch. This level does not allow you to change security and password settings.
READ WRITE ALL	6	Permits you to have all the rights of read-write access and the ability to change security settings, including command line interface (CLI) and web-based management user names and passwords, and the SNMP community strings.

Switch access level	TACACS+ privilege level	Description
NONE	7 to 14	If the TACACS+ server returns an access level of 7 to 14, the user is denied access. You cannot log into the device if you have an access level of 7 to 14.
READ WRITE ALL	15	Permits you to have all the rights of read-write access and the ability to change security settings, including command line interface (CLI) and web-based management user names and passwords, and the SNMP community strings. Note: Access level 15 is internally mapped to access level 6, which ensures consistency with other vendor implementations. The switch does not differentiate between an access level of 6 and an access level of 15.



Note

If you enable enhanced secure mode with the **boot config flags enhancedsecure-mode** command, you enable different access levels, along with stronger password complexity, length, and minimum change intervals. With enhanced secure mode enabled, the switch supports the following access levels for RADIUS authentication:

- Administrator
- Privilege
- Operator
- Auditor
- Security

The switch associates each user name with a certain role and appropriate authorization rights to view and configure commands. For more information, see [Enhanced Secure Mode](#) on page 266.

TACACS+ Command Authorization

After you enable TACACS+ authorization, the current privilege-level to command mapping on the switch is no longer relevant because the TACACS+ server has complete

responsibility for command authorization. TACACS+ authorization provides access to the system based on user name, not based on privilege level.

After you enable TACACS+ command authorization for a particular privilege level, and a user with that privilege level logs on, the user can access commands based on his user name.

TACACS+ Switch Level and TACACS+ Switch Back Commands

The user can only issue the **tacacs switch level** command after TACACS+ authenticates the user. Locally authenticated users, which means users authenticated only by the switch and not by the TACACS+ server, cannot use the **tacacs switch level** command.

Consider a user, called X, with a privilege level of 4, who uses the **tacacs switch level <1-15>** command to change the privilege level from 4 to 6.

If user X successfully changes the switch level to 6, the user name changes from X to "\$enab6\$", and the privilege level changes from 4 to 6. If TACACS+ command authorization is enabled for privilege level 6, then the TACACS+ server authorizes commands issued based on the rules defined for (dummy) user "\$enab6\$".

If TACACS+ command authorization is not enabled for privilege level 6, then the switch locally authorizes the user X based on the privilege level of the user.

The user can return to his previous privilege level using the **tacacs switch back** command. In the preceding scenario, if the user issues the **tacacs switch back** command, the user name changes for user X from "\$enab6\$" to X, and the privilege level changes from 6 to 4.

TACACS+ switch level supports up to eight levels, and TACACS+ switch level allows a user to switch level up to eight times from his original privilege level. The switch stores all of the previous privilege levels in the same order in which the user switches levels. After switching eight times, if the user tries to switch a level the ninth time, the following error message displays:

```
Only allowed to switch level 8 times!
```

The user can switch back to his previous privilege levels using the **tacacs switch back** command. The **tacacs switch back** command switches back in the reverse order in which you issued the **tacacs switch level** command. Consider a user who switched levels from 4 to 5, and then to 6. If the user used the **tacacs switch back** command, the user first moves from 6 to 5, and then using the **tacacs switch back** command again moves from 5 to 4.



Note

If you want to switch to a privilege level 'X' using **tacacs switch level <1-15>** command, you must create a user "\$enabX\$" on the TACACS+ server. X is the privilege level that you want to change.

TACACS+ Switch Level Functionality:

The following table explains TACACS+ switch level functionality.

User logs in with	TACACS+ server available	Result
TACACS+ authentication	Yes	The user can issue the tacacs switch level <1-15> command.
Local authentication	No	The user cannot issue the tacacs switch level <1-15> command.
Local authentication	Yes	Even if a TACACS+ server becomes reachable, the user remains locally authenticated and cannot issue the tacacs switch level <1-15> command.

TACACS+ Command Authorization Functionality:

The following table explains TACACS+ command authorization functionality.

User logs in with	Command authorization	Result
Local authentication		The switch authorizes the user locally.
TACACS+ authentication	Not enabled for the logged-in level.	The switch authorizes the user locally. If the server connection is lost, the switch authorizes the user locally.
TACACS+ authentication	Enabled for the logged-in level.	The TACACS+ server authorizes the user. If the server connection is lost, the user can only issue exit and logout commands.

**Note**

A user who configures TACACS+ is locally authenticated and authorized by the switch, so even after the user configures TACACS+, the switch continues to locally authorize the user.

TACACS+ and RADIUS Differences

TACACS+ and RADIUS are security protocols that you can use on network devices.

You can enable TACACS+ and RADIUS together. However, TACACS+ has a higher priority. If the TACACS+ server is not available the authentication is sent to RADIUS, if RADIUS is enabled. However, if TACACS+ authentication fails, then requests are not sent to RADIUS.

Following is a list of differences between TACACS+ and RADIUS.

TACACS+	RADIUS
Separates Authorization, Authentication and Accounting (AAA). As a result, you can selectively implement one or more TACACS+ services. With TACACS+ you can use different servers for each service.	Combines authentication and authorization.
Uses TCP. TCP is connection-oriented. TCP immediately indicates if a server crashes or is not running. TCP offers an acknowledgement that a request has been received.	Uses UDP. UDP is best-effort delivery. RADIUS uses re-transmit attempts and timeouts to make up for the support TCP has.
Encrypts the entire body of the packet, which includes the password and user name.	Encrypts only the password from the client to the server.
Used for administrator access. Usually used for administrator access to network devices.	Used for subscriber access. Usually used to authenticate remote users to a network.
Can control which access level of commands a user or group can access.	Cannot control which access level of commands can be used.

TACACS+ Feature Limitations

TACACS+ does not support the following features:

- Point-to-Point Protocol (PPP) authentication and accounting
- IPv6 for TACACS+
- S/KEY (One Time Password) authentication
- PAP/CHAP/MSCHAP authentication methods
- The FOLLOW response of a TACACS+ server, in which the AAA services are redirected to another server. The response is interpreted as an authentication failure.
- User capability to change passwords at runtime over the network. The system administrator must change user passwords locally, on the server.
- TACACS+ command authorization when the user accesses the switch through EDM and SNMP.
- Restriction of command authorization for a specific kind of access. After you enable command authorization, command authorization applies for Telnet, SSH, and serial-port access. You cannot restrict command authorization to just one kind of access.

If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not permit the user to execute any command that has privilege level command authorization enabled.

TACACS+ configuration using CLI

This section provides procedures to configure TACACS+ using the command line interface (CLI).

Enable TACACS+

Enable TACACS+ globally on the switch.

The switch supports the TACACS+ client. TACACS+ is a security application implemented as a client and server-based protocol that provides centralized validation of users who attempt to gain access to a router or network access server (the switch).

By default, TACACS+ is disabled.

Before You Begin

- You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch are available.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable TACACS+ globally:
`tacacs protocol enable`
3. Disable TACACS+ globally:
`no tacacs protocol enable`
`default tacacs protocol enable`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs protocol enable
```

Add a TACACS+ Server

Add a primary and secondary TACACS+ server and specify the authentication process.

If you have a backup server configured, the AAA request goes to the backup server if the primary server is not available.

About This Task

The TACACS+ server and the switch must have the same:

- Encryption key
- Connection mode (single connection or per-session connection. Per-session connection is the same as multi-connection mode)
- TCP port number

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Add a primary TACACS+ server with an encryption key:

```
tacacs server host {A.B.C.D} key WORD<0-128>
```

3. (Optional) Configure the parameters for the primary TACACS+ server as required.

- a. (Optional) Specify a single connection to maintain a constant connection between the switch and the TACACS+ daemon:

```
tacacs server host {A.B.C.D} single-connection
```



Note

The TACACS+ daemon must also support this mode. If you do not configure this, the switch uses the default connection type, which is the per-session connection. Per-session is the same as multi-connection mode.

- b. (Optional) Specify the TCP port to use when the switch connects to the TACACS+ daemon:

```
tacacs server host {A.B.C.D} port <1-65535>
```

The default port is 49.

- c. (Optional) Specify the period of time (in seconds) the switch waits for a response from the TACACS+ daemon before it times out and shows an error:

```
tacacs server host {A.B.C.D} timeout <10-30>
```

4. Specify the IP address of the secondary TACACS+ server and specify an encryption key:

```
tacacs server secondary-host {A.B.C.D} key WORD<0-128>
```

5. (Optional) Configure the optional parameters on the secondary TACACS+ server as required.

- a. (Optional) Specify a single connection for the secondary TACACS+ server to maintain a constant connection between the switch and the TACACS+ daemon:

```
tacacs server secondary-host {A.B.C.D} single-connection
```



Note

The TACACS+ daemon must also support this mode. If you do not configure this, the switch uses the default connection type, which is the per-session connection. Per-session is the same as multi-connection mode.

- b. (Optional) Specify the TCP port to use when the switch connects to the TACACS+ daemon:

```
tacacs server secondary-host {A.B.C.D} port <1-65535>
```


- c. (Optional) Specify the period of time (in seconds) the switch waits for a response from the TACACS+ daemon before it times out and shows an error:

```
tacacs server secondary-host {A.B.C.D} timeout<10-30>
```
6. Display the status of the TACACS+ configuration:

```
show tacacs
```
7. (Optional) Delete a primary TACACS+ server:

```
no tacacs server host{A.B.C.D} [single-connection]
```
8. (Optional) Delete a backup TACACS+ server:

```
no tacacs server secondary-host{A.B.C.D} [single-connection]
```
9. (Optional) Configure a primary TACACS+ server or secondary TACACS+ server to the default settings:

```
default tacacs server {A.B.C.D} [port][single-connection][timeout]
```

Example

Configure the primary server with the IP address 192.0.2.1 and the encryption key 1dt4ly. Configure the secondary server with the IP address 198.51.100.2 with the same encryption key 1dt4ly. Display the configuration to ensure proper configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs server host 192.0.2.1 key 1dt4ly
Switch:1(config)#tacacs server secondary-host 198.51.100.2 key 1dt4ly
Switch:1(config)#show tacacs
```

Global Status:

```

global enable : true

authentication enabled for : cli

accounting enabled for : none

authorization : disabled

User privilege levels set for command authorization : None

```

Server:

```

create :

```

Prio	Status	Key	Port	IP address	Timeout	Single	Source
SourceEnabled							
Primary	Conn	*****	49	192.0.2.1	10	false	0.0.0.0
false							
Backup	NotConn	*****	49	198.51.100.2	10	false	0.0.0.0
false							

```

Switch:1(config)#no tacacs server host 192.0.2.1
Switch:1(config)#no tacacs server secondary-host 198.51.100.2

```

Variable Definitions

The following table defines parameters for the **tacacs server host** and the **tacacs server secondary-host** commands.

Variable	Value
<i>{A.B.C.D}</i>	Specifies the IP address of the TACACS+ server you want to add. Only IPv4 addresses are valid.
<i>key WORD</i> <0-128>	Configures the authentication and encryption key for all TACACS+ communications between the device and the TACACS+ server. If the key length is zero, that indicates no encryption is used. You must configure the same encryption key for the TACACS+ server and the switch.
<i>port</i> <1-65535>	Configures the TCP port, on which the client establishes a connection to the server. A value of 0 indicates the system specified default value is used. The default is 49. You must configure the same TCP port for the TACACS+ server and the switch.
<i>single-connection</i>	Specifies if the TCP connection between the device and the TACACS+ server is a single connection. If you specify the single connection parameter, the connection between the switch and the TACACS+ daemon remains open, which is more efficient because it allows the daemon to handle a higher number of TACACS+ operations. The single-connection is torn down if TACACS+ is disabled due to inactivity. If you do not configure this, the switch uses the default connection type, which is the multi-connection. With the multi-connection, the connection opens and closes each time the switch and TACACS+ daemon communicate. Note: You must configure the same connection mode for the TACACS+ server and the switch. To enable single-connection, the TACACS+ daemon has to support this mode as well.
<i>timeout</i> <10-30>	Configures the maximum time, in seconds, to wait for this TACACS+ server to reply before it times out. The default value is 10 seconds.

Configure TACACS+ Authentication

Configure what application TACACS+ authenticates: CLI, web, or all.

TACACS+ authentication provides control of authentication through login and password.

By default, CLI authentication is enabled.

Before You Begin

- You must enable TACACS+ globally for TACACS+ authentication to function.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure TACACS+ authentication:

```
tacacs authentication <all/cli/web>
```
3. (Optional) Disable TACACS+ authentication:

```
no tacacs authentication <all/web>
```
4. (Optional) Configure TACACS+ authentication to the default settings (default is cli authentication enabled):

```
default tacacs authentication <all/cli/web>
```
5. Display the configuration:

```
show tacacs
```

Example

Configure TACACS+ to authenticate CLI and display the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs authentication cli
Switch:1(config)#show tacacs
Global Status:

    global enable : true

    authentication enabled for : cli

    accounting enabled for : none

Server:

    create :

Prio    Status  Key      Port  IP address  Timeout  SingleSource  Source  Enabled
Primary Conn    * 49    192.0.2.1    10    false        0.0.0.0 false
Backup  NotConn * 49    198.51.100.2 10    false        0.0.0.0 false
```

Variable Definitions

The following table defines parameters for the **tacacs authentication** command.

Variable	Value
<i>all</i>	Specifies TACACS+ authentication for all applications. By default, CLI authentication is enabled.
<i>cli</i>	Specifies TACACS+ authentication for command line connections. By default, CLI authentication is enabled.
<i>web</i>	Specifies TACACS+ authentication for web connections. By default, CLI authentication is enabled.

Configure TACACS+ Accounting

Determines for which applications TACACS+ collects accounting information. Use TACACS+ accounting to track the services that users access and the amount of network resources that users consume. If unassigned, TACACS+ does not perform the accounting function.

If enabled, TACACS+ accounting logs the following events:

- User log on and log off
- Log off generated because of activity timeout
- Unauthorized command
- Telnet session closed (not logged off)

If unassigned, TACACS+ does not perform the accounting function. No default value exists.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable TACACS+ accounting:
`tacacs accounting enable cli`
3. (Optional) Disable TACACS+ accounting:
`no tacacs accounting cli`

`tacacs accounting disable [cli]`

Example

Enable TACACS+ accounting:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs accounting enable cli
```

Configure Command Authorization with TACACS+

Use this procedure to enable TACACS+ authorization for a particular privilege level. Use this option to limit the use of certain commands to certain users.

If command authorization fails, the following log message displays: Command <command> not authorized for user <username>.

By default, command authorization is disabled on the switch. The default for the command authorization level is none.

Before You Begin

- You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch are available. You must verify that the switch can reach the TACACS+ server and that you configure TACACS+ properly before you

enable command authorization. If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow you to issue any command that has privilege level command authorization enabled. If the switch cannot reach the TACACS+ server, you can only issue logout and exit commands.

- To use TACACS+ authorization, you must enable TACACS+ authentication.

About This Task

Two kinds of authorization requests exist:

1. Login authorization: Login authorization happens immediately after authentication when the user logs on to the device, authorization provides the user access level. You cannot configure login authorization.
2. Command authorization: When you configure command authorization for a particular level, all commands that you issue are sent to the TACACS+ server for authorization. You need to configure command authorization globally and at individual access levels.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Enable TACACS+ authorization:
`tacacs authorization enable`
3. Configure TACACS+ privilege level for TACACS+ command authorization:
`tacacs authorization level <1-6>`

`tacacs authorization level all`

`tacacs authorization level none`
4. (Optional) Disable TACACS+ authorization:
`tacacs authorization disable`

`default tacacs authorization`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs authorization enable
Switch:1(config)#tacacs authorization level 6
```

Variable Definitions

The following table defines parameters for the **tacacs authorization** command.

Variable	Value
<code>level <1-6></code>	Enables command authorization for a specific privilege level. The default for the command authorization level is none.
<code>level all</code>	Enables command authorization for all privilege levels. The default for the command authorization level is none.
<code>level none</code>	Disables command authorization for all privilege levels. The default for the command authorization level is none.

Change Privilege Levels at Runtime

Users can change their privilege levels at runtime. The privilege level determines what commands a user can access through TACACS+ server authorization.

A user can only use the **tacacs switch level** command, after TACACS+ authenticates the user. Locally authenticated users, which means users authenticated only by the switch and not by the TACACS+ server, cannot use the **tacacs switch level** command.

Before You Begin

- You need to configure separate profiles in the TACACS+ server configuration file for switch level. As part of the profile, you specify a user name, level, and password.

About This Task

After you enable TACACS+ authorization, the current privilege-level to command mapping on the switch is no longer relevant because the TACACS+ server has complete responsibility for command authorization. TACACS+ authorization provides access to the system based on user name, not based on privilege level.

After you enable TACACS+ command authorization for a particular privilege level, and a user with that privilege level logs on, the user can access commands based on his user name.



Note

If you want to switch to a privilege level 'X' using **tacacs switch level <1-15>** command, you must create a user "\$enabX\$" on the TACACS+ server. X is the privilege level to which you want to change.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
- Change the privilege level for a user at runtime:

```
tacacs switch level <1-15>
```
- Return to the original privilege level:

```
tacacs switch back
```

Example

Change the privilege level for a user at runtime. Return to the original privilege level:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs protocol enable
Switch:1(config)#tacacs switch level 5
Password:*****
```

Return to the original privilege level:

```
Switch:1(config)#tacacs switch back
```

Variable Definitions

The following table defines parameters for the **tacacs switch** command.

Variable	Value
<i>level</i> <1-15>	<p>Specifies the privilege level you want to access. You can change your privilege level at runtime by using this parameter. You are prompted to provide the required password. If you do not specify a level in the command, the administration level is selected by default.</p> <p>Note: For switch level, you need to configure separate profiles in the TACACS+ server configuration file. As part of the profile, you specify a user name, level, and password. To preconfigure a dummy user for that level on the TACACS+ daemon, the format of the user name for the dummy user is \$enab<n>\$, where <n> is the privilege level to which you want to allow access.</p>
<i>back</i>	Specifies that you want to return to the original privilege level.

TACACS+ configuration using EDM

This section provides procedures to configure TACACS+ using Enterprise Device Manager (EDM).

Configure TACACS+ Globally

Before You Begin

- You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch (network access server) are available.

You must verify that the switch can reach the TACACS+ server and that you configure TACACS+ properly before you enable command authorization.

- If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow the user to issue any command that has privilege level command authorization enabled. In such a case, the user can only issue logout and exit commands.

- You must enable TACACS+ globally for TACACS+ authentication to function.
- You must enable TACACS+ authentication for TACACS+ authorization to function.

About This Task

Enable TACACS+ globally on the switch. TACACS+ is a security application implemented as a client and server-based protocol that provides centralized validation of users. By default, TACACS+ is disabled.

Configure what application TACACS+ authenticates. TACACS+ authentication provides control of authentication through login and password dialog, challenge and response. By default, CLI authentication is enabled.

After authentication is complete, the switch starts the authorization process. By default, command authorization is disabled on the switch. The default for the command authorization level is none. If command authorization fails, the following log message displays: `Command <command> not authorized for user <username>.`

Two kinds of authorization requests exist:

1. Login authorization: Login authorization happens immediately after authentication when the user logs on to the device, authorization provides the user access level. You cannot configure login authorization.
2. Command authorization: When you configure command authorization for a particular level, all commands that you issue are sent to the TACACS+ server for authorization. You need to configure command authorization globally and at individual access levels.

Enable TACACS+ accounting function and determine which application TACACS+ accounts. After you enable accounting, the switch reports user activity to the TACACS+ server in the form of accounting records. The default for accounting is none.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **TACACS+**.
3. Select the **TACACS+ Globals** tab.
4. Select **GlobalEnable** to enable TACACS+ globally.
5. Select **cli** to enable the **Accounting** option.
6. Select **cli** or **web** to enable the **Authentication** option.
7. Select **CliCommandAuthorizationEnabled** to enable TACACS+ authorization.
8. Select the level in **CliCommandAuthorizationLevels**.
9. Select **Apply**.

TACACS+ Globals Field Descriptions

Use the data in the following table to use the **TACACS+ Globals** tab.

Name	Description
GlobalEnable	Enables or disables the TACACS+ feature globally.
Accounting	<p>Determines for which applications TACACS+ collects accounting information. Use TACACS+ accounting to track the services that users access and the amount of network resources that users consume. If unassigned, TACACS+ does not perform the accounting function. The default is none.</p> <p>If enabled, TACACS+ accounting logs the following events:</p> <ul style="list-style-type: none"> • User log on and log off • Log off generated because of activity timeout • Unauthorized command • Telnet session closed (not logged off)
Authentication	<p>Configures what application TACACS+ authenticates. The options include:</p> <ul style="list-style-type: none"> • cli • web <p>TACACS + authentication provides control of authentication through login and password dialog, challenge and response. By default, CLI authentication is enabled.</p>
LastUserName	Displays the last user for which the system attempted authentication.
LastAddressType	Displays the type of address to access the TACACS+ server.
LastAddress	Displays the last address to access the TACACS+ server.
CliCommandAuthorizationEnabled	<p>Enables TACACS+ authorization for a particular privilege level. Use this option to limit the use of certain commands to certain users. To use TACACS+ authorization, you must also use TACACS+ authentication.</p> <p>The switch allows the user to access the switch according to the access level. The default is disabled.</p>
CliCommandAuthorizationLevels	<p>Enables command authorization for a specific privilege level.</p> <p>The default for the command authorization level is none.</p>

Add a TACACS+ Server

Add a TACACS+ server, configure the TACACS+ server, and specify the authentication process.

If you have a secondary server configured, the AAA request goes to the backup server if the primary server is not available.

Before You Begin

You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch are available.

About This Task

The TACACS+ server and the switch must have the same:

- Encryption key
- Connection mode (single connection or per-session connection. Per-session is the same as multi-connection mode.)
- TCP port number

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **TACACS+**.
3. Select the **TACACS+ Servers** tab.
4. Select **Insert**.
5. In **AddressType**, select **ipv4**.
6. In **Address**, type the IP address of the TACACS+ server.
7. (Optional) In **PortNumber**, type the TCP port on which the client establishes a connection to the TACACS+ server.
8. (Optional) In **ConnectionType**, select either **singleConnection** or **perSessionConnection** to specify the TCP connection type between the switch and TACACS+ server.
9. (Optional) In **Timeout**, type the period of time (in seconds) the switch waits for a response from the TACACS+ server.
10. In **Key**, enter the key that the switch and the TACACS+ server share.
11. In **Priority**, select either **primary** or **backup** to determine the order the switch uses the TACACS+ servers.
12. Select **Insert**.

TACACS+ Servers Field Descriptions

Use the data in the following table to use the **TACACS+ Servers** tab.

Name	Description
AddressType	Specifies the type of IP address to use on the TACACS+ server. You must set the value to IPv4.
Address	Specifies the IP address of the TACACS+ server.
PortNumber	<p>Configures the TCP port on which the client establishes a connection to the server. The default is 49. A value of 0 indicates that the system specified default value is used.</p> <p>You must configure the same TCP port for the TACACS+ server and the switch.</p>
ConnectionType	<p>Specifies if the TCP connection between the device and the TACACS+ server is a single connection. If you specify the single connection parameter, the connection between the switch and the TACACS+ daemon remains open, which is more efficient because it allows the daemon to handle a higher number of TACACS+ operations. The single-connection session is torn down if TACACS+ is disabled due to inactivity.</p> <p>If you do not configure this parameter, the switch uses the default connection type, which is the multi-connection. With the multi-connection, the connection opens and closes each time the switch and TACACS+ daemon communicate.</p> <p>Note: You must configure the same connection mode for the TACACS+ server and the switch. To enable single-connection, the TACACS+ daemon has to support this mode as well.</p>
ConnectionStatus	Specifies if the TCP connection between the device and TACACS+ server is connected or not connected.
Timeout	Configures the maximum time, in seconds, to wait for this TACACS+ server to reply before it times out. The default value is 10 seconds.

Name	Description
Key	Configures the authentication and encryption key for all TACACS+ communications between the device and the TACACS+ server. If the key length is zero, that indicates no encryption is used. You must configure the same encryption key for the TACACS+ server and the switch.
Priority	Determines the order in which the switch uses the TACACS+ servers, where 1 is the highest priority. The priority values are primary and backup. If more than one server shares the same priority, the device uses the servers in the order they exist in the table.

Modify a TACACS+ Configuration

About This Task

Modify an existing TACACS+ configuration to customize the server.

Procedure

1. In the navigation pane, expand **Configuration > Security > Control Path**.
2. Select **TACACS+**.
3. Select **TACACS+ Servers** tab.
4. Double-click in the fields that you want to modify.
In some of the fields, the text becomes bold, which indicates that you can edit them.
In other fields, the system displays a list.
5. In the fields that you can edit, type the desired values.
6. In the fields with lists, select the desired option.
7. Select **Apply**.

TACACS+ Configuration Example

This section provides a configuration example to configure the switch to use TACACS+.




Figure 238: Switch connects to the TACACS+ server

TACACS+ Configuration on the Switch

The following section shows the steps required to configure TACACS+ on the switch.

The example displays how to:

- Configure a key to be used by the TACACS+ server and the switch. In the example, the key is configured to the word `secret`.
- Configure an IP address for the TACACS+ server. In the example the IP address for the primary server is 192.0.2.8, which is accessible by the management interface.
- Configure the TACACS+ server to authenticate CLI sessions.
- Enable TACACS+.

Switch

```
TACACS CONFIGURATION

tacacs server host 192.0.2.8 key *****
tacacs protocol enable
tacacs accounting enable cli
tacacs authorization enable
tacacs authorization level 6
```

Verify your configuration

The **show tacacs** output must show as `global enable: true` to confirm TACACS is enabled.

The output for the **show tacacs** command must display the IP addresses for the TACACS+ server. The IP addresses must be accessible to the management interface on the switch.

If you want to use the TACACS+ server to authenticate sessions in CLI, the output must display as `authentication enabled for: cli`. If you want to authenticate EDM sessions, the output must display as `authentication enabled for: web`.

Ensure the other parameters match what you have configured.

```
Global Status:
```

```
global enable : true

authentication enabled for : cli

accounting enabled for : cli

authorization : enabled

User privilege levels set for command authorization : rwa

Server:
        create :

Prio      Status  Key      Port  IP address  Timeout Single Source
SourceEnabled
Primary   Conn    *        49    192.0.2.8   10     false  0.0.0.0
false
```



Traffic Filtering

[Traffic Filtering Fundamentals](#) on page 3681

[Access Control List Configuration using CLI](#) on page 3708

[Access Control List Configuration using EDM](#) on page 3721

[Access Control Entry Configuration using CLI](#) on page 3734

[Access Control Entry Configuration using EDM](#) on page 3758

Traffic filtering can generally provide a mechanism to accurately manage and secure network flows or prioritize crucial information over other network traffic.

The following topics provide necessary concepts and procedures to configure an access control list (ACL) and access control entry (ACE) to filter traffic.

Traffic Filtering Fundamentals

Use the information in this section to help you understand filtering. This section describes a range of features that you can use with the switch to allocate network resources to apply filters.

In a large and busy network, traffic management is very important and can be complex. Traffic filtering can generally provide a mechanism to accurately manage and secure network flows or prioritize crucial information over other network traffic. Some of the primary uses of filtering are:

- Manage traffic flows.
- Implement security permissions on network traffic.
- Prioritize mission critical traffic flows.
- Redirect traffic to firewalls or other devices to efficiently manage bandwidth.

Overview of Traffic Filtering

Traffic filtering on the switch is based on ACL filter implementation. Access Control List (ACL) based filters are a means to provide predictable and flexible traffic filtering. ACL Traffic filters can be configured using the Command Line Interface (CLI) or the Enterprise Device Manager (EDM). ACL filters set a list of criteria for the network traffic to be matched against, performing a predefined set of actions. Access Control Lists and Action Control Entries provide traffic filtering services on the switch.

Traffic filtering supports IPv6 ingress and egress port/vlan security and QoS ACL/filters. IPv6 egress QoS ACL/filters are not supported.

QoS and Filters

The switch has functions you can use to provide appropriate QoS levels to traffic for each customer, application, or packet. These functions include port-based shapers, DiffServ access or core port settings, and ingress port-rate limiting or policing. The switch also provides access control list (ACL)-based filters. You do not need to use filters to provide QoS; however, filters aid in prioritizing customer traffic. Filters also provide protection by blocking unwanted traffic.

Port rate limiting or policing apply at ingress; shapers apply at egress. ACL-based filters apply at ingress and egress.

The switch supports two ingress filter groups, where each type can hold both Security and QoS actions in both Primary Bank and Secondary Bank ranges.

Filters help you provide QoS by permitting or dropping traffic based on the parameters you configure. You can use filters to mark packets for specific treatment.

Typically, filters act as firewalls or are used for Layer 3 redirection. In more advanced cases, traffic filters can identify Layer 3 and Layer 4 traffic streams. The filters cause the streams to be re-marked and classified to attain a specific QoS level at both Layer 2 (802.1p) and Layer 3 (DSCP).

Traffic filtering is a key QoS feature. The switch, by default, determines incoming packet 802.1p or DiffServ markings, and forwards traffic based on their assigned QoS levels. However, situations exist where the markings are incorrect, or the originating user application does not have 802.1p or DiffServ marking capabilities. Also, you can give a higher priority to select users (executive class). In these situations, use filters to prioritize specific traffic streams.

You can use filters to assign QoS levels to devices and applications. To help you decide whether to use a filter, key questions include:

1. Does the user or application have the ability to mark QoS information on data packets?
2. Is the traffic source trusted? Are the QoS levels configured appropriately for each data source?

Users can maliciously configure QoS levels on their devices to take advantage of higher priority levels.

3. Do you want to prioritize traffic streams?

This decision-making process is outlined in the following figure.

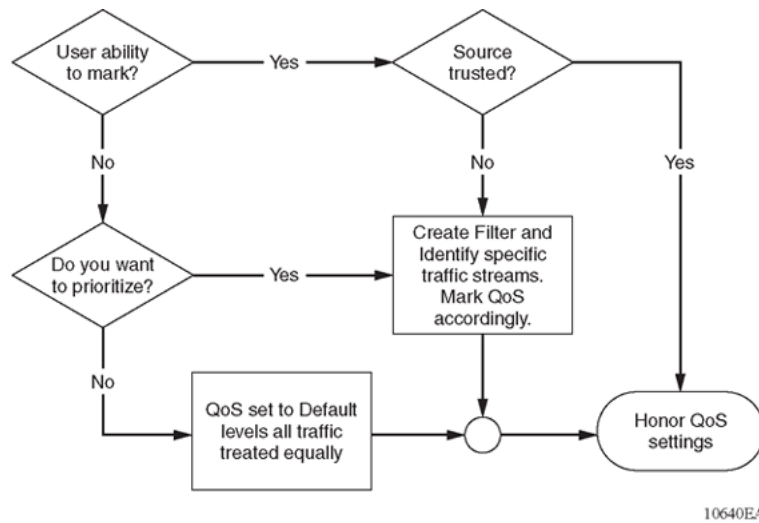


Figure 239: Filter decision-making process

Configure filters through the use of Access Control Lists (ACL) and Access Control Entries (ACE), which are implemented in hardware. An ACL can include both security and QoS type ACEs.

The following steps summarize the filter configuration process:

1. Determine your desired match fields.
2. Create an ACL.
3. Create an ACE within the ACL.
4. Configure the desired precedence, traffic type, and action.

You determine the traffic type by creating an ingress or egress ACL.

5. Modify the parameters for the ACE.

Access Control Lists

Table 309: Access Control List product support

Feature	Product	Release introduced
IPv4 Ingress and Egress port ACLs	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
		Note: Egress ACLs are not supported.
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Ingress VLAN	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Ingress VSN	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 309: Access Control List product support (continued)

Feature	Product	Release introduced
IPv6 ingress filters	4220 Series	Fabric Engine 9.2 IPv6 management only
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
IPv6 egress filters	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
IPv6 ACL DSCP Remarking	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

Apply rules to incoming and outgoing traffic. The total number of ACLs that you can configure differs depending on the switch.

An ACL can filter either IPv6 or non-IPv6 packets. By default, an ACL filters non-IPv6 packets. You must specify the packet type as IPv6 at the ACL level to enable IPv6 filtering. You cannot change the packet type for the ACL after you configure it. If you need a different packet type, you must delete the ACL, and then re-create it with the other packet type.

**Note**

All ACLs are enabled when a DvR Leaf node is created. When creating the DvR Leaf node, or when you enable IS-IS globally, the message `dvr leaf [x]` displays in CLI. If the DvR Leaf node is reset or reconnects to the DvR domain, this message does not display but the same principle applies. As a best practice, delete ACLs not used on DvR Leaf nodes, rather than disable them, because the ACLs re-enable without warning after a DvR domain transition caused by reset, or protocol or link bounce.

You can associate an ACL with the following interfaces:

- Ingress port (inPort)
- Ingress VLAN (inVLAN)

**Note**

VLAN-based ACL filters are not supported on a DvR Leaf node.

- Ingress VSN (inVSN)
- Egress port (outPort)

**Note**

Exception: does not apply to 7830 Series.

The ingress VLAN ACL associations apply to all active port members of a VLAN. An ACL is created in the enabled state by default.

The InVSN Filter is an Access Control List (ACL) that can be used with MAC-in-MAC (MIM) encapsulated packets that are received on the network-to-network interface (NNI) ingress ports and are routed or bridged to UNI ports or terminated on the fabric node. The InVSN Filter matches and filters IPv4 and IPv6 packet headers coming on UNI ports only, NNI ports only, or both UNI and NNI ports. The InVSN Filter does not filter packets that arrive on NNI ingress ports but are bridged to other NNI ports or are for transit traffic.

An ACL can contain multiple filter rules called Access Control Entries (ACE). ACEs provide match criteria and rules for ACL-based filters. An ACE can provide actions such as dropping a packet, monitoring a packet, or remarking QoS on a packet. Complete lists of actions are provided in the Access Control Entries section. After an ingress or egress packet meets the match criteria specified in ACEs within an ACL, the system executes the predefined action.

ACLs provide the ability to configure default and global actions. A default action is applied when no filter rule (ACE) matches on a packet flow. The global action runs after

any filter rule (ACE) matches on a packet flow. The default action mode for ACLs is permit. ACL global actions are:

- monitor-dst-mlt
- monitor-dst-ports

The following figure shows the relationships between ACEs and VLAN- and port-based ACLs.

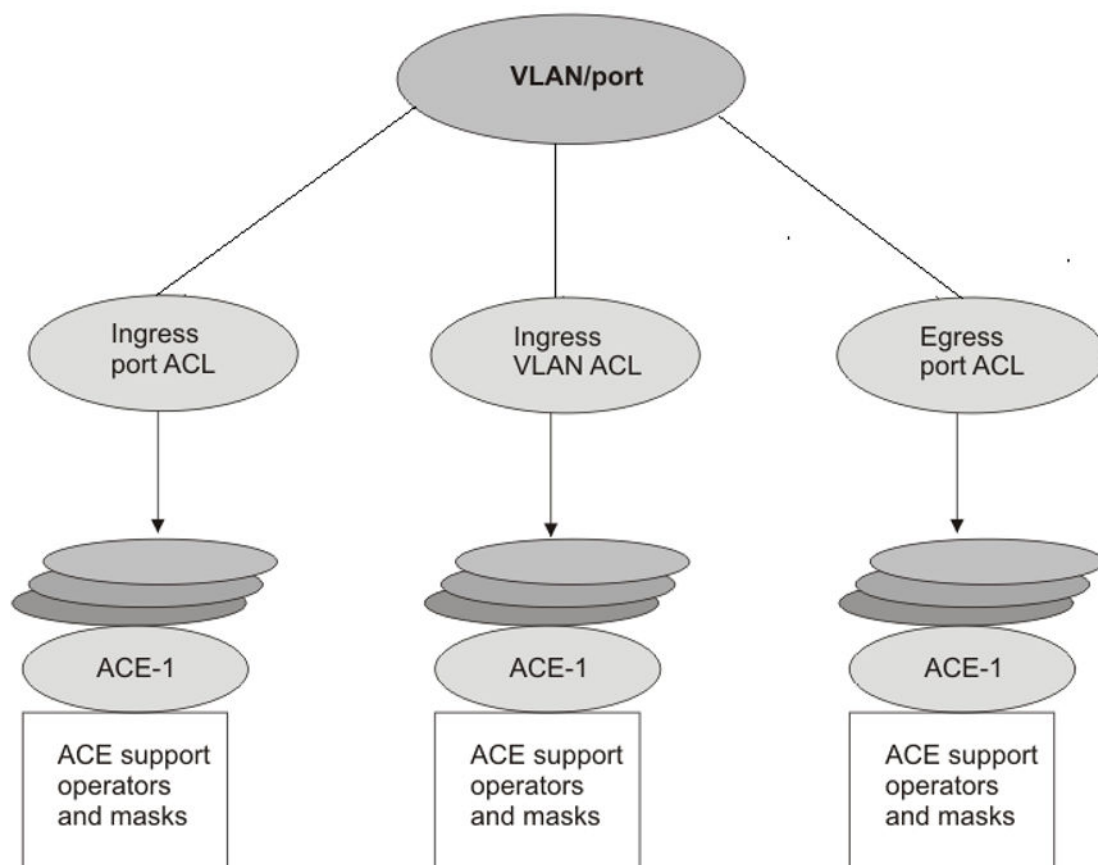


Figure 240: ACE and ACL relationships

Access Control Entries

Table 310: Access Control Entry product support

Feature	Product	Release introduced
QoS Access Control Entries (ACE)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Security ACEs	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
IPv4 ACL filter enhancement - Apply ACE with both Security and QoS actions	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	Fabric Engine 8.7
	5520 Series	Fabric Engine 8.7
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.7

Table 310: Access Control Entry product support (continued)

Feature	Product	Release introduced
Filter enhancement - Apply ACE to Routed Packets only	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.4
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.4
	VSP 7400 Series	VOSS 8.4
Policy Based Routing (redirect-next-hop)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 310: Access Control Entry product support (continued)

Feature	Product	Release introduced
Policy Based Routing (redirect-next-hop) with VRF support	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The switch filter rules are defined using Access Control Entries (ACE). An ACE is an ordered set of filter rules contained in an Access Control List (ACL). ACE rules are divided into the following three components:

- Operators
- Attributes
- Actions

An ACE generally operates on fields in a packet. If a packet field matches an ACE rule, the system executes the action specified. As each packet enters through an interface with an associated ACL, the system scans the ACE list configured on that ACL and matches on the packet fields. If multiple ACE rules are associated with the ACL, the lower ACE ID will have a higher precedence.

Operators

ACEs use operators to match on packet fields. The switch supports the following operators:

- Equal-to

This rule operator looks for an exact match with the field defined. If the field matches exactly with the rule, the system will return a match (hit). If the rule does not match, the search continues and at the end of the search a miss is returned.

- Mask

ACL-based filters provide the mask operator to match on Layer 2, Layer 3, and Layer 4 packet fields. The mask operator is used to mask bits in packet fields during a search or to match on a partial value of a packet field. This section provides examples of the mask operator.

If a mask bit is set to 1, it means it is not part of the match criteria (treated as do not care), and a mask bit of 0 means that the value represented is part of the match criteria. You can use the mask operator for the following attributes:

- source MAC address
- destination MAC address
- VLAN ID
- Dot1p
- IPv4/IPv6 source address
- IPv4/IPv6 destination address
- destination IP address
- DSCP
- Layer 4 source port
- Layer 4 destination port
- TCP flags

**Note**

MAC Address cannot be configured as attributes for IPv6 filters.

The syntax for ACL and ACE configuration of a mask is similar to the use of equal operator, except that you provide the mask value. You can specify a mask value (number) to represent the bits to mask in the attribute. You can define a mask in different ways depending on the attribute you need to mask:

- If you use a decimal number for an IP address mask, it specifies the most significant bits of the provided IP address to match on. For example, a mask of 24 used with an IP address is the same as a mask of 0.0.0.255, and a mask of 8 used with an IP address is the same as a mask of 0.255.255.255.
- If you use a decimal number for a MAC address mask, it specifies the least significant bits of the provided MAC address to ignore. For example, a mask of 32 used with a MAC address is the same as a mask of 0x0000ffffff, and a mask of 16 used with a MAC address is the same as a mask of 0x00000000ffff.

**Note**

Unlike the standard convention, for ACL filter configuration, a mask bit value of '1' specifies a do-not-care bit, and value of '0' signifies must-match bit.

The following table explains the mask operator for MAC addresses.

Table 311: Mask operator for MAC address

Rule	Result
<pre>filter acl ace ethernet 10 10 dst-mac mask 01:00:5e:00:00:01 0x000000FFFFFF</pre>	The rule matches only on the most significant 24 bits as they are not masked, for example, 01:00:5e, and does not care about the least significant 24 bits because they are masked; the least significant 24 bits can have a value of 00:00:00 - FF:FF:FF.
<pre>filter acl ace ethernet 10 10 dst-mac mask 01:00:5e:00:00:01 0xFFFFFFFF0000</pre>	The rule matches only on the least significant 16 bits because they are not masked, for example, 00:01, and does not care about the most significant 32 bits because they are masked; the most significant 32 bits can have a value of 00:00:00:00 – FF:FF:FF:FF.
<pre>filter acl ace ethernet 10 10 dst-mac mask 01:00:5e:00:00:01 0xFF00FF0000FF</pre>	The rule matches only on the unmasked bits, for example, 0xXX:00:XX:00:00:XX. The rule matches only on the bits not masked, for example, all the zeroes and the x represents a do not care (0xXX:00:XX:00:00:XX)

The following table explains the mask operator for IP addresses.

Table 312: Mask operator for IP address

Rule	Result
<pre>filter acl ace ip 10 10 src-ip mask 2.10.10.12 0.255.255.255</pre>	The rule matches only the most significant 8 bits, and does not care about the value of the remaining 24 bits as they are considered masked. For example, 10.10.12. Packets with a source IP address of 2.15.16.122 or 2.3.4.5 match on the filter rule while packets with a source IP address of 3.10.10.12 and 4.10.10.12 do not match on the filter rule.
<pre>filter acl ace ip 10 10 src-ip mask 3.4.5.6 255.255.255.0</pre>	The rule matches only the least significant 8 bits, for example, 6, and does not care about the most significant 24 bits, 3.4.5. Packets with a source IP address of 17.16.5.6 or 192.168.1.6 match on the filter rule while packets with a source IP address of 3.4.5.4 or 3.4.5.7 do not match on the filter rule.

The following table explains the mask operator for Layer 4 source port.

Table 313: Mask operator for Layer 4 source port

Rule	Result
<pre>filter acl ace protocol 10 10 src- port mask 80 0xF</pre>	The filter rule matches on Layer 4 source port 80 (1010000). The mask value 0xF (1111) masks the least significant 4 bits, which means source port 81 (1010001) through 95 (1011111) also match this filter rule. This means the range 80–95 is a match on this rule.

The following table demonstrates the resulting action based on mask configuration and example packets.

Table 314: Mask operator configuration examples

Filter configuration	Address examples that match the filter	Address examples that do not match the filter
<p>Ethernet mask:</p> <pre>filter acl 1000 type inport filter acl port 1000 1/5,1/11 filter acl ace 1000 12 filter acl ace ethernet 1000 12 src-mac mask 00:00:11:11:16:00 0x00ff000000f0 filter acl ace action 1000 12 permit count filter acl ace 1000 12 enable</pre>	<p>Source MAC:</p> <pre>00:01:11:11:16:10 00:10:11:11:16:f0 00:1f:11:11:16:10 00:ff:11:11:16:f0 00:00:11:11:16:60 00:e6:11:11:16:e0</pre>	<p>Source MAC:</p> <pre>00:00:11:11:16:01 00:ff:11:11:16:f1</pre>
<pre>filter acl ace 1000 1000 filter acl ace ethernet 1000 1000 dst-mac mask 00:00:00:64:16:00 0x00000060001f filter acl ace action 1000 1000 deny count filter acl ace 1000 1000 enable</pre>	<p>Destination MAC:</p> <pre>00:00:00:64:16:01 00:00:00:04:16:01 00:00:00:24:16:1f 00:00:00:64:16:1f 00:00:00:44:16:10 00:00:00:04:16:05</pre>	<p>Destination MAC:</p> <pre>00:00:00:24:16:20 00:00:00:64:16:20 00:00:00:63:16:01 00:00:00:65:16:01</pre>
<p>IP mask (dotted decimal notation):</p> <pre>filter acl 10 type outport filter acl port 10 1/13 filter acl ace 10 11 filter acl ace ethernet 10 11 ether- type eq ip filter acl ace ip 10 11 src-ip mask 192.168.4.0 0.0.0.31 filter acl ace action 10 11 permit count filter acl ace 10 11 enable</pre>	<p>Source IP: 192.168.4.1 192.168.4.10 192.168.4.30 192.168.4.31</p>	<p>Source IP: 192.168.3.1 192.168.4.32</p>

Table 314: Mask operator configuration examples (continued)

Filter configuration	Address examples that match the filter	Address examples that do not match the filter
<pre>filter acl ace 10 12 filter acl ace ethernet 10 12 ether- type eq ip filter acl ace ip 10 12 dst-ip mask 192.168.7.0 0.0.0.3 filter acl ace action 10 12 deny count filter acl ace 10 12 enable</pre>	Destination IP: 192.168.7.1 192.168.7.3	Destination IP: 192.168.7.4 192.168.7.5
<p>IP mask (decimal notation):</p> <pre>filter acl 10 type outport filter acl port 10 1/13 filter acl ace 10 11 filter acl ace ethernet 10 11 ether- type eq ip filter acl ace ip 10 11 src-ip mask 192.168.4.0 255.255.255.31 filter acl ace action 10 11 permit count filter acl ace 10 11 enable</pre>	Source IP: 192.168.4.1 192.168.4.10 192.168.4.30 192.168.4.31	Source IP: 192.168.3.1 192.168.4.32
<pre>filter acl ace 10 12 filter acl ace ethernet 10 12 ether- type eq ip filter acl ace ip 10 12 dst-ip mask 192.168.7.0 255.255.255.3 filter acl ace action 10 12 deny count filter acl ace 10 12 enable</pre>	Destination IP: 192.168.7.1 192.168.7.3	Destination IP: 192.168.7.4 192.168.7.5
<p>Protocol mask:</p> <pre>filter acl 901 type inport filter acl port 901 1/2 filter acl ace 901 1 filter acl ace ip 901 1 ip-protocol- type eq tcp filter acl ace protocol 901 1 src- port mask 256 0xff filter acl ace action 901 1 deny count filter acl ace 901 1 enable</pre> <p>This mask implies packets with TCP source port 256–511 match the filter, while 0–255 and > 511 miss the filter.</p>	TCP source port 256 TCP source port 356 TCP source port 511	TCP source port 255 TCP source port 512

Attributes

Attributes are fields in a packet (Layer 2, Layer 3, Layer 4) or other information related to the packet on which an ACE rule is applied like slot/port. [Table 315](#) provides a list of all the attributes and the operators that you can apply to them.

If you want to configure IPv6 attributes, you must configure an ACL to filter either IPv6 or non-IPv6 traffic. You can only configure IPv6 attributes for IPv6 packets. You cannot configure IPv6 attributes for non-IPv6 packets.

Table 315: Attribute list

Attribute Name	Operator
Slot/Port	Equal
Destination MAC (IPv4 filters only)	Equal, Mask
Source MAC (IPv4 filters only)	Equal, Mask
VLAN ID	Equal, Mask
.1p bits	Equal, Mask
Ether Type	Equal
ARP Opcode	Equal
Source IP	Equal, Mask
Destination IP	Equal, Mask
Protocol Type	Equal
Type of Service	Equal, Mask
IP Fragmentation	Equal
Note: Exception: not supported on 7830 Series.	
IP Options	Equal
Layer 4 Destination Port	Equal, Mask
Layer 4 Source Port	Equal, Mask
TCP Flags	Equal, Mask
ICMP Message Type	Equal
Source IPv6 (IPv6 only)	Equal, Mask
Destination IPv6 (IPv6 only)	Equal, Mask
Next header (IPv6 only)	Equal
Traffic class (IPv6 only)	Equal
Routed only	Equal

4220 Series Restrictions

Note the following restrictions to the attribute list:

- 4220 Series supports the following for IPv4 ACLs only.
 - .1p bits
 - TCP flags
 - ICMP Message Type

- 4220 Series models are restricted to a maximum of 15 distinct values for each source/destination port. The following list identifies the reserved entries in the 15 number set:
 - 67-68 [DHCP]
 - 546-547 [DHCPv6]
 - 53 [DNS]
 - 23, 2323 [Telnet]

5320 Series Restrictions

Note the following restrictions to the attribute list:

- 48-port 5320 Series models support .1p bits for both IPv4 and IPv6 ACLs. 16-port and 24-port 5320 Series models support .1p bits for IPv4 ACLs only.
- 16-port and 24-port 5320 Series models are restricted to a maximum of 15 distinct values for each source/destination port. The following list identifies the reserved entries in the 15 number set:
 - 67-68 [DHCP]
 - 546-547 [DHCPv6]
 - 53 [DNS]
 - 23, 2323 [Telnet]
- 48-port 5320 Series models support TCP flags for both IPv4 and IPv6 ACLs. 16-port and 24-port 5320 Series models support TCP flags for IPv4 ACLs only.
- 48-port 5320 Series models support ICMP Message Type for both IPv4 and IPv6 ACLs. 16-port and 24-port 5320 Series models support ICMP Message Type for IPv4 ACLs only.
- Only the 48-port 5320 Series models support the attributes identified as IPv6 only.

Actions

Actions occur when the filter rule is hit or missed. The types of actions that the filter configuration can execute are split into two categories:

- security actions
- QoS actions



Note

- Ingress ACLs support security and QoS ACE actions. Egress ACLs do not support QoS ACEs.

For Fabric Engine switches, except for 4220 Series and 5320 Series, the ACE ID range for the Primary Bank is 1-1000 and for the Secondary Bank, the ACE ID range is 1001-2000. The switch performs a parallel search on both ACE lists. If actions do not conflict, both actions apply. If actions conflict, the action from the range with higher priority applies.

The 4220 Series and 5320 Series switches use a modified implementation for IPv4 ACL filters. In this implementation, there is no distinction between ACE IDs used for Security

or QoS actions. You can configure Security and QoS actions in the ACE ID range 1-2000 and hence the switch does not perform a parallel search on the two ACE types.

If you apply multiple ACE rules, the lower ACE ID has a higher precedence.



Note

The ACE ID range for both security and QoS actions is different for different hardware platforms. For more information, see [ACL Filters Behavior Differences](#) on page 3706.

The following tables show the supported switch actions:

Table 316: Security ACE Actions

Security ACE Actions	User supplied parameters	Comments
count	None	Collect ACE statistics. Applies to Ingress and Egress ACLs.
mode	Permit or Deny	Applies to both Ingress and Egress ACLs.
monitor-dst-mlt	mlt-id	Applies to Ingress ACLs only. Note: This option is not supported on 7830 Series.
monitor-dst-ports	Port	Applies to Ingress ACLs only.
monitor-isid-offset	Offset-id	Determines the actual monitor I-SID value to which packets are mirrored. Note: This option is not supported on 7830 Series. Note: This option is not supported on Multi-area SPB Boundary Nodes.

Table 316: Security ACE Actions (continued)

Security ACE Actions	User supplied parameters	Comments
redirect-next-hop	IP address, Mode	Redirects the packet to the user supplied IP address. If the switch cannot resolve ARP for the user-specified next-hop, packets that match the filter are dropped. Applies to ingress IPv4 ACLs only (routed and Layer 2).
unreachable	permit or deny	Denies or permits packet dropping when the next hop for the packet is unreachable.

Table 317: QoS ACE Actions

QoS ACE Actions	User supplied parameters	Comments
count	None	Applies to Ingress and Egress ACLs.
<ul style="list-style-type: none"> internal-qos remark-dot1p remark-dscp 	<ul style="list-style-type: none"> Internal-qos dot1p (ingress only) DCSP 	<p>Applies to Ingress ACLs.</p> <p>Note: remark-dot1p and internal-qos do not apply to IPv6 filtering.</p> <p>Note: Some hardware platforms do not support remark-dot1p and support remark-DSCP for Layer 3 routed packets only.</p>

When you configure an IPv6 ACL with an ACE action of remark DSCP for a mirrored packet, the mirrored copy does not include the remark DSCP value. Because of port-mirroring functionality, the mirrored copy does not reflect the changes that occur in the switch to the outgoing packet. As a result, the mirrored copy is not identical to the outgoing packet. For more information, see [Port Mirroring](#) on page 3870.

Internal QoS Level and Remarking

Setting the internal QoS level is an ingress action. Remarking is an egress action.

The internal-qos action assigns a new value to the packet internal-qos. It determines the packet egress queue, outgoing packet dot1p value and egress-DSCP value.

The remark-dot1p action assigns a new dot1p value to the outgoing packet. The remark-DSCP action assigns a new DSCP value to the outgoing packet.

If a packet is filtered by a rule set to internal-qos action only, then the packet internal qos, egress queue, egress dot1p and egress DSCP will be derived from the filter internal-qos value.

If a packet is filtered by a rule set to remark-dot1p only or remark-DSCP only or both remark actions, then the packet will be remarked with the new dot1p or DSCP, or both. However, these remarked values will not have any impact on the internal-qos packet. It will be based on the native packet coming into the switch.

If a packet is filtered by a rule set with all three qos actions, then the internal-qos will determine the egress queue, but the remark-dot1p determines the egress dot1p and the remark-DSCP determines the egress DSCP.

If you want to change the internal QoS for remarked incoming packets, you have to add the **permit internal-qos** command as shown in the following ACL filter example.

```
filter acl 10 type inPort name "ACL-CTI"
filter acl port 10 1/2-1/50
filter acl ace 10 1302 name "CIFS-SCCM Source"
filter acl ace action 10 1302 permit remark-dscp phbaf11 remark-dot1p 1 count
filter acl ace action 10 1302 permit internal-qos 0
filter acl ace ethernet 10 1302 ether-type eq ip
filter acl ace ip 10 1302 src-ip mask 0.0.0.0 255.255.255.255
filter acl ace ip 10 1302 ip-protocol-type eq tcp
filter acl ace protocol 10 1302 src-port mask 0 0xffff
```

When a packet goes through the switch, the internal QoS level governs which queue the packet uses on egress. To verify which queue the packets are egressing on, use the **show qos cosq-stats interface [value]** command. For more information, see [View Port Egress CoS Queue Statistics](#) on page 3027 or [View Port Egress CoS Queue Statistics](#) on page 3037.

Conflict and Precedence

The switch supports both port-based and VLAN-based ACLs. A port can be associated with both Port-based ACL and a VLAN-based ACL, as shown in [Access Control Lists](#) on page 3684. Within an ACL, a rule match can generate security actions and QoS actions. The system goes through a set of precedence levels to resolve any conflicting actions between port-based ACL and VLAN-based ACL lookup.

The following table provides a list of search results and actions for all possible conflicts between port and VLAN-based ACLs and security and QoS ACE for each ACL.

Table 318: Conflict and Precedence resolution

Port-based ACL look up		Actions performed on Port-based ACL		If VLAN-based ACL is enabled		Actions performed on VLAN-based ACL search	
Security	QoS	Security action	QoS action	Security	QoS	Security action	QoS action
Security ACE search is a Miss and ACL mode is Permit.	QoS ACE search is a Miss	Default security statistics collected	Default QoS statistics collected	Security ACE search is a Miss and mode is set to Permit	QoS ACE search is a Miss	Collect default Miss statistics	Collect default Miss statistics
				Security ACE search is a Miss and mode is set to Permit	QoS ACE search returns a Hit	Collect default Miss statistics	Execute configured ACE and default ACL actions
				Security ACE search is a Miss and mode is set to Deny	Search result is invalid, since security mode is set to Deny	Drop packet and collect default Miss statistics	No action is executed
				Security ACE search is a Hit and mode is set to Permit	QoS ACE search returns a Miss	Execute configured ACE and default ACL actions	Collect default Miss statistics
				Security ACE search is a Hit and mode is set to Permit	QoS ACE search is a Hit	Execute configured ACE and default ACL actions	Execute configured ACE and default ACL actions
				Security ACE search is a Hit and mode is set to Deny	QoS ACE search returns a Hit	Discard the packet and execute configured ACE and global actions	No action is executed
Security ACE is Miss and ACL mode is Deny	Search result is invalid since security mode is set to Deny	Discard the packet and collect default statistics	No action is executed	VLAN-based ACL is not configured	VLAN-based ACL is not configured	No action is executed	No action is executed

Table 318: Conflict and Precedence resolution (continued)

Port-based ACL look up		Actions performed on Port-based ACL		If VLAN-based ACL is enabled		Actions performed on VLAN-based ACL search	
Security	QoS	Security action	QoS action	Security	QoS	Security action	QoS action
Security ACE search is a Miss and ACL mode is set to Permit	QoS ACE search is a Hit	Default search statistics collected	Execute configured ACE and default ACL actions	Security ACE search is a Miss and mode is set to Permit	Port-based ACL's QoS action take precedence. QoS search result invalid.	Collect default Miss statistics	No action is executed
				Security ACE search is a Miss and mode is set to Deny	Port-based ACL's QoS action take precedence. QoS search result invalid.	Drop packet and collect default Miss statistics	No action is executed
				Security ACE search is a Hit and mode is set to Permit	Port-based ACL's QoS action take precedence. QoS search result invalid.	Execute configured ACE and default ACL actions	No action is executed
				Security ACE search is a Hit and mode is set to Deny	Port-based ACL's QoS action take precedence. QoS search result invalid.	Discard the packet and execute configured ACE and global Actions	No action is executed
Security ACE search is a Hit and ACE mode is Permit	QoS ACE search is a Miss	Execute configured ACE and default ACL actions	Collect default Miss statistics	Port-based ACL's Security action take precedence. Security search result invalid	QoS ACE search returns a Miss	No action is executed	Collect default Miss statistics
				Port-based ACL's Security action take precedence. Security search result invalid.	QoS ACE search returns a Hit	No action is executed	Execute configured ACE and default ACL actions

Table 318: Conflict and Precedence resolution (continued)

Port-based ACL look up		Actions performed on Port-based ACL		If VLAN-based ACL is enabled		Actions performed on VLAN-based ACL search	
Security	QoS	Security action	QoS action	Security	QoS	Security action	QoS action
Security ACE search is a Hit and ACE mode is Permit	QoS ACE search is a Hit	Execute configured ACE and default ACL actions	Execute configured ACE and default ACL actions.	Port-based ACL's Security action take precedence. Security search result invalid	Port-based ACL's QoS action take precedence. QoS search result invalid.	No action is executed	No action is executed
Security ACE search is a Hit and ACE mode is Deny	Search result is invalid since Security mode is set to Deny	Discard the packet and collect default statistics	No action is executed	Port-based ACL's Security action take precedence. Security search result invalid	Port-based ACL's QoS action take precedence. QoS search result invalid.	No action is executed	No action is executed

Common ACE Uses and Configuration

The following table describes configurations you can use to perform common actions.

Table 319: Common ACE uses and configurations

Function	ACE configuration
Permit a specific host to access the network	<ul style="list-style-type: none"> Use action permit. Configure the source IP address to be the host IP address. <pre>filter acl ace 1 5 name "Permit_access_to_198.51.100.0" filter acl ace action 1 5 permit filter acl ace ethernet 1 5 ether-type eq ip filter acl ace ip 1 5 src-ip eq 198.51.100.0 filter acl ace 1 5 enable</pre>
Deny a specific host from accessing the network	<ul style="list-style-type: none"> Use action deny. Configure the source IP address to be the host IP address. <pre>filter acl ace 1 5 name "Deny_access_to_198.51.100.0" filter acl ace action 1 5 deny filter acl ace ethernet 1 5 ether-type eq ip filter acl ace ip 1 5 src-ip eq 198.51.100.0 filter acl ace 1 5 enable</pre>

Table 319: Common ACE uses and configurations (continued)

Function	ACE configuration
Permit a specific range of hosts to access the network	<ul style="list-style-type: none"> Use action permit. Configure the source IP address to be the range of host IP addresses. <pre>filter acl ace 1 5 name "Permit_access_to_1.2.3.4-1.2.3.7" filter acl ace action 1 5 permit filter acl ace ethernet 1 5 ether-type eq ip filter acl ace ip 1 5 src-ip mask 1.2.3.4 0.0.0.3 filter acl ace 1 5 enable</pre>
Deny Telnet traffic	<ul style="list-style-type: none"> Use action deny. Configure the protocol as TCP and the TCP destination port to be 23. <pre>filter acl ace 1 5 name "Deny_telnet" filter acl ace action 1 5 deny filter acl ace ethernet 1 5 ethertype eq ip filter acl ace ip 1 5 ip-protocol-type eq tcp filter acl ace protocol 1 5 dst-port eq 23 filter acl ace 1 5 enable</pre>
Deny FTP traffic	<ul style="list-style-type: none"> Use action deny. Configure the protocol as TCP and the TCP destination port to be 21. <pre>filter acl ace 1 5 name "Deny_ftp" filter acl ace action 1 5 deny filter acl ace ethernet 1 5 ethertype eq ip filter acl ace ip 1 5 ip-protocoltype eq tcp filter acl ace protocol 1 5 dst-port eq 21 filter acl ace 1 5 enable</pre>

Switched UNI ACL Filters

InPort and OutPort filters are available on Switched UNI (S-UNI) and Fabric Attach ports for traffic mapped to an I-SID without an associated platform VLAN. The Customer VLAN-ID (CVID) can be used as a VLAN-ID qualifier in both InPort and OutPort filters. For traffic mapped to an I-SID with an associated platform VLAN, InPort, OutPort, and inVLAN filters are supported on S-UNI and Fabric Attach ports. In this case, the platform VLAN should be used as the VLAN-ID in InPort and inVLAN filters, while the CVID should be used as the VLAN-ID in the OutPort filter.

Traffic Filter Configuration

Traffic filtering manages traffic by defining filtering conditions and associating these conditions with specific actions. The following steps summarize the filtering configuration process:

1. Determine your desired match fields.
2. Configure an ACL and associate it with Ingress or Egress traffic flow.
3. Configure an ACE within the ACL.
4. Configure the desired precedence, attributes, and action.
5. Enable the ACE.

Ingress Policer and Port Rate Limiter

Table 320: Port-Based Rate Limiting, Policing, and Shaping product support

Feature	Product	Release introduced
Egress port shaper	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Ingress policer and port rate limiter	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5
QoS ingress port-based rate limiter	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Not Supported
	7720 Series	Not Supported
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	Not Supported

Ingress policer and port rate limiter functionality handles congestion, meters traffic, and takes action based on the ACL filter rules. A policer can be attached directly to an ACL ACE entry to perform flow-based rate limiting. You can configure the service rate

and peak rate. All traffic that qualifies for a filter is metered. Every packet that exceeds the service rate is marked as yellow, and every packet that exceeds the peak rate is marked as red and drops. The line rate of the port is the maximum rate that you can configure.

The ingress policer and port rate limiter feature is available on the following interfaces:

- Ingress port (inPort)
- Ingress VLAN (inVLAN)
- Ingress VSN (inVSN)

Operation Considerations

The following section describes operational considerations for ingress policer and port rate limiter functionality:

- IPv4 policer statistics are available for red, yellow, and green packets and bytes.
- Egress port (outPort) interface is not supported.
- Policers are not supported on IPv6 ACLs.
- Supports QoS ACEs only.
- You can configure the default action policer within an ACL for permit action mode only.
- Drop action mode is not supported.
- **monitor-dst-ports**, **monitor-dst-mlt**, or **monitor-isid-offset** actions are supported. All traffic received is mirrored.
- The filter with policer and redirect-next-hop actions redirect only the traffic that passes the policer (if the stream is 100 Mbps, and the policer peak-rate is 50 Mbps, only 50 Mbps is redirected).

ACL and ACE Configuration Guidelines

To find the maximum number of ACLs and ACEs that the switch supports, see the *Fabric Engine Release Notes*.

ACL Filters Behavior Differences

The implementation of ACL filters is similar in all switches but there are some differences as summarized in the following tables.



Note

The InVSN Filter shares the port-based groups in the following table.

Table 321: Action behavior

Filter	4220 Series 5320 Series	5420 Series 5520 Series 5720 Series 7520 Series 7720 Series 7830 Series
ACE ID ranges supported	<ul style="list-style-type: none"> IPv4 and IPv6 filters: <ul style="list-style-type: none"> ACEs: 1–2000 support both Security and QoS actions 	<ul style="list-style-type: none"> IPv4 filters support both Security and QoS actions in both Primary Bank and Secondary Bank ranges: <ul style="list-style-type: none"> Primary Bank: 1-1000 Secondary Bank: 1001-2000 IPv6 filters: <ul style="list-style-type: none"> ACEs: 1–2000 support both Security and QoS actions
redirect-next-hop support	<p>Supported in both the Global Routing Table and VRF contexts.</p> <p>The 4220 Series switches support only a single active VRF with IP configuration.</p> <p>For the 5320 Series, only the following models support more than one VRF with IP configuration:</p> <ul style="list-style-type: none"> 5320-16P-2MXT-2X 5320-24T-4X-XT 5320-24T-24S-4XE-XT 	<p>Supported in both the Global Routing Table and VRF contexts.</p>

Table 321: Action behavior (continued)

Filter	4220 Series 5320 Series	5420 Series 5520 Series 5720 Series 7520 Series 7720 Series 7830 Series
	<ul style="list-style-type: none"> • 5320-48P-8XE • 5320-48T-8XE 	

Table 322: ACL statistics behavior

4220 Series 5320 Series	5420 Series 5520 Series 5720 Series 7520 Series 7720 Series 7830 Series
Does not support viewing ACL statistics by the ACE type, Security and QoS. The output displays N/A.	Supports viewing ACL statistics by the ACE type, Primary Bank and Secondary Bank.

Table 323: ACE match criteria

4220 Series 5320 Series	5420 Series 5520 Series 5720 Series 7520 Series 7720 Series 7830 Series
<p>The 16-port and 24-port models support the following ACE match criteria for IPv6 ACLs:</p> <ul style="list-style-type: none"> • ethernet ACE: <ul style="list-style-type: none"> ◦ ether-type ◦ port ◦ vlan-id • IPv6 ACE: <ul style="list-style-type: none"> ◦ dst-ipv6 ◦ nxt-hdr ◦ routed-only ◦ src-ipv6 	<p>Supports the following ACE match criteria for IPv6 ACLs:</p> <ul style="list-style-type: none"> • ethernet ACE: <ul style="list-style-type: none"> ◦ ether-type ◦ port ◦ vlan-id ◦ vlan-tag-prio • IPv6 ACE: <ul style="list-style-type: none"> ◦ dst-ipv6 ◦ nxt-hdr ◦ routed-only ◦ src-ipv6 ◦ traffic-class

Table 323: ACE match criteria (continued)

4220 Series 5320 Series	5420 Series 5520 Series 5720 Series 7520 Series 7720 Series 7830 Series
<ul style="list-style-type: none"> protocol ACE: <ul style="list-style-type: none"> dst-port src-port <p>Note: 4220 Series and 16-port and 24-port 5320 Series models are restricted to a maximum of 15 distinct values for each source/destination port. For more information, see Attributes on page 3694.</p> <p>Support on the 48-port models is the same as 5420 Series and 5520 Series.</p>	<p>Note: 7830 Series does not support IPv6 ACE.</p> <ul style="list-style-type: none"> protocol ACE: <ul style="list-style-type: none"> dst-port icmpv6-msg-type src-port tcp-flags

For QoS scaling and filter scaling information, see *Fabric Engine Release Notes*.

Access Control List Configuration using CLI

Use an access control list (ACL) to specify an ordered list of access control entries (ACEs), or filter rules. The ACEs provide specific actions that you want the filter to perform.

Create an IPv4 ACL

Create an ACL to specify an ordered list of ACEs, or filter rules.

About This Task

Do not configure IPv4 egress ACL filters on network-to-network interface (NNI) ports because the system-generated egress vIST filter rules and the user-created IPv4 egress rules use the same filter hardware.

Procedure

- Enter Global Configuration mode:

```
enable

configure terminal
```
- Create an ACL:

```
filter acl <1-2048> type <inVlan|inPort|outPort|inVsn> [matchType
<both|terminatingNNIOnly|uniOnly> ] [name WORD<0-32>] [enable]
```

3. Enable an ACL:

```
filter acl [enable]
```

4. Ensure the configuration is correct:

```
show filter acl [<1-2048>]
```

Variable Definitions

The following table defines parameters for the **filter acl** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<i>enable</i>	Enables the ACL state, and all associated ACEs. Enabled is the default state.
<i>matchType</i> < <i>both</i> <i>terminatingNNIOnly</i> <i>uniOnly</i> >	For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are: <ul style="list-style-type: none"> <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node <i>uniOnly</i> for traffic ingressing on UNI ports only The default value is <i>both</i> .
<i>name</i> WORD<0-32>	Specifies an optional descriptive name for the ACL.
<i>type</i> < <i>inVlan</i> <i>inPort</i> <i>outPort</i> <i>inVsn</i> >	Specifies the ACL type. The values inVlan, inPort, and inVsn are ingress ACLs, and outPort is an egress ACL. A port-based ACL has precedence over a VLAN-based ACL.

Create an IPv6 ACL

Create an IPv6 ACL to specify an ordered list of ACEs, or filter rules.

You must specify the packet type as IPv6 at the ACL level to enable IPv6 filtering. By default, an ACL filters non IPv6 packets.

**Note**

You cannot change packet type for the ACL after you have configured it. If you want a different packet type, you must delete the ACL and re-create it using the other packet type.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create an IPv6 ACL:

```
filter acl <1-2048> type <inVlan|inPort|outPort|inVsn> [matchType
<both|terminatingNNIOnly|uniOnly> ] [name WORD<0-32>] [pktType ipv6]
[name <0-32>]
```

**Note**

IPv6 egress QoS ACL/Filters are not supported.

3. Enable the ACL:

```
filter acl <1-2048> enable
```

4. Ensure the configuration is correct:

```
show filter acl [<1-2048>]
```

Variable Definitions

The following table defines parameters for the **filter acl** command.

Variable	Value
<code><1-2048></code>	Specifies the ACL ID.
<code>enable</code>	Enables the ACL state, and all associated ACEs. Enabled is the default state.
<code>matchType <both terminatingNNIOnly uniOnly></code>	For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are: <ul style="list-style-type: none"> <code>both</code> for traffic ingressing on both UNI ports and NNI ports terminating on this node <code>terminatingNNIOnly</code> for traffic ingressing on NNI ports only and terminating on this node <code>uniOnly</code> for traffic ingressing on UNI ports only The default value is <code>both</code> .
<code>name WORD<0-32></code>	Specifies an optional descriptive name for the ACL.
<code>type <inVlan inPort outPort inVsn></code>	Specifies the ACL type. The values <code>inVlan</code> , <code>inPort</code> , and <code>inVsn</code> are ingress ACLs. The value <code>outPort</code> configures IPv6 egress filters. A port-based ACL has precedence over a VLAN-based ACL.
<code>pktType <ipv6></code>	Specifies the IP version as IPv6. The default is <code>nonipv6</code> . <p>Note: You cannot change packet type for the ACL after you have configured it. If you want a different packet type, you must delete the ACL and re-create it using the other packet type.</p>

Associate VLANs with an ACL

Associate VLANs with an ACL to apply filters to VLAN traffic.

A VLAN can be part of two different ACLs of different types: IPv6 and non-IPv6.

Before You Begin

- The ACL exists.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Add VLAN interfaces to an ACL:

```
filter acl vlan <1-2048> <1-4059>
```
3. Remove specified VLAN interfaces from an ACL:

```
no filter acl vlan <1-2048> <1-4059>
```

Variable Definitions

The following table defines parameters for the **filter acl vlan** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Associate Ports with an ACL

Associate ports with an ACL to apply filters to port traffic.

A port can be part of two different ACLs of different types: IPv6 and non-IPv6.

Before You Begin

- The ACL exists.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Associate port interfaces with a particular ACL:

```
filter acl port <1-2048> {slot/port[-slot/port]}[,...][slot/all][all]}
```
3. Remove port interfaces from a particular ACL:

```
no filter acl port <1-2048> {slot/port[-slot/port]}[,...][slot/all][all]}
```

Variable Definitions

The following table defines parameters for the **filter acl port** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Associate an I-SID with an ACL

Before You Begin

- The inVsn ACL exists.
- This I-SID is already configured on the fabric node.

About This Task

For inVsn ACL types, specify the I-SID associated with the customer VLAN (Layer 2 VSN), the customer VRF (Layer 3 VSN), or the IP Shortcut.



Note

For IP Shortcut traffic, the inVsn ACL match type must be *both*. In this case, the I-SID is zero (0).

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Specify the I-SID.


```
filter acl i-sid <1-2048> <0-15999999>
```

Variable definitions

Use the data in the following table to use the **filter acl i-sid** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<0-15999999>	<p>Specifies the I-SID associated with the customer VLAN (Layer 2 VSN) or the customer VRF (Layer 3 VSN). This I-SID must already be configured on the fabric node.</p> <p>The InVSN Filter supports IP Shortcut traffic if the inVsn ACL match type is <i>both</i>. In this case, the I-SID is zero (0).</p> <p>Important: You can specify a Switched UNI I-SID if the I-SID is associated with a platform VLAN.</p>

Configure Global and Default Actions for an ACL

Configure the default action to specify packet treatment if a packet does not match an ACE.

Configure the global action to specify packet treatment if a packet does match an ACE. You can only configure the global action for ingress ACLs.

Before You Begin

Ensure that the ACL exists.

Procedure

- Enter Global Configuration mode:


```
enable
configure terminal
```
- Configure the global action for an ACL:


```
filter acl set <1-2048> global-action [monitor-dst-ports {slot/port[-slot/port][,...][slot/all][all]] [monitor-dst-mlt <1-512>]
```
- Configure the default action for an ACL using one of the following commands:
 - Configure the default action to permit:

```
filter acl set <1-2048> default-action permit [policer svc-rate <0-4000000000> peak-rate <8-4000000000>]
```
 - Configure the default action to deny:

```
filter acl set <1-2048> default-action deny [control-packet-action <deny | permit>]
```

Variable Definitions

The following table defines parameters for the **filter acl set** commands.

Variable	Value
<1-2048>	Specifies the ACL ID.
<i>control-packet-action</i> <deny permit>	Specifies the action to apply on control packets when none of the ACEs match. The default is permit. To use this optional parameter, you must configure the <i>default-action</i> to deny .
<i>default-action</i> <permit deny>	Specifies the default action to take when none of the ACEs match. Options are <permit deny>. The default is permit.
<i>monitor-dst-mlt</i> <1-512>	Configures mirroring to a destination MLT in the range of 1 to 512.
<i>monitor-dst-ports</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}	Specifies the global action to take for matching ACEs: <ul style="list-style-type: none"> monitor destination ports—Configures mirroring to a destination port or ports. Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> a single slot and port (slot/port) a range of slots and ports (slot/port-slot/port) a series of slots and ports (slot/port,slot/port,slot/port) all ports on the same slot (slot/all) all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<i>policer svc-rate</i> <0-4000000000> <i>peak-rate</i> <8-4000000000>	Specifies the policer for filter with service rate and peak transfer rate of packets. The service rate value specifies the rate of traffic committed to be delivered. Packets above the specified peak rate value are dropped on ingress. To use this optional parameter, you must configure the <i>default-action</i> to permit .

Rename an ACL

Perform this procedure to change the name of an existing ACL.

Procedure

- Enter Global Configuration mode:

```
enable

configure terminal
```
- Rename an ACL:

```
filter acl <1-2048> name WORD<0-32>
```


3. Reset the ACL name to the default name:

```
default filter acl <1-2048> name
```

Variable Definitions

The following table defines parameters for the **filter acl** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<i>enable</i>	Enables the ACL state, and all associated ACEs. Enabled is the default state.
<i>matchType</i> < <i>both</i> <i>terminatingNNIOnly</i> <i>uniOnly</i> >	For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are: <ul style="list-style-type: none"> <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node <i>uniOnly</i> for traffic ingressing on UNI ports only The default value is <i>both</i> .
<i>name</i> WORD<0-32>	Specifies an optional descriptive name for the ACL.
<i>type</i> < <i>inVlan</i> <i>inPort</i> <i>outPort</i> <i>inVsn</i> >	Specifies the ACL type. The values inVlan, inPort, and inVsn are ingress ACLs, and outPort is an egress ACL. A port-based ACL has precedence over a VLAN-based ACL.

Disable an ACL

Perform this procedure to disable an ACL and all ACEs that belong to it.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable an ACL:

```
no filter acl <1-2048> enable
```

Variable Definitions

The following table defines parameters for the **filter acl** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<i>enable</i>	Enables the ACL state, and all associated ACEs. Enabled is the default state.

Variable	Value
<i>matchType</i> < <i>both</i> <i>terminatingNNIOnly</i> <i>uniOnly</i> >	For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are: <ul style="list-style-type: none"> • <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node • <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node • <i>uniOnly</i> for traffic ingressing on UNI ports only The default value is <i>both</i> .
<i>name</i> WORD<0–32>	Specifies an optional descriptive name for the ACL.
<i>type</i> < <i>inVlan</i> <i>inPort</i> <i>outPort</i> <i>inVsn</i> >	Specifies the ACL type. The values inVlan, inPort, and inVsn are ingress ACLs, and outPort is an egress ACL. A port-based ACL has precedence over a VLAN-based ACL.

Reset an ACL to Default Values

Reset an ACL to change the ACL name to the default name and the filter ACL mode to a default of enable.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Reset an ACL to default values:

```
default filter acl <1-2048>
```

Variable Definitions

The following table defines parameters for the **filter acl** command.

Variable	Value
<1–2048>	Specifies the ACL ID.
<i>enable</i>	Enables the ACL state, and all associated ACEs. Enabled is the default state.

Variable	Value
<i>matchType</i> < <i>both</i> <i>terminatingNNIOnly</i> <i>uniOnly</i> >	For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are: <ul style="list-style-type: none"> <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node <i>uniOnly</i> for traffic ingressing on UNI ports only The default value is <i>both</i> .
<i>name</i> WORD<0-32>	Specifies an optional descriptive name for the ACL.
<i>type</i> < <i>inVlan</i> <i>inPort</i> <i>outPort</i> <i>inVsn</i> >	Specifies the ACL type. The values inVlan, inPort, and inVsn are ingress ACLs, and outPort is an egress ACL. A port-based ACL has precedence over a VLAN-based ACL.

Delete an ACL

Delete an ACL to remove an ordered list of filter rules.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Delete an ACL:

```
no filter acl <1-2048>
```

The system displays the following message:

```
WARNING: All ACE entries under this ACL will be Deleted.
Do you wish to delete this ACL? (y/n)?
```

3. Enter *y*.

Variable Definitions

The following table defines parameters for the **filter acl** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<i>enable</i>	Enables the ACL state, and all associated ACEs. Enabled is the default state.

Variable	Value
<i>matchType</i> < <i>both</i> <i>terminatingNNIOnly</i> <i>uniOnly</i> >	For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are: <ul style="list-style-type: none"> <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node <i>uniOnly</i> for traffic ingressing on UNI ports only The default value is <i>both</i> .
<i>name</i> WORD<0-32>	Specifies an optional descriptive name for the ACL.
<i>type</i> < <i>inVlan</i> <i>inPort</i> <i>outPort</i> <i>inVsn</i> >	Specifies the ACL type. The values inVlan, inPort, and inVsn are ingress ACLs, and outPort is an egress ACL. A port-based ACL has precedence over a VLAN-based ACL.

Clear ACL Statistics

Clear default ACL statistics if you no longer require previous statistics.



Note

The ACL statistics do not support security action on some hardware platforms.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Enter the following command to clear default ACL statistics:
clear filter acl statistics default <1-2048>
3. Enter the following command to clear global ACL statistics:
clear filter acl statistics global <1-2048>
4. Enter the following command to clear all ACL statistics:
clear filter acl statistics all
5. Enter the following command to clear statistics associated with a particular ACL, ACE, or ACE type:
clear filter acl statistics <1-2048> <1-2000> [qos] [security]

Variable Definitions

The following table defines parameters for the **clear filter acl statistics** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<1-2000>	Specifies the ACE ID.

Enable IPv6 Egress Filters

Use the **boot config flags** command to enable IPv6 egress filters and to add IPv6 egress qualifiers at startup.



Note

This procedure does not apply to 4220 Series.

Before You Begin

If more than 200 IPv4 egress entries exist in the configuration file, make a backup of the configuration file before you enable IPv6 egress filters. Only a maximum of 200 IPv4 egress entries are saved in the configuration file after you use the **save config** command.

For example, you can enter more than 200 IPv4 egress entries in the configuration file prior to enabling IPv6 egress filters. However, the entries are stored in ascending numerical order with ACL ID and ACE ID respectively, and not in the order in which they were added. Therefore, after you enable IPv6 egress filters and restart, and because the configuration file is read in ascending order, you receive an error message after the 200 maximum has been reached, such as:

```
CP1 [2017-09-28T00:44:24.077+05:30] 7K-Fi-94-I6:1 0x001049d4 00000000
GlobalRouter FILTER ERROR Unable to allocate data path resources for ACL
ID 12.
```

About This Task

This flag is disabled by default.



Note

Product Notice: For 5320 Series and 5420 Series platforms, the **boot config flags ipv6-egress-filter** and **boot config flags macsec** commands are mutually exclusive.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Enable the IPv6-egress-filter boot config flag:
`boot config flags ipv6-egress-filter`
3. Save the configuration, and then restart the switch for the change to take effect.
4. After you restart the switch, verify that the IPv6-egress-filter boot config flag is configured to true:
`show boot config flags`

Example

```
Switch:1>enable
Switch:1#configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags ipv6-egress-filter
Warning: Please save the configuration and reboot the switch
        for this configuration to take effect.
```



Note

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

Display Filter Access Control List Configuration

About This Task

Perform this procedure to display the filter Access Control List (ACL) configuration on the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display filter ACL configuration:
show filter acl

Example

```
Switch:1(config)#show filter acl
```

```
=====
                        Vlan/VSN ACL Table
=====
Acl  Type   AclName          PktType   State   Origin # of   Default   CtrPkt   Service-rate Peak-rate   Vlan/I-sid
Id                                     ACEs      Action    Rule                                     Id
=====
```

Vlan ACL Global-Action Table											
Acl Id	Type	Ipfix	Monitor Dst-Mlt	Monitor Dst-Port							
Port ACL Table											
Acl Id	Type	AclName	PktType	State	Origin	# of	Default	CtrPkt	Service-rate ACEs	Peak-rate Action	Port
Rule											
1	Ingress	ACL-1	nonipv6	enabled	config	2	permit	permit	100	200	1/1
3	Ingress	ACL-3	nonipv6	enabled	config	1	permit	permit	800	1000	1/7-1/8
Port ACL Global-Action Table											
Acl Id	Type	Ipfix	Monitor Dst-Mlt	Monitor Dst-Port							
1	Ingress	Disable	0								
3	Ingress	Disable	0								
Displayed 2 of 2 Entries											

Access Control List Configuration using EDM

Use traffic filtering to provide security by blocking unwanted traffic and prioritizing other traffic.

Configure an Access Control List

About This Task


Use an access control list (ACL) to specify an ordered list of access control entries (ACE), or filter rules. The ACEs provide specific actions for the filter to perform.

Do not configure IPv4 egress ACL filters on network-to-network interface (NNI) ports because the system-generated egress vIST filter rules and the user-created IPv4 egress rules use the same filter hardware.

You cannot change a parameter that the system displays it dimmed; in this case, delete the ACL, and then configure a new one.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select **Insert**.

5. In **AcId**, type an ACL ID, or accept the default value .
 6. In **Type**, specify the type of ACL.
 7. In **Name**, specify a name for the ACL.
 8. Perform one of the following if the ACL is VLAN-based or port-based:
 - a. If the ACL is VLAN-based, select the **VlanList** ellipsis, and then choose a VLAN list.
 - b. If the ACL is port-based, select the **PortList** ellipsis, and then choose a port list.
 9. Select the desired ports, and then select **Ok**.
 10. Configure the **DefaultAction**.
 11. Configure the **ControlPktAction**.
-  **Note**
There is no control packet action support for the InVSN Filter. Control packets go to the CPU after termination.
12. Enable or disable the **State**, as required.
 13. In **PktType**, select the packet type to create either IPv4 or IPv6 ACLs.
 14. If the ACL type is inVsn, do the following:
 - a. In **MatchType**, select the match type to associate with the ACL that the traffic is ingressing on.
 - b. In **Isid**, type the I-SID associated with the customer VLAN (Layer 2 VSN) or the customer VRF (Layer 3 VSN) or enter 0 for IP shortcut.
 15. Configure the remaining fields, as appropriate.
 16. Select **Insert**.
 17. (Optional) To delete an ACL, select the ACL, and then select **Delete**.

ACL Field Descriptions

Use the data in the following table to use the **ACL** tab.

Name	Description
AcId	Specifies a unique identifier for the ACL.
Type	<p>Specifies the ACL type. Valid options are</p> <ul style="list-style-type: none"> • inVlan • inPort • outPort • inVsn <p>Important: The inVlan ACLs drop packets if you add a VLAN after ACE creation.</p> <p>Important: You can insert an inVsn ACL type for a Switched UNI only if the Switched UNI I-SID is associated with a platform VLAN.</p>
Name	Specifies a descriptive user-defined name for the ACL.

Name	Description
VlanList	For inVlan ACL types, specifies all VLANs to associate with the ACL.
PortList	For inPort and outPort ACL types, specifies the ports to associate with the ACL.
DefaultAction	Specifies the action taken when no ACEs in the ACL match. Valid options are deny and permit, with permit as the default. Deny means the system drops the packets; permit means the system forwards packets.
ControlPktAction	Specifies the action taken for control packets. Valid options are deny and permit.
State	Enables or disables all of the ACEs in the ACL. The default value is enable.
PktType	Indicates the packet type to which this ACL applies.
MirrorMltId	Configures mirroring to a destination MLT.
MirrorDstPortList	Configures mirroring to a destination port or ports.
MatchType	<p>For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are:</p> <ul style="list-style-type: none"> • <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node • <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node • <i>uniOnly</i> for traffic ingressing on UNI ports only <p>The default value is <i>both</i></p>
Isid	<p>For inVsn ACL types, specifies the I-SID associated with the customer VLAN (Layer 2 VSN) or the customer VRF (Layer 3 VSN). This I-SID should already be configured on the fabric node.</p> <p>The InVSN Filter supports IP Shortcut traffic if the inVsn ACL match type is <i>both</i>. In this case, the I-SID is zero (0).</p> <p>Important: You can specify a Switched UNI I-SID if the I-SID is associated with a platform VLAN.</p>
Origin	<p>Indicates the origin of the ACL:</p> <ul style="list-style-type: none"> • config - ACL created by the user. • eap - ACL created by Extensible Authentication Protocol (EAP) through Remote Authentication Dial-In User Service (RADIUS) response.
DefaultSvcRate	Specifies the service rate limit in kbps {8-4000000000}.The granularity is 8 kbps.
DefaultPeakRate	Specifies the value when exceeded causes packets to drop on ingress. Peak rate limit in kbps {8-4000000000}.The granularity is 8 kbps.

View ACL Statistics

About This Task

Graph statistics for a specific ACL ID to view default statistics.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select an ACL.
5. Select **Graph**.
6. (Optional) Select **Clear Counters** to clear the data.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.



Note

Based on your hardware platform, the output can display the ACL packets by ACE type or Primary Bank or Secondary Bank.

Name	Description
AcId	Specifies the ACL ID.
MatchDefaultSecurityPkts	Shows a security packet count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultSecurityOctets	Shows a security byte count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultQosPkts	Shows a QoS packet count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultQosOctets	Shows a QoS byte count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalSecurityPkts	Shows a security packet count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalSecurityOctets	Shows a security byte count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalQosPkts	Shows a QoS packet count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalQosOctets	Shows a QoS byte count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.

Name	Description
MatchDefaultAllOctets	Shows byte count hits on Default Action if count is configured.
MatchDefaultAllPackets	Shows packet count hits on Default Action if count is configured.
MatchGlobalAllOctets	Shows byte count matching any of the ACE rules if count is configured.
MatchGlobalAllPackets	Shows packet count matching any of the ACE rules if count is configured.
MatchDefaultPrimaryBankPkts	Shows a Primary Bank packet count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultPrimaryBankOctets	Shows a Primary Bank byte count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultSecondaryBankPkts	Shows a Secondary Bank packet count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchDefaultSecondaryBankOctets	Shows a Secondary Bank byte count of traffic that does not match an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalPrimaryBankPkts	Shows a Primary Bank packet count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalPrimaryBankOctets	Shows a Primary Bank byte count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalSecondaryBankPkts	Shows a Secondary Bank packet count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
MatchGlobalSecondaryBankOctets	Shows a Secondary Bank byte count of traffic that matches an ACE rule or hits the default action if count is configured in the ACL global action.
DefaultGreenBytes	Shows byte count of traffic sent with a rate lower than committed-rate matching any of the ACE rules or hits on Default Action.
DefaultGreenPackets	Shows packet count of traffic sent with a rate lower than committed-rate matching any of the ACE rules or hits on Default Action.
DefaultYellowBytes	Shows byte count of traffic sent with a rate lower than peak-rate and higher than committed-rate matching any of the ACE rules or hits on Default Action.

Name	Description
DefaultYellowPackets	Shows packet count of traffic sent with a rate lower than peak-rate and higher than committed-rate matching any of the ACE rules or hits on Default Action.
DefaultRedBytes	Shows byte count of traffic sent with a rate higher than peak-rate matching any of the ACE rules or hits on Default Action.
DefaultRedPackets	Shows packet count of traffic sent with a rate higher than peak-rate matching any of the ACE rules or hits on Default Action.

Clear ACL Statistics

About This Task

Clear ACL statistics when you want to gather a new set of statistics.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select a field.
5. Select **ClearStats**.

Enable IPv6 Egress Filters

Enable IPv6 egress filters to add IPv6 egress qualifiers at startup.



Note

This procedure does not apply to 4220 Series

Before You Begin

If more than 200 IPv4 egress entries exist in the configuration file, make a backup of the configuration file before you enable IPv6 egress filters. Only a maximum of 200 IPv4 egress entries are saved in the configuration file after you save the configuration.

For example, you can enter more than 200 IPv4 egress entries in the configuration file prior to enabling IPv6 egress filters. However, the entries are stored in ascending numerical order with ACL ID and ACE ID respectively, and not in the order in which they were added. Therefore, after you enable IPv6 egress filters and restart, and because the configuration file is read in ascending order, you receive an error message after the 200 maximum has been reached, such as:

```
CP1 [2017-09-28T00:44:24.077+05:30] 7K-Fi-94-I6:1 0x001049d4 00000000
GlobalRouter FILTER ERROR Unable to allocate data path resources for ACL
ID 12.
```

About This Task

This flag is disabled by default.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **Boot Config** tab.
4. Select **EnableIpv6EgressFilterMode**.
5. Select **Apply**.
6. Save the configuration, and then restart the switch for the change to take effect.

Boot Config *Field Descriptions*

Use the data in the following table to use the **Boot Config** tab.

Name	Description
SwVersion	Specifies the software version that currently runs on the switch.
LastRuntimeConfigSource	Specifies the last source for the run-time image.
PrimaryConfigSource	Specifies the primary configuration source.
PrimaryBackupConfigSource	Specifies the backup configuration source to use if the primary does not exist.

Name	Description
EnableFactoryDefaultsMode	<p>Specifies whether the switch uses the factory default settings after you reboot the switch. The options are:</p> <ul style="list-style-type: none"> • false: The node does not use factory default settings. • fabric: This mode is not supported. • noFabric: The node uses the factory default mode settings. <p>Removes the configuration files, primary and secondary config file names, all user account passwords, digital certificates, IKE/OSPF/IS-IS keys, and SNMP communities. All ports are disabled and assigned to VLAN 1. License files are not removed.</p> <ul style="list-style-type: none"> • config-only <p>Boots the switch with a blank configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities. All ports are disabled and assigned to VLAN 1.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p> <ul style="list-style-type: none"> • reset-all-files: <p>Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset. If required, you must revoke the license file. The 30-day factory license does not apply to 4220 Series.</p> <p>Important:</p> <p>Only CLI provides two options to achieve the same behavior; there is no option other than reset-all-files in EDM.</p> <ul style="list-style-type: none"> • zero-touch:

Name	Description
	<p>Boots the switch with a default configuration that enables Auto-sense. This parameter resets secure files but keeps the security mode and performs a ZTP+ configuration after reboot. License files are not removed.</p> <ul style="list-style-type: none"> • zero-touch-config-only <p>Boots the switch with a blank configuration that enables Auto-sense and performs a ZTP+ configuration.</p> <p>This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, license files, and SNMP communities.</p> <p>Note: Feature support can vary across platforms. For more information, see <i>Fabric Engine and VOSS Feature Matrix</i>.</p> <p>Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.</p> <p>The default value is false. This flag is automatically reset to the default value after the switch restarts.</p> <p>If you select an option but then want to select a different option and have not yet rebooted the switch, you must disable the first option before you select the second option.</p>
EnableDebugMode	<p>Enabling the debugmode allows a user to enable TRACE on any port by prompting the selection on the console during boot up. This allows the user start trace for debugging earlier on specified port. It only works on console connection. The default value is disabled.</p> <p>Important: Do not change this parameter.</p>

Name	Description
EnableRebootOnError	Activates or disables automatic reboot on a fatal error. The default value is activated. Important: Do not change this parameter.
EnableTelnetServer	Activates or disables the Telnet server service. The default value is disabled.
EnableFtpServer	Activates or disables the FTP server on the switch. The default value is disabled. To enable FTP, ensure that the TFTPD flag is disabled.
EnableTftpServer	Activates or disables Trivial File Transfer Protocol server service. The default value is disabled.
EnableSshServer	Activates or disables the SSH server service. The default value is disabled.
EnableSpbmConfigMode	Enables you to configure SPB and IS-IS, but you cannot configure PIM and IGMP either globally or on an interface. The boot flag is enabled by default.
EnableIpv6Mode Note: Exception: does not apply to 4220 Series	Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits. This flag is disabled by default.
EnableEnhancedsecureMode	Enables or disables the enhanced secure mode. Select either jitc or non-jitc to enable the enhanced secure mode in one of these sub-modes. The default is disabled. Note: As a best practice, enable the enhanced secure mode in the non-JITC sub-mode because the JITC sub-mode is more restrictive and prevents the use of some troubleshooting utilities.
EnableUrpMode Note: Exception: does not apply to 4220 Series	Enables Unicast Reverse Path Forwarding (uRPF) globally. You must enable uRPF globally before you configure it on a port or VLAN. The default is disabled.

Name	Description
EnableVxlanGwFullInterworkingMode Note: Exception: only supported on 7520 Series and 7720 Series.	Enables VXLAN Gateway in Full Interworking Mode, which supports SPB, SMLT, and vIST. By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes by enabling this boot configuration flag. In Base Interworking Mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.
EnableFlowControlMode	Enables or disables flow control globally. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames. The default is disabled.

Name	Description
<p>AdvancedFeatureBwReservation</p> <p>Note: Exception: high is supported on 4220 Series, 5720 Series, 7520 Series, and 7720 Series . Exception: low is not supported on 5720 Series. Exception: vim is only supported on 5520 Series and 5720 Series.</p>	<p>Enables the switch to support advanced features by reserving ports as loopback ports. When disabled, you can use all ports on the switch, but advanced features do not work.</p> <p>When you select low or high, the corresponding bandwidth of the reserved Ethernet ports is allocated to the loopback required for advanced feature functionality.</p> <p>When you select vim, the bandwidth is allocated from the unused vim slot capacity; you cannot select this parameter if the VIM is present.</p> <p>The default varies depending on the platform:</p> <ul style="list-style-type: none"> • The default for 4220 Series is enabled with high level. • The default for 5320 Series, 5420 Series, 7520 Series, and 7720 Series is enabled with low level. • The default for 5520 Series is enabled with vim level if Versatile Interface Module (VIM) is not installed, else low level is enabled. • The default for 5720 Series is enabled with vim level if Versatile Interface Module (VIM) is not installed, else high level is enabled. <p>If you change this parameter, you must restart the switch.</p>
<p>EnableDvrLeafMode</p> <p>Note: Exception: does not apply to 4220 Series</p>	<p>Enables the switch to be configured as a DvR Leaf.</p> <p>When enabled, you cannot configure the switch to operate as a DvR Controller.</p>
<p>EnablevrfScaling</p> <p>Note: Exception: does not apply to 4220 Series</p>	<p>Changes the maximum number of VRFs and Layer 3 VSNs that the switch supports. If you select this check box, the maximum number increases. The default is disabled.</p> <p>Important: If you select both this check box and the EnableSpbmConfigMode check box, the switch reduces the number of configurable VLANs. For more information about maximum scaling numbers, see <i>Fabric Engine Release Notes</i>.</p>

Name	Description
EnableSyslogRfc5424Format	Enables or disables the RFC 5424 syslog format. The default is enabled. If the pre-existing configuration file is for a release prior to this enhancement, then the flag is disabled automatically.
NniMstp Note: Exception: does not apply to 4220 Series	Enables MSTP, and allows non SPBM B-VLAN configuration on SPBM network-to-network interface (NNI) ports. The default is disabled. Note: Spanning Tree is disabled on all SPBM NNIs. You cannot add an SPBM NNI port or MLT port to any non SPBM B-VLAN.
EnableIpv6EgressFilterMode Note: Exception: does not apply to 4220 Series	Enables IPv6 egress filters. The default is disabled. If you change this parameter, you must restart the switch. For 5320 Series, 5420 Series, and 5720 Series platforms, EnableIpv6EgressFilterMode and EnableMacsec are mutually exclusive.
EnableMacsec Note: Exception: only required for MACsec supporting models of 5320 Series and 5420 Series.	Enables Media Access Control Security (MACsec) mode globally. To enable MACsec mode, you must configure the boot flag. EnableIpv6EgressFilterMode and EnableMacsec are mutually exclusive.
EnableSpbmNodeScaling Note: Exception: only applies to 5320 Series (not 5320-16P-2MXT-2X or 5320-24T-4X-XT) and 5420 Series.	Enables the switch to increase the number of supported SPB nodes per area. By default, the switch supports up to 350 SPB nodes per area. The default is disabled. If you change this parameter, you must restart the switch.
Slot	Specifies the slot number.
TftpHash	Enables TFTP hashing.
TftpRetransmit	Set TFTP retransmit timeout counter.
TftpTimeout	Set TFTP timeout counter.
User	Configure host user.
Password	Configure host password.

Access Control Entry Configuration using CLI

Use an access control entry (ACE) to provide an ordered list of traffic filtering rules. You can configure up to 1000 ACEs in a single ACL.



Note

For information about the supported ACE IDs ranges on all hardware platforms, see [ACL Filters Behavior Differences](#) on page 3706.

Configure ACEs

Use an ACE to define packet attributes and the desired behavior for packets that carry the attribute or list of attributes.

Before You Begin

- The ACL exists. If you want to use IPv6 filters, you must specify the packet type as IPv6 at the ACL level to enable IPv6 filtering.

About This Task

ACLs are by default created in enabled state while ACEs are by default created in disabled state. Use CLI commands to enable an ACE.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create and name an ACE:

```
filter acl ace <1-2048> <1-2000> [name WORD<0-32>]
```

The ACE ID determines ACE precedence (that is, the lower the ID, the higher the precedence).

3. Configure the mode as deny or permit:

```
filter acl ace action <1-2048> <1-2000> <deny|permit>
```

4. Configure ACE actions as required.

5. Ensure the configuration is correct:

```
show filter acl ace <1-2048> <1-2000>
```

6. Ensure the filter is enabled:

```
filter acl ace <1-2048> <1-2000> enable
```

7. Optionally, reset an ACE to default values (reset the ACE name to the default name and the administrative state to the default value of disable):

```
default filter acl ace <1-2048> <1-2000>
```

8. Optionally, delete an ACE ID:

```
no filter acl ace <1-2048> <1-2000>
```

Variable Definitions

Use the data in the following table to use the **filter acl ace** and the **filter acl ace action** commands.

Variable	Value
<code><1-2048></code>	Specifies the ACL ID.
<code><1-2000></code>	Specifies the ACE ID.
<code><deny permit></code>	Configures the action mode for security ACEs. Note: For each Security ACE, you must define one or more actions as well as the associated action mode (permit or deny). Otherwise, the security ACE cannot be enabled. There is no default configuration for Security ACEs. With QoS ACE, the action mode is not configurable. QoS ACEs are always set to action mode permit.
<code>enable</code>	Enables an ACE within an ACL. After you enable an ACE, to make changes, first disable it.
<code>name WORD<0-32></code>	Specifies an optional descriptive name for the ACE that uses 0-32 characters.

Configure ACE Actions

Configure ACE actions to determine the process that occurs after a packet matches an ACE.

Before You Begin

- Create ACE and ACL.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure ACE actions:

```
filter acl ace action <1-2048> <1-2000> <permit | deny>
```
3. (Optional) Configure ACE actions to count matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> count
```

4. (Optional) Configure the QoS level for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> internal-qos
<0-7>
```



Note

This step does not apply to IPv6 filtering.

5. (Optional) Enable mirroring on destination MLT for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> monitor-dst-
mlt <1-512>
```

6. (Optional) Enable mirroring on a port for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> monitor-
dst-ports {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all]
[all]}
```

7. (Optional) Enable mirroring on destination I-SID for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> monitor-isid-
offset <1-1000>
```

8. (Optional) Configure the next hop IPv4 or IPv6 address for redirect mode for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> redirect-next-
hop WORD<1-45> [count | unreachable | vrf {WORD <1-16>}]
```



Important

Ensure you configure the ACE match rules so that you only collect the desired traffic. For example, routed packets.

9. (Optional) Configure the next hop IPv4 or IPv6 address for redirect mode for matching packets for a VRF:



Note

If the next hop is unreachable, you can also configure ACE actions to permit or deny packet dropping within the VRF.

```
filter acl ace action <1-2048> <1-2000> <permit | deny> redirect-next-
hop WORD<1-45> vrf WORD <1-16> unreachable <permit | deny>
```

10. (Optional) Configure the next hop IPv4 or IPv6 address for redirect mode for matching packets for a VRF:



Note

If the next hop is unreachable, you can also configure ACE actions to count matching packets, or to permit or deny packet dropping within the VRF

```
filter acl ace action <1-2048> <1-2000> <permit | deny> redirect-next-
hop WORD<1-45> vrf WORD <1-16> unreachable <permit | deny> count
```

11. (Optional) Configure the QoS dot1p priority for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> remark-dot1p
<0-7>
```

**Note**

This step does not apply to IPv6 filtering.

12. (Optional) Configure the QoS phb and dscp for matching packets:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> remark-dscp
[phbcs0 | phbcs1 | phbaf11 | phbaf12 | phbaf13 | phbcs2 | phbaf21 |
phbaf22 | phbaf23 | phbcs3 | phbaf31 | phbaf32 | phbaf33 | phbcs4 |
phbaf41 | phbaf42 | phbaf43 | phbcs5 | phbef | phbcs6 | phbcs7]
```

13. (Optional) Configure the mode when next hop is unreachable:

```
filter acl ace action <1-2048> <1-2000> <permit | deny> unreachable
[permit | deny]
```

14. Ensure the configuration is correct:

- `show filter acl action <1-2048> <1-2000>`
- `show filter acl config`
- `show filter acl ace`

Example

Configure ACE actions:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#filter acl ace action 1 47 permit redirect-next-hop 192.0.2.5
unreachable deny count
```

Display the configuration using the **show filter ace action** command:

```
Switch:1(config)#show filter acl action
```

=====								
Ace Action Table (Part I)								
=====								
Acl Id	Ace Id	AceName	Admin State	Oper State	Mode	Mlt Id	Remark DSCP	Remark Dot1p

1	47	ace47	Disable	Down	permit	0	disable	disable
=====								
Ace Action Table (Part II)								
=====								
Acl Id	Ace Id	Redirect Next-Hop	Vrf name	Unreach -able			Police	Internal Qos

1	47	2.0.0.0	GlobalRouter	deny			0	0
=====								
Ace Action Table (Part III)								
=====								
Acl Id	Ace Id	Ipfix	Count	Log	CopyTo Pcap	Monitor Dst-Mlt	Monitor Dst-Vlan	Monitor Dst-Port

```

1    47    disable enable    disable disable 1        0

=====
                        Ace Action Table (Part IV)
=====

Acl  Ace  Monitor          Dscp    Ttl      Monitor   Isid   QoS Remove-Tag
Id   Id   Dst-IP                  Isid      Offset
-----
1    47    0.0.0.0          ----    ----    ---      ---    ---

Displayed 1 of 1 Entries

```

Display the configuration using the **show filter acl config** command:

```

Switch:1(config)#show filter acl config

=====
                        Filter ACL-ACE Configuration
=====

filter acl 1 type inPort name "ACL-1"
filter acl set 1 policer svc-rate 100 peak-rate 200
filter acl port 1 1/1
filter acl ace 1 1
filter acl ace action 1 1 permit count
filter acl ace ethernet 1 1 src-mac eq aa:bb:cc:dd:ee:ff
filter acl ace policer 1 1 svc-rate 300 peak-rate 400
filter acl ace 1 1 enable
filter acl ace 1 2
filter acl ace action 1 2 permit count
filter acl ace ethernet 1 2 dst-mac eq ff:ff:ff:ff:ff:ff
filter acl ace policer 1 2 svc-rate 500 peak-rate 600
filter acl ace 1 2 enable
filter acl 3 type inPort name "ACL-3"
filter acl set 3 policer svc-rate 800 peak-rate 1000

filter acl port 3 1/7-1/8

filter acl ace 3 2

filter acl ace action 3 2 permit count

filter acl ace ethernet 3 2 dst-mac eq 00:00:00:00:00:33

filter acl ace policer 3 2 svc-rate 1000 peak-rate 4000

filter acl ace 3 2 enable

```

Display the configuration using the **show filter acl ace** command:

```

Switch:1(config)#show filter acl ace

=====
                        Ace Action Table (Part I)
=====

Acl  Ace  AceName          Admin  Oper   Mode   Mlt Remark  Remark
Id   Id                  State  State  Mode   Id  DSCP   Dot1p
-----
1    1    ACE-1              Enable Up      permit 0   disable disable
1    2    ACE-2              Enable Up      permit 0   disable disable

```


3	2	ACE-2	Enable	Up	permit 0	disable	disable
4	4	ACE-4	Disable	Down	permit 0	disable	disable
6	10	ACE-10	Disable	Down	permit 0	disable	disable
15	15	ACL15	Enable	Up	permit 0	phbaf23	disable

=====

Ace Action Table (Part II)

=====

Acl Id	Ace Id	Redirect Next-Hop	Vrf name	Unreach-able	Police	Internal Qos
1	1	0.0.0.0	GlobalRouter	deny	0	0
1	2	0.0.0.0	GlobalRouter	deny	0	0
3	2	0.0.0.0	GlobalRouter	deny	0	0
4	4	0.0.0.0	GlobalRouter	deny	0	0
6	10	0:0:0:0:0:0:0:0	GlobalRouter	deny	0	0
15	15	0:0:0:0:0:0:0:0	GlobalRouter	deny	0	0

=====

Ace Action Table (Part III)

=====

Acl Id	Ace Id	Ipfix	Count	Log	CopyTo Pcap	Monitor Dst-Mlt	Monitor Dst-Vlan	Monitor Dst-Port
1	1	disable	enable	disable	0	0		
1	2	disable	enable	disable	0	0		
3	2	disable	enable	disable	0	0		
4	4	disable	disable	disable	0	0		
6	10	disable	disable	disable	0	0		
15	15	disable	enable	disable	0	0		

=====

Ace Action Table (Part IV)

=====

Acl Id	Ace Id	Monitor Dst-IP	Dscp	Ttl	Monitor Isid	Isid Offset	QoS	Remove-Tag
1	1	0.0.0.0	----	----	---	---	---	
1	2	0.0.0.0	----	----	---	---	---	
3	2	0.0.0.0	----	----	---	---	---	
4	4	0.0.0.0	----	----	---	---	---	
6	10	0.0.0.0	----	----	---	---	---	
15	15	0.0.0.0	----	----	---	---	---	

Displayed 3 of 3 Entries

=====

ACE Arp Table

=====

AclId	AceId	Operation
1	1	
1	2	
3	2	
4	4	
6	10	
15	15	

Displayed 3 of 3 entries

```
=====
```

ACE Ethernet Table (Part I)				
=====				
Acl Id	Ace Id	Operator/SourceMac	Operator/DestMac	Operator/PortList

1	1	eq aa:bb:cc:dd:ee:ff		
1	2	eq ff:ff:ff:ff:ff:ff		
3	2	eq 00:00:00:00:00:33		
4	4			
6	10			
15	15			

```
=====
```

ACE Ethernet Table (Part II)				
=====				
Acl Id	Ace Id	Operator/EtherType	Operator/VlanId	Operator/VlanTagPrio

1	1			
1	2			
3	2			
4	4			
6	10			
15	15		eq 10	

Displayed 3 of 3 entries

```
=====
```

ACE Ip Table (Part I)					
=====					
Acl Id	Ace Id	Operator/SourceIp	SourceIp mask	Operator/DestIp	DestIp mask

1	1				
1	2				
3	2				
4	4				
6	10				
15	15				

```
=====
```

ACE Ip Table (Part II)					
=====					
Acl Id	Ace Id	Ip Option	Operator/IpFragFlag	Operator/IpProtoType	Operator/Dscp

1	1				
1	2				
3	2				
4	4				
6	10				
15	15				

Displayed 3 of 3 entries

```
=====
```

ACE Ipv6 Table (Part I)					
=====					

Acl Id	Ace Id	Operator/ SrcIpv6	SrcIpv6 mask
1	1		
1	2		
3	2		
4	4		
6	10		
15	15		

=====

ACE Ipv6 Table (Part II)

=====

Acl Id	Ace Id	Operator/ DstIpv6	DstIpv6 mask
1	1		
1	2		
3	2		
4	4		
6	10		
15	15		

=====

ACE Ipv6 Table (Part III)

=====

Acl Id	Ace Id	Operator/ NxtHdr	Operator/ Traffic-Cls
1	1		
1	2		
3	2		
4	4		
6	10		
15	15		

Displayed 3 of 3 entries

=====

ACE Policer Table

=====

Acl Id	Ace Id	Service-rate	Peak-rate
1	1	300	400
1	2	500	600
3	2	1000	4000
4	4		
6	10		
15	15		

Displayed 3 of 3 entries

=====

ACE Protocol Table (Part I)

=====

Acl Id	Ace Id	Operator/ SrcPort	Operator/ DstPort
1	1		

```

1      2
3      2
4      4      eq aa:bb:cc:dd:ee:ff
6      10     eq ff:ff:ff:ff:ff:ff
15     15     eq 00:00:00:00:00:33

```

=====

ACE Protocol Table (Part II)

=====

Acl Id	Ace Id	Operator/ TcpFlags	Operator/ IcmpMsgType
1	1		
1	2		
3	2		
4	4		
6	10		
15	15		

=====

ACE Protocol Table (Part III)

=====

Acl Id	Ace Id	Operator/ Routing-Type
1	1	
1	2	
3	2	
4	4	
6	10	
15	15	

Displayed 3 of 3 entries

Variable Definitions

Use the data in the following table to use the **filter acl ace action** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<1-2000>	Specifies the ACE ID.
<i>count</i>	Enables the ability to count matching packets. Use this parameter with either a security or QoS ACE. The default is disabled.
<deny permit>	Configures the action mode for security ACEs. Note: For each Security ACE, you must define one or more actions as well as the associated action mode (permit or deny). Otherwise, the security ACE cannot be enabled. There is no default configuration for Security ACEs. With QoS ACEs, the action mode is not configurable. QoS ACEs are always set to action mode permit.

Variable	Value
<code>monitor-isid-offset <1-1000></code>	Specifies the offset ID which will be mapped to the actual monitor I-SID where packets are mirrored. Monitor I-SID = base monitor I-SID + offset ID. The base monitor I-SID is 16776000.
<code>internal-qos <0-7></code>	This variable is a QoS action. The default value is 1. Note: This does not apply to IPv6 filtering.
<code>monitor-dst-ports {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	Configures mirroring to a destination port or ports. This action is a security action. Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code>monitor-dst-mlt <1-512></code>	Configures mirroring to a destination MLT in the range of 1 to 512.
<code>redirect-next-hop WORD <1-45></code>	Specifies the nexthop IPv4 or IPv6 address for redirect node. This action is a security action.
<code>unreachable <permit deny></code>	Denies or permits packet dropping when the next hop for the packet is unreachable. The default value is deny. This action is a security action.
<code>vrf WORD <1-16></code>	Specifies the direct next hop VRF name. The name must be in the ranger of 1 to 16 characters.

Variable	Value
<code>remark-dscp <phbcs0 phbcs1 phbaf11 phbaf12 phbaf13 phbcs2 phbaf21 phbaf22 phbaf23 phbcs3 phbaf31 phbaf32 phbaf33 phbcs4 phbaf41 phbaf42 phbaf43 phbcs5 phbcs6 phbef phbcs7></code>	<p>Specifies the new Per-Hop Behavior (PHB) for matching packets: phbcs0, phbcs1, phbaf11, phbaf12, phbaf13, phbcs2, phbaf21, phbaf22, phbaf23, phbcs3, phbaf31, phbaf32, phbaf33, phbcs4, phbaf41, phbaf42, phbaf43, phbcs5, phbef, phbcs6, phbcs7.</p> <p>This action is a QoS action.</p> <p>Note: This action applies to IPv6 filtering.</p>
<code>remark-dot1p <0-7></code>	<p>Specifies the new 802.1 priority bit for matching packets: zero, one, two, three, four, five, six, or seven.</p> <p>This action is a QoS action.</p> <p>Note: This does not apply to IPv6 filtering.</p>

Configure ARP ACEs

Use ACE Address Resolution Protocol (ARP) entries to ensure the filter looks for ARP requests or responses.

You cannot configure ARP attributes for IPv6 filters.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Configure an ACE for ARP packets:


```
filter acl ace arp <1-2048> <1-2000> operation eq <arprequest|
arpresponse>
```
3. Ensure the configuration is correct:


```
show filter acl arp <1-2048> <1-2000>
```
4. Optionally, delete the individual attributes from the ARP portion of the ACE:


```
no filter acl ace arp <1-2048> <1-2000> [operation]
```
5. Optionally, delete all the attributes from the ARP portion of the ACE:


```
default filter acl ace arp <1-2048> <1-2000>
```

Variable Definitions

Use the data in the following table to use the **filter acl ace arp** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<1-2000>	Specifies the ACE ID.
<i>operation</i> eq <arprequest arpresponse>	Specifies the type of ARP operation to filter: arpRequest or arpResponse.

Configure an Ethernet ACE

Configure an Ethernet ACE to filter on Ethernet parameters.

You do not need to configure Ethertype for IPv6 filters. If you try to configure an Ethertype other than 0x86dd or IPv6 the device displays an error.

Before You Begin

- The ACL exists.
- The ACE exists.

About This Task

The *eq* and *mask* parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure an ACE for the destination or source MAC address attribute:

```
filter acl ace ethernet <1-2048> <1-2000> <dst-mac|src-mac> eq WORD<1-1024>
```

OR

```
filter acl ace ethernet <1-2048> <1-2000> <dst-mac|src-mac> mask  
WORD<1-1024> WORD<1-1024>
```



Note

This is supported only for IPv4 filters.

3. Configure an ACE for an Ethernet type attribute:

```
filter acl ace ethernet <1-2048> <1-2000> ether-type eq WORD<1-200>
```
4. Configure an ACE for a port attribute:

```
filter acl ace ethernet <1-2048> <1-2000> port eq {slot/port[sub-port]}
```

5. Configure an ACE for a VLAN attribute:

```
filter acl ace ethernet <1-2048> <1-2000> vlan-id eq <1-4059>
```

OR

```
filter acl ace ethernet <1-2048> <1-2000> vlan-id mask <1-4059>  
<0-0xFFF>
```

6. Configure an ACE for a VLAN tagged priority attribute:

```
filter acl ace ethernet <1-2048> <1-2000> vlan-tag-prio eq <0-7>
```

OR

```
filter acl ace ethernet <1-2048> <1-2000> vlan-tag-prio mask <0-7> <0-  
0x7>
```

7. Ensure the configuration is correct:

```
show filter acl ethernet <1-2048> <1-2000>
```

8. (Optional) Delete the individual attributes from the Ethernet portion of the ACE:

```
no filter acl ace ethernet <1-2048> <1-2000>
```

9. (Optional) Delete all the attributes from the Ethernet portion of the ACE:

```
default filter acl ace ethernet <1-2048> <1-2000>
```

Variable Definitions

The following table defines parameters for the **filter acl ace ethernet** command.

Variable	Value
<0-7>	Specifies the priority bits (3-bit field) from the 802.1Q/p tag.
<0-0x7>	Specifies the mask value for VLAN tagged priority attribute.
<0-0xFFF>	Specifies the mask value for a VLAN attribute. For example: filter acl ace ethernet 10 10 vlan-id eq 10 filter acl ace ethernet 10 10 vlan-id mask 1025 0xF
<1-2048>	Specifies the ACL ID.
<1-2000>	Specifies the ACE ID.
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Variable	Value
<code>{slot/port[/sub-port]}</code>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
<code>WORD<1-200></code>	<p>Specifies an ether-type name or number:</p> <ul style="list-style-type: none"> • 0x0–0xffff • ip, arp, ipx802dot3, ipx802dot2, ipxSnap, ipxEthernet2, appleTalk, decLat, decOther, sna802dot2, snaEthernet2, netBios, xns, vines, ipv6, rarp, or PPPoE <p>Note: Ethernet ACE filter configured with ether-type eq ipx802dot3 does not match the packet with format destination MAC address, source MAC address, length, 0xFFFF, payload and FCS. Ethernet ACE filter configured with ether-type eq ipx802dot2 does not match the packet with format destination MAC address, source MAC address, length, 0xE0E0, payload and FCS.</p>
<code>WORD<1-1024></code>	<p>If the operator is mask, the <code>WORD<1-1024></code> parameter is {"" 1..48 , mac address mask 0x0..FFFFFFFFFFFFFF}}</p> <p>If the operator is eq, the <code>WORD<1-1024></code> parameter is the destination or source MAC address: AA:BB:CC:DD:EE:FF</p> <p>For example:</p> <pre>filter acl ace ethernet 10 10 dst-mac eq 01:00:5:00:00:01 filter acl ace ethernet 10 10 dst-mac mask 01:00:5:00:00:01 24 filter acl ace ethernet 10 10 src-mac mask 01:00:5:00:00:01 0xFFFFFFFFF0000</pre>

Configure an IP ACE

Configure an IP ACE to filter on the source IP address, destination IP address, DiffServ Code Point (DSCP), protocol, IP options, IP fragmentation, and routed packets only.

Before You Begin

- The ACL exists.
- The ACE exists.

About This Task

The `eq` and `mask` parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure an ACE for the DSCP attribute:

```
filter acl ace ip <1-2048> <1-2000> dscp
eq {<0..63>|<0x00..0x3f>|phbcs0|phbcs1|phbaf11|phbaf12|phbaf13|phbcs2|
phbaf21|phbaf22|phbaf23|phbcs3|phbaf31|phbaf32|phbaf33|phbcs4|phbaf41|
phbaf42|phbaf43|phbcs5|phbef|phbcs6|phbcs7}
```

OR

```
filter acl ace ip <1-2048> <1-2000>
dscp mask {<0..63>|<0x00..0x3f>|phbcs0|phbcs1|phbaf11|
phbaf12|phbaf13|phbcs2|phbaf21|phbaf22|phbaf23|phbcs3|phbaf31|phbaf32|
phbaf33|phbcs4|phbaf41|phbaf42|phbaf43|phbcs5|phbef|phbcs6|phbcs7}
WORD<0x0-0x40>
```

3. Configure an ACE for the destination or source IP address attribute:

```
filter acl ace ip <1-2048> <1-2000> <dst-ip|src-ip> eq WORD<1-1024>
```

OR

```
filter acl ace ip <1-2048> <1-2000> <dst-ip|src-ip> mask WORD<1-1024>
{<0-32>|null|<A.B.C.D>}
```

4. Configure an ACE for the IP fragmentation attribute:

```
filter acl ace ip <1-2048> <1-2000> ip-frag-flag eq <noFragment|
anyFragment>
```

5. Configure an ACE for the IP options attribute:

```
filter acl ace ip <1-2048> <1-2000> ip-options any
```

6. Configure an ACE for the protocol type attribute:

```
filter acl ace ip <1-2048> <1-2000> ip-protocol-type eq WORD<1-256>
```

7. Configure an ACE for routed packets only:

```
filter acl ace ip <1-2048> <1-2000> routed-only
```

8. Ensure the configuration is correct:

```
show filter acl ip <1-2048> <1-2000>
```

9. (Optional) Delete all the attributes from the IP (Layer 3) portion of the ACE:

```
default filter acl ace ip <1-2048> <1-2000>
```

Example

```
Switch:1(config)#filter acl ace ip 1 12 dst-ip eq 198.51.100.0
```

Variable definitions

Use the data in the following table to use the **filter acl ace ip** command.

Variable	Value
<1-2000>	Specifies the ACE ID.
<1-2048>	Specifies the ACL ID.
{<0-32> <i>null</i> <A.B.C.D>}	Specifies the mask value for the destination or source IP address For example: <pre>filter acl ace ip 10 10 dst-ip mask 198.51.100.0 25 filter acl ace ip 10 10 dst-ip mask 198.51.100.1 203.0.113.0 filter acl ace ip 10 10 src-ip mask 198.51.100.2 22 filter acl ace ip 10 10 src-ip mask 198.51.100.3 203.0.113.1</pre>
< <i>noFragment</i> <i>anyFragment</i> >	Specifies a match option for IP fragments <i>noFragment</i> or <i>anyFragment</i> .

Variable	Value
<code>{<0..63> <0x00..0x3f> phbcs0 phbcs1 phbaf11 phbaf12 phbaf13 phbcs2 phbaf21 phbaf22 phbaf23 phbcs3 phbaf31 phbaf32 phbaf33 phbcs4 phbaf41 phbaf42 phbaf43 phbcs5 phbcs6 phbef phbcs7}</code>	<p>Specifies the DSCP value using one of the following formats:</p> <ul style="list-style-type: none"> Enter as an integer (0–63) or hex (0x00–0x3f), or as a string: <ul style="list-style-type: none"> phbcs0 — Enter as string “phbcs0”, integer 0 or hex 0x00 phbcs1 — Enter as string “phbcs1”, integer 8 or hex 0x08 phbaf11 — Enter as string “phbaf11” integer 10 or hex 0x0a phbaf12 — Enter as string “phbaf12”, integer 12 or hex 0x0c phbaf13 — Enter as string “phbaf13”, integer 14 or hex 0x0e phbcs2 — Enter as string “phbcs2”, integer 16 or hex 0x10 phbaf21 — Enter as string “phbaf21”, integer 18 or hex 0x12 phbaf22 — Enter as string “phbaf22”, integer 20 or hex 0x14 phbaf23 — Enter as string “phbaf23”, integer 22 or hex 0x16 phbcs3 — Enter as string “phbcs3”, integer 24 or hex 0x18 phbaf31 — Enter as string “phbaf31”, integer 26 or hex 0x1a phbaf32 — Enter as string “phbaf32”, integer 28 or hex 0x1c phbaf33 — Enter as string “phbaf33”, integer 30 or hex 0x1e phbcs4 — Enter as string “phbcs4”, integer 32 or hex 0x20 phbaf41 — Enter as string “phbaf41”, integer 34 or hex 0x22 phbaf42 — Enter as string “phbaf42”, integer 36 or hex 0x24 phbaf43 — Enter as string “phbaf43”, integer 38 or hex 0x26 phbcs5 — Enter as string “phbcs5”, integer 40 or hex 0x28 phbef — Enter as string “phbef”, integer 46 or hex 0x2e phbcs6 — Enter as string “phbcs6”, integer 48 or hex 0x30 phbcs7 — Enter as string “phbcs7”, integer 56 or hex 0x38
<code>WORD<0x0-0x40></code>	<p>Specifies the mask value, for example, <code>filter acl ace ip 10 10 dscp mask 129 0x40</code></p>

Variable	Value
<i>WORD</i> <1-256>	Specifies one or more IP protocol types: (1-256), or tcp, udp, ipsecesp, vrrp, snmp or undefined.
<i>WORD</i> <1-1024>	Specifies the destination or source IP address (a.b.c.d).

Configure an IPv6 ACE

Configure an IPv6 ACE to filter traffic based on Source IPv6 address, Destination IPv6 address, IPv6 next header, and IPv6 traffic class, and routed packets only.

Source IPv6 and destination IPv6 support equal (eq) and mask operators. Next header and traffic class attributes support the equal (eq) operator. The equal to rule operator looks for an exact match with the field defined. If the field matches exactly with the rule, the system will return a match (hit). ACL-based filters provide the mask operator to match on Layer 2, Layer 3, and Layer 4 packet fields. The mask operator is used to mask bits in packet fields during a search or to match on a partial value of a packet field.

Before You Begin

- The ACL exists with the IPv6 packet type. You can only configure ACE IPv6 attributes to filter on an IPv6 packet.
- The ACE exists.

About This Task

The *eq* and *mask* parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Create and name an ACE:

```
filter acl ace <1-2048> <1-2000> [name Word<1-32>]
```
3. Configure an ACE for the destination IPv6 address attribute:

```
filter acl ace ipv6 <1-2048> <1-2000> dst-ipv6 eq WORD<0-255>
```

OR

```
filter acl ace ipv6 <1-2048> <1-2000> dst-ipv6 mask WORD<1-128> WORD<0-255>
```
4. Configure an ACE for the source IP address attribute:

```
filter acl ace ipv6 <1-2048> <1-2000> src-ipv6 eq WORD<0-255>
```

OR

```
filter acl ace ipv6 src-ipv6 <1-2048> <1-2000> mask WORD<1-128> WORD<0-255>
```

- Specify the next header of the IP header:

```
filter acl ace ipv6 <1-2048> <1-2000> nxt-hdr eq {fragment|hop-by-hop|
icmpv6|ipsecah|ipsecesp|noHdr|routing|tcp|udp|undefined}
```

You must configure next header to configure the protocol attributes.

- Specify the traffic class attribute of the IPv6 header:

```
filter acl ace ipv6 <1-2048> <1-2000> traffic-class eq WORD<0-255>
```

- Configure an ACE for routed packets only:

```
filter acl ace ipv6 <1-2048> <1-2000> routed-only
```

- Ensure that your configuration is correct:

```
show filter acl ipv6 <1-2048> <1-2000>
```

Example

```
Switch:1(config)#filter acl ace ipv6 15 15 dst-ipv6 eq 30:0:0:0:0:0:ffff/64
```

Configure a Protocol ACE

Configure a protocol ACE to filter on the source port, destination port, ICMP and ICMPv6 message type, or TCP flags.



Note

For IPv6 filters, you must configure next header to configure the protocol attributes.

Before You Begin

- The ACL exists.
- The ACE exists.

About This Task

The *eq* and *mask* parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

- Configure an ACE for destination port attributes:

```
filter acl ace protocol <1-2048> <1-2000> dst-port eq WORD<1-60>
```

OR

```
filter acl ace protocol <1-2048> <1-2000> dst-port mask WORD<1-60>
WORD<1-256>
```

3. Configure an ACE for source port attributes:

```
filter acl ace protocol <1-2048> <1-2000> src-port eq WORD<1-65535>
```

OR

```
filter acl ace protocol <1-2048> <1-2000> src-port mask WORD<1-65535>
WORD<1-256>
```

4. Configure an ACE for ICMP message type attributes:

```
filter acl ace protocol <1-2048> <1-2000> icmp-msg-type eq WORD<1-200>
```

5. Configure an ACE for TCP flags attributes:

```
filter acl ace protocol <1-2048> <1-2000> tcp-flags eq WORD<1-50>
```

OR

```
filter acl ace protocol <1-2048> <1-2000> tcp-flags mask {0-0x3F|
0-0x3F}
```

6. Ensure the configuration is correct:

```
show filter acl protocol <1-2048> <1-2000>
```

7. (Optional) Delete the individual attributes from the protocol portion of the ACE:

```
no filter acl ace protocol <1-2048> <1-2000> [dst-port] [icmp-msg-
type] [icmpv6-msg-type] [routing-type] [src-port] [tcp-flags]
```

8. (Optional) Delete all the attributes from the protocol portion of the ACE:

```
default filter acl ace protocol <1-2048> <1-2000>
```

Examples

Specify ICMP packets:

```
Switch:1(config)#filter acl ace protocol 1 12 icmpv6-msg-type eq echoRequest
```

Table 324: TCP Flags Order in Packet

32 (decimal)	16 (decimal)	8 (decimal)	4 (decimal)	2 (decimal)	1 (decimal)
Urgent	Ack	Push	Reset	Syn	Fin

Configure an ACE for TCP flags attributes: Example 1

The mask is set for an 'ack' tcp flag bit regardless of whether any other tcp flag bits are also set:

```
Switch:1(config)#filter acl ace protocol 1 1 tcp-flags mask ack ?
<0-0x3F | 0-63> Mask value <Hex | Decimal>: This six bit mask is a reverse mask where
0:care
                                about, 1:do not care about

Switch:1(config)#filter acl ace protocol 1 1 tcp-flags mask ack 0x2f
Hex Value 20 10 8 4 2 1
TCP Flags _ack _ _ _ _
Binary Value 1 0 1 1 1 1 or in hex = 0x2F
```

Configure an ACE for TCP flags attributes: Example 2

A packet will match this filter if the 3 tcpflag bits are set in the tcp header (and only those 3 bits).

```
Switch:1(config)#filter acl ace protocol 1 1 tcp-flags eq ?
WORD<1-50> Tcp flags
{none | fin | syn | rst | push | ack | urg | undefined}
Switch:1(config)#filter acl ace protocol 1 1 tcp-flags eq syn,push,urg
```

You can configure a functionally equivalent filter with the mask operator as follows:

```
Switch:1(config)#filter acl ace protocol 1 1 tcp-flags mask syn,push,urg 0x0
```

Configure an ACE for TCP flags attributes: Example 3

The mask operator provides more flexibility. For example a packet will match the following filter if the 'syn,push,urg' tcpflag bits are set, regardless of whether any other tcpflag bits are also set:

```
Switch:1(config)#filter acl ace protocol 1 1 tcp-flags mask syn,push,urg ?
<0-0x3F | 0-63> Mask value <Hex | Decimal>: This six bit mask is a reverse mask where
0:care
                                about, 1:do not care about
Switch:1(config)#filter acl ace protocol 1 1 tcp-flags mask syn,push,urg 0x15
```

Configure an ACE for ICMP message type: Example 4

```
filter acl 1 type inPort name "ICMP_TRAFFIC_FILTER"
filter acl port 1 1/3
filter acl ace 1 1
filter acl ace action 1 1 deny count
filter acl ace ethernet 1 1 ether-type eq ip
filter acl ace ip 1 1 src-ip mask 194.183.100.64 0.0.0.15
filter acl ace ip 1 1 dst-ip eq 146.97.137.42
filter acl ace ip 1 1 ip-protocol-type eq icmp
filter acl ace protocol 1 1 icmp-msg-type eq echo-request
filter acl ace 1 1 enable
filter acl ace 1 2
filter acl ace action 1 2 deny count
filter acl ace ethernet 1 2 ether-type eq ip
filter acl ace ip 1 2 src-ip mask 194.183.100.64 0.0.0.15
filter acl ace ip 1 2 dst-ip eq 146.97.137.42
filter acl ace ip 1 2 ip-protocol-type eq icmp
filter acl ace protocol 1 2 icmp-msg-type eq echoreply
filter acl ace 1 2 enable
```

Variable Definitions

Use the data in the following table to use the **filter acl ace protocol** command.

Variable	Value
{0-0x3F}	Specifies the mask value.
<1-2000>	Specifies the ACE ID.
<1-2048>	Specifies the ACL ID.
WORD<1-50>	Specifies one or more TCP flags—none, fin (finish connection), syn (synchronize), rst (reset connection), push, ack (acknowledge), urg (urgent), and undefined.

Variable	Value
<i>WORD<1-60></i>	Specifies the destination port: (0-65535), or echo, ftpdata, ftpcontrol, ssh, telnet, dns, http, hdot323, bootpServer, bootpClient, tftp, rtp, rtcp, or undefined.
<i>WORD<1-200></i>	Specifies the ICMP message type: Icmpmsg type (0-255), or echoreply, destunreach, sourcequench, redirect, echo-request, routeradv, routerselect, time-exceeded, param-problem, timestamp-request, timestamp-reply, addressmask-request, addressmask-reply, or traceroute.
<i>WORD<1-200></i>	Specifies the ICMPv6 message type: Icmpmsg type (0-255), or destUnreach, pktTooBig, timeExceeded, paramProblem, echoRequest, echoReply, mcastListenReq, mcastListenRpt, mcastListenDone, routerSolicit, routerAdvert, neighborSolicit, neighborAdvert, redirectMsg, nodeInfoReq, nodeInfoRsp, or v2McastListenRpt.
<i>WORD<1-256></i>	Specifies the mask parameter, {0-0xFFFF}.
<i>WORD<0-65535></i>	Specifies the source port (0-65535).

Configure Ingress Policer and Port Rate Limiter

Before You Begin

Ensure that both the ACL and the ACE exist.

About This Task

Perform this procedure to limit the number of packets that can be attached to an ACL ACE. If you want to modify the ingress bandwidth rate limiter values, you must first disable the ACE.



Note

You can attach one policer to ACL and ACE. Every policer is independent.

Procedure

1. Enter the Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure an ACL and ACE policer:

```
filter acl ace policer <1-2048> <1-2000> svc-rate <0-4000000000> peak-  
rate <8-4000000000>
```



Note

You must configure both svc-rate and peak-rate.

3. Ensure the configuration is correct:

```
show filter acl policer <1-2048> <1-2000>
```

Variable Definitions

The following table defines parameters for the **filter acl ace policer** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<1-2000>	Specifies the ACE ID.
<i>svc-rate</i> <0-4000000000>	Specifies the rate of the traffic committed to be delivered. Service rate limit in kbps {0-4000000000}.The granularity is 8 kbps.
<i>peak-rate</i> <8-4000000000>	Specifies the value when exceeded causes packets to drop on ingress. Peak rate limit in kbps {8-4000000000}.The granularity is 8 kbps.

View ACL and ACE Configuration Data

View your configuration to review the information and ensure it is correct.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View ACL information:
`show filter acl <1-2048>`
3. View IPv6 ACL information:
`show filter acl ipv6 <1-2048> <1-2000>`
4. View the running configuration for an ACL and corresponding ACE:
`show filter acl config <1-2048> <1-2000>`

Variable Definitions

Use the data in the following table to use the **show filter acl** and **show filter acl config** commands.

Variable	Value
<1-2000>	Specifies the ACE ID.
<1-2048>	Specifies the ACL ID.

View ACE Statistics

View ACE statistics to ensure that the filter operates correctly.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View ACE statistics for a specific ACL, ACE, or ACE type:
show filter acl statistics <1-2048> <1-2000> [qos] [security]
3. View all ACE statistics:
show filter acl statistics all
4. View default ACE statistics:
show filter acl statistics default [<1-2048>]
5. View global statistics for ACEs:
show filter acl statistics global [<1-2048>]
6. View filter policer statistics for ACL and ACE:
show filter acl statistics <1-2048> <1-2000> policer

Examples

View ACE statistics:

```
Switch:1#show filter acl statistics all
=====
Acl Default Statistics Table
=====
Acl Id  Acl Name  Acl Type  Primary Bank  Primary Bank  Secondary Bank  Secondary Bank  Acl  Acl
Packets Bytes  Packets  Bytes  Packets  Bytes  Packets  Bytes
-----
1      ACL-1    inVlan   1360         92480         1360         92480         2720    184960
=====
Displayed 1 of 1 entries
```

View filter acl ace policer statistics:



Note

IPv4 policer statistics are available for red, yellow, and green packets. Statistics display as numbers, which represent the hit number of packets and bytes.

```
Switch:1#show filter acl statistics 3 3 policer
=====
Acl Ace Policer Table
=====
Acl Id  Ace Id  Packets  Bytes  Green  Green  Yellow  Yellow  Red  Red
Packets Bytes  Packets Bytes  Packets Bytes  Packets Bytes
-----
3       3      29964026 106591916 142584 14254 142574 277225 28869904 10275810600
=====
```

View filter acl ace statistics with no policer configured.

```
Switch:1(config)#show filter acl statistics 3 3
=====
Acl Ace Policer Table
=====
Acl Id  Ace Id  Packets  Bytes  Green  Green  Yellow  Yellow  Red  Red
Packets Bytes  Packets Bytes  Packets Bytes  Packets Bytes
-----
10      1001    0         0      N/A    N/A    N/A    N/A    N/A    N/A
=====
Displayed 1 of 1 entries
```

View filter acl ace statistics with filter ACE enabled and policer configured, but the filter is not hit or there is no traffic running.

```
Switch:1(config)#show filter acl statistics 10 1011
=====
```

```
=====
Acl Ace Policer Table
=====
```

Acl Id	Ace Id	Packets	Bytes	Green Packets	Green Bytes	Yellow Packets	Yellow Bytes	Red Packets	Red Bytes
3	3	0	0	0	0	0	0	0	0

```
-----
Displayed 1 of 1 entries
```

View filter acl ace statistics with filter ACE not enabled and policer not configured

```
Switch:1(config)#show filter acl statistics
=====
Acl Ace Policer Table
=====
```

Acl Id	Ace Id	Packets	Bytes	Green Packets	Green Bytes	Yellow Packets	Yellow Bytes	Red Packets	Red Bytes

```
-----
Displayed 1 of 1 entries
```

Variable Definitions

The following table defines parameters for the **show filter acl statistics** command.

Variable	Value
<1-2048>	Specifies the ACL ID.
<1-2000>	Specifies the ACE ID.
<i>policer</i>	Specifies the policer parameter for the command.

Access Control Entry Configuration using EDM

Use an access control entry (ACE) to define a pattern (found in a packet) and the desired behavior for packets that carry the pattern. You can configure up to 1000 ACEs in a single ACL.

As a best practice, create access control lists (ACL) with a default action of permit, and with an ACE mode of deny. For deny or permit ACLs or ACEs, the default action and the mode must be opposite for the ACE (filter) to have meaning.

Configure an ACE

Before You Begin

- The ACL exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the ACL to which to add an ACE.
5. Select **ACE**.
6. Select the **ACE Common** tab.

7. Select **Insert**.
8. Configure the ACE ID.
9. Name the ACE.
10. Choose the mode: **deny** (drop packets) or **permit** (forward packets).
11. Configure the ACE actions as required.
12. Select **Insert**.
13. Configure the ACE attributes as required.
14. To enable the ACE, in the **ACE Common** tab, configure **AdminState** to enable, and then select **Apply**.
15. To delete an ACE Common entry, select the entry, and then select **Delete**.

ACE Common Field Descriptions

Use the data in the following table to use the **ACE Common** tab.

Name	Description
AclId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
Name	Specifies a descriptive user-defined name for the ACE. The system automatically assigns a name if you do not type one.
AdminState	Indicates the status of the ACE as enabled or disabled. You can modify an ACE only if you disable it.
OperState	Indicates the current operational state of the ACE.
Mode	Indicates the operating mode for this ACE. Valid options are deny and permit, with deny as the default.
RedirectNextHop	Redirects matching IPv4/IPv6 traffic to IPv4/IPv6 nexthop.
RedirectNextHopVrfname	Specifies the direct next hop VRF name. The name must be in the range of 1 to 16 characters.
RedirectUnreach	Denies or permits packet dropping when the next hop for the packet is unreachable. The default value is deny. This action is a security action.
InternalQos	This variable is a QoS action. The default value is 1.
RemarkDscp	Specifies whether the DSCP parameter marks nonstandard traffic classes and local-use Per-Hop Behavior. The default is disable. Use this option to create a QoS ACE.
RemarkDot1Priority	Specifies whether Dot1 Priority, as described by Layer 2 standards (802.1Q and 802.1p) is enabled. The default is disable. Use this option to create a QoS ACE.
Count	Specifies whether the count action in the ACE is enabled or disabled. The default is disabled.

Configure ACE Actions

Configure ACE actions to determine the process that occurs after a packet matches (or does not match) an ACE. Use debug actions (flags) to use filters for troubleshooting and monitoring procedures.

Before You Begin

- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select an **AceId**.
7. Select **Action**.
8. Configure the actions as required, and then select **Apply**.

Action Field Descriptions

Use the data in the following table to use the **Action** tab.



Note

The table lists the options for both Security ACEs and QoS ACEs. Dependent upon the ACE, the system displays different options on the EDM interface.

Name	Description
AcId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
Mode	Configures the action mode for security ACEs. The default value is deny.
RemarkDscp	Specifies the new Per-Hop Behavior (PHB) for matching packets: phbcs0, phbcs1, phbaf11, phbaf12, phbaf13, phbcs2, phbaf21, phbaf22, phbaf23, phbcs3, phbaf31, phbaf32, phbaf33, phbcs4, phbaf41, phbaf42, phbaf43, phbcs5, phbef, phbcs6, phbcs7. This action is a QoS action.
RemarkDot1Priority	Specifies the new 802.1 priority bit for matching packets: zero, one, two, three, four, five, six, or seven. This action is a QoS action. Note: This does not apply to IPv6 filtering.

Name	Description
InternalQoS	This variable is a QoS action. The default value is 1. Note: This does not apply to IPv6 filtering.
RedirectNextHop	Specifies the next-hop IPv4 address (a.b.c.d) or IPv6 address (aaaa:bbbb:cccc:dddd:eeee:ffff:gggg:hhhh) for redirect mode. Applies to ingress ACLs (routed and Layer 2 packets).
Count	Enables the ability to count matching packets. Use this parameter with either a security or QoS ACE. The default is disabled.
DstPortList	Configures mirroring to a destination port or ports. This action is a security action.
DstMltd	Configures mirroring to a destination MLT. This action is a security action.
MonitoringIsidOffset	Configures the monitoring I-SID offset value. The offset ID is mapped to the actual monitor I-SID value to which the packets are mirrored.
MirroringQoS	Defines the Quality of Service (QoS) profiles for the mirrored packet into monitoring I-SID.

Configure ACE ARP Entries

Use ACE Address Resolution Protocol (ARP) entries so that the filter looks for ARP request or response packets.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select a parameter for the appropriate ACL.
5. Select **ACE**.
6. Select a parameter for the appropriate ACE.
7. Select **Arp**.
8. Select **Insert**.

9. Select ARP request or response.
10. Select **Insert**.

ARP Field Descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
AclId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
Type	Specifies the ACE ARP operation. The only option is operation.
Oper	Specifies the operator for the ACE ARP operation. The only option is eq (equal).
Value	Specifies the ARP packet type. Valid options are arpRequest and arpResponse.

View ACE ARP Entries for an ACL

View all of the ACE ARP entries associated with an ACL.

Procedure

1. In the navigation tree, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **Arp**.
6. Select the parameter and modify as you need.
7. Select **Apply**.

ARP Field Descriptions

Use the data in the following table to use the **ARP** tab.

Name	Description
AclId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
Type	Specifies the ACE ARP operation. The only option is operation.
Oper	Specifies the operator for the ACE ARP operation. The only option is eq (equal).
Value	Specifies the ARP packet type. Valid options are arpRequest and arpResponse.

Configure an ACE Ethernet Source Address

Use this procedure to filter on specific Ethernet source addresses.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Eth**.
8. Select the **Source Address** tab.
9. Select **Insert**.
10. Specify the ACE Ethernet operation.
11. In the **List** dialog box, specify the Ethernet source address.
12. Select **Insert**.

Source Address Field Descriptions

Use the data in the following table to use the **Source Address** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
List	Specifies the MAC address to match.
OperMask	Specifies the MAC Address mask value in hexadecimal format. The value for this variable is empty or 000000000000 if the Oper variable is eq.

Configure an ACE Ethernet Destination Address

Perform this procedure to filter on specific Ethernet destination addresses.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Eth**.
8. Select the **Destination Address** tab.
9. Select **Insert**.
10. Specify the ACE Ethernet operation.
11. In the **List** dialog box, specify the Ethernet source address.
12. Select **Insert**.

Destination Address Field Descriptions

Use the data in the following table to use the **Destination Address** tab.

Variable	Value
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
List	Specifies the MAC address to match.
OperMask	Specifies the MAC address mask value in hexadecimal format if the Oper variable is mask. The value of this variable is empty or 000000000000 if Oper is eq.

Configure an ACE LAN Traffic Type

Perform this procedure to filter for specific LAN traffic packets.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Eth**.

8. Select the **Ethernet Type** tab.
9. Select **Insert**.
10. Specify the operation type.
11. In the **TypeList** box, type the Ethernet types.
12. Select **Insert**.

Ethernet Type Field Descriptions

Use the data in the following table to use the **Ethernet Type** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
TypeOper	The eq parameter specifies an operator for a field match condition: equal to.
TypeList	<p>Specifies the Ethernet type. Entries include: 0 to 0xffff or ip, arp, ipx802.3, ipx802.2, ipxSnap, ipxEthernet2, appleTalk, appleTalk-ARP, sna802.2, snaEthernet2, netBios, xns, vines, ipv6, rarp, PPPoE-discovery, and PPPoE-session.</p> <p>Note: Ethernet ACE filter configured with Ethernet Type ipx802.3 does not match the packet with format destination MAC address, source MAC address, length, 0xFFFF, payload and FCS. Ethernet ACE filter configured with Ethernet Type ipx802.2 does not match the packet with format destination MAC address, source MAC address, length, 0xE0E0, payload and FCS.</p>

Configure an ACE Ethernet VLAN Tag Priority

Perform this procedure to filter for specific VLAN tag priorities.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Eth**.

8. Select the **Vlan Tag Priority** tab.
9. Select **Insert**.
10. Specify the operation type.
11. In the **VlanTagPrio** box, select the priority bits.
12. Select **Insert**.

VLAN Tag Priority *Field Descriptions*

Use the data in the following table to use the **Vlan Tag Priority** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
VlanTagPrio	Specifies the priority bits (3-bit field) from the 802.1Q/p tag: <ul style="list-style-type: none"> • zero • one • two • three • four • five • six • seven
OperMask	Specifies the mask value in hexadecimal format if the Oper value is mask.

Configure an ACE Ethernet Port

Use ACE Ethernet port entries so that the filter looks for traffic on specific ports. You can only insert an ACE Common Ethernet port for VLAN ACL types.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Eth**.

8. Select the **Port** tab.
9. Select **Insert**.
10. Specify the operation type.
11. Select the **Port** ellipses (...).
12. Choose the ports.
13. Select **OK**.
14. Select **Insert**.

Port Field Descriptions

Use the data in the following table to use the **Port** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq parameter specifies an operator for a field match condition: equal to.
Port	Specifies the port or port list on which to perform a match.

Configure an ACE Ethernet VLAN ID

Use ACE Ethernet VLAN ID entries so that the filter looks for traffic on specific VLANs. You can insert an ACE Ethernet VLAN ID only for ACL VLAN types.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Eth**.
8. Select the **Vlan Id** tab.
9. Select **Insert**.
10. Specify the operation type.
11. Enter the VLAN ID or select from a list.
12. Select **Insert**.

VLAN ID *Field Descriptions*

Use the data in the following table to use the **Vlan Id** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq parameter specifies an operator for a field match condition: equal to.
VlanId	Specifies the VLAN ID on which to perform a match.
OperMask	Specifies the mask value for a VLAN attribute.

View All ACE Ethernet Entries for an ACL

View all of the ACE Ethernet entries associated with an ACL.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **Eth**.

Ethernet *Field Descriptions*

Use the data in the following table to use the **Ethernet** tab.

Name	Description
AcId	Shows the ACL ID.
AcId	Specifies the ACE ID.
SrcAddrList	Shows the list of Ethernet source addresses to match.
SrcAddrOper	Shows the operators for the ACE Ethernet source MAC address.
SrcAddrOperMask	Shows the source MAC address mask value in hexadecimal format if the SrcAddrOper variable is mask. The value of this field is empty or 000000000000 if the SrcAddrOper field is eq.
DstAddrList	Shows the list of Ethernet destination addresses to match.
DstAddrOper	Shows the operators for the ACE Ethernet destination MAC address.
DstAddrOperMask	Shows the destination MAC address mask value in hexadecimal format if the DstAddrOper variable is mask. The value for this field is empty or 000000000000 if the DstAddrOper field is eq.

Name	Description
EtherTypeList	Shows the EtherType value from the Ethernet header. For example, ARP uses 0x0806 and IP uses 0x0800. Platform support determines the behavior for 802.1Q/p tagged packets. The EtherType for 802.1Q tagged frames is 0x8100. The range is 0–65535 and supports lists and ranges of values. An invalid Ether-type of 65536 indicates that you do not want the parameter in the match criteria.
EtherTypeOper	Shows the Ethernet type operators.
VlanTagPrio	Shows the priority bits (3-bit field) from the 802.1Q/p tag.
VlanTagPrioOper	Shows the operators for the ACE Ethernet VLAN tag priority.
VlanTagPrioOperMask	Shows the VLAN tag priority mask value in hexadecimal format if the VlanTagPrioOper field is mask.
Port	Shows the port number or port list to match.
PortOper	Shows the operator for the ACE Ethernet port.
VlanId	Shows the VLAN ID to match.
VlanIdOper	Shows the operator for the ACE Ethernet VLAN ID.
VlanIdOperMask	Shows the VLAN ID mask value in hexadecimal format if the VlanIdOper field is mask.

Configure an ACE IP Source Address

Configure ACE IP source address entries to have the filter look for specific source IP addresses.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IP**.
8. Select the Source Address tab.
9. Select **Insert**.

10. Specify the operation type.
11. In the **IPAddr** box, enter the source IP address.
12. Select **Insert**.

Source Address *Field Descriptions*

Use the data in the following table to use the **Source Address** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
IPAddr	Specifies the source IP address.
OperMask	Specifies the mask value for the source IP address.

Configure an ACE IP Destination Address

Configure ACE IP destination address entries to have the filter look for specific destination IP addresses.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IP**.
8. Select the **Destination Address** tab.
9. Select **Insert**.
10. Specify the operation type.
11. In the **IPAddr** box, enter the destination IP address.
12. Select **Insert**.

Destination Address *Field Descriptions*

Use the data in the following table to use the **Destination Address** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
IPAddr	Specifies the destination IP address.
OperMask	Specifies the mask value for the destination IP address.

Configure an ACE IP DSCP

Configure ACE IP DSCP entries to have the filter look for packets with specific DSCP markings.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IP**.
8. Select the **DSCP** tab.
9. Select **Insert**.
10. Specify the operation type.
11. In the **List** box, enter the count for the DSCP values.
12. Select **Insert**.

DSCP *Field Descriptions*

Use the data in the following table to use the **DSCP** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.

Name	Description
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
List	Specifies a count for the number of discrete ranges entered for the DSCP values. Entries include 0–256, disable, phbcs0, phbcs1, phbaf11, phbaf12, phbaf13, phbcs2, phbaf21, phbaf22, phbaf23, phbcs3, phbaf31, phbaf32, phbaf33, phbcs4, phbaf41, phbaf42, phbaf43, phbcs5, phbef, phbcs6, and phbcs7.
OperMask	Specifies the mask value.

Configure an ACE IP Protocol

Configure ACE IP protocol entries to have the filter look for packets of specific protocols.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IP**.
8. Select the **Protocol** tab.
9. Select **Insert**.
10. Specify the operation type.
11. In the **List** box, enter the IP protocol type.
12. Select **Insert**.

Protocol Field Descriptions

Use the data in the following table to use the **Protocol** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq parameter specifies an operator for a field match condition: equal to.
List	Specifies the IP protocol type. Entries include 0–256, undefined, tcp, udp, ipsecesp, vrrp, and undefined.

Configure ACE IP Options

Configure ACE IP option entries to have the filter look for packets with an IP option specified.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IP**.
8. Select the **Options** tab.
9. Select **Insert**.
10. Specify the logical operator.
Any is the only choice.
11. Select **Insert**.

Options Field Descriptions

Use the data in the following table to use the **Options** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	Specifies the logical operator for the ACE IP options. Any is the only option.

Configure ACE IP Fragmentation

Configure ACE IP fragmentation entries to have the filter look for packets with the fragmentation flag.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.

4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IP**.
8. Select the **Fragmentation** tab.
9. Select **Insert**.
10. Specify the operator for IP fragmentation.
Eq is the only choice.
11. Specify the fragmentation bits to match from the IP header.
12. Select **Insert**.

Fragmentation Field Descriptions

Use the data in the following table to use the **Fragmentation** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	Specifies the logical operator for the ACE IP options. Any is the only option.
Fragmentation	Specifies the IP fragmentation bits to match from the IP header: <ul style="list-style-type: none"> • noFragment • anyFragment The default is noFragment.

View all ACE IP Entries for an ACL

View all of the ACE IP entries associated with an ACL.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **IP**.

IP Field Descriptions

Use the data in the following table to use the **IP** tab.

Name	Description
AcId	Shows the ACL IP ID.
AcId	Shows the ACE ID.
SrcAddrOper	Shows the operators for the ACE IP source address.

Name	Description
SrcAddrIpAddr	Shows the IP source address to match from the IP header.
SrcAddrOperMaskRange	Shows the IP mask value if SrcAddrOper is set to mask, or the highest IP address if SrcAddrOper is set to range.
DstAddrOper	Shows the operators for the ACE IP destination address.
DstAddrIpAddr	Shows the IP destination address to match from the IP header.
DstAddrOperMaskRange	Shows the IP mask value if DstAddrIpAddr is set to mask, or the highest IP address if DstAddrIpAddr is set to range.
DscpList	Shows how the 6-bit DSCP parameter from the TOS byte in the IPv4 header encodes PHB information following RFC 2474.
DscpOper	Shows the operators for the ACE IP DSCP.
DscpOperMask	Shows the mask value in hexadecimal format when the mask option is selected in DscpOper .
ProtoList	Shows the IP protocol type from the IP header to match. The range is 0–255.
ProtoOper	Shows the operators for the ACE IP protocols.
Options	Shows the IP options to match from the IP header.
OptionsOper	Shows the logical operator. Any is the only option.
Fragmentation	Shows the IP fragmentation bits to match from the IP header.
FragOper	Shows the operator for IP fragmentation.

Configure an ACE IPv6 Source Address

Configure ACE IPv6 source address entries to have the filter look for specific source IP addresses.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.

7. Select **IPv6**.
8. Select the **Source Address** tab.
9. Select **Insert**.
10. In **Oper**, select the operation type.
11. In **List**, enter the source IP address.
12. In **OperMask**, enter the operation mask value.
13. Select **Insert**.

Source Address *Field Descriptions*

Use the data in the following table to use the **Source Address** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
List	Specifies the source IP address.
OperMask	Specifies the mask value for the source IP address.

Configure an ACE IPv6 Destination Address

Configure ACE IPv6 destination address entries to have the filter look for specific destination IP addresses.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IPv6**.
8. Select the **Destination Address** tab.
9. Select **Insert**.
10. In **Oper**, select the operation type.
11. In **List**, enter the destination IP address.
12. In **OperMask**, enter the operation mask value.
13. Select **Insert**.

Destination Address Field Descriptions

Use the data in the following table to use the **Destination Address** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
List	Specifies the destination IP address.
OperMask	Specifies the mask value for the destination IP address.

Configure an ACE IPv6 Next Header

Configure ACE IPv6 next header entries to have the filter look for specific next headers.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IPv6**.
8. Select the **Next Hdr** tab.
9. Select **Insert**.
10. In **Oper**, select the operation type.
11. In **NextHdr**, select the next header type.
12. Select **Insert**.

Next Hdr Field Descriptions

Use the data in the following table to use the **Next Hdr** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.

Name	Description
Oper	The eq parameter specifies an operator for an “equal to” field match condition.
NextHdr	Specifies the next header of the IPv6 header. Specifies hop-by-hop, tcp, udp, routing, fragment, ipsecESP, ipsecAH, icmpv6, noNxtHdr, or undefined.

Configure an ACE IPv6 Traffic Class

Configure ACE IPv6 traffic class.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IPv6**.
8. Select the **Traffic Class** tab.
9. Select **Insert**.
10. In **Oper**, select the operation type.
11. In **TrafficCls**, type the traffic class number.
12. Select **Insert**.

Traffic Class Field Descriptions

Use the data in the following table to use the **Traffic Class** tab.

Name	Description
AclId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
Oper	The eq parameter specifies an operator for an “equal to” field match condition.
TrafficCls	Specifies the traffic class attribute of the IPv6 header. Traffic class identifies different classes or priorities of IPv6 packets. The range is 0–255.

View all ACE IPv6 Entries for an ACL

View all of the ACE IPv6 entries associated with an ACL.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **IPv6**.
8. Select the **IPv6** tab.

IPv6 Field Descriptions

Use the data in the following table to use the **IPv6** tab.

Name	Description
AcId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
SrcAddrList	Shows the source IP address.
SrcAddrOper	Shows the operators for the ACE IP source address.
DstAddrList	Shows the destination IP address.
DstAddrOper	Shows the operators for the ACE IP destination address.
NxtHdrNxtHdr	Shows the next header of the IPv6 header.
NxtHdrOper	Shows the operators for the next header.
TrafficClsOper	Shows the operators for the traffic class.
TrafficCls	Shows the traffic class attribute of the IPv6 header.
SrcAddrMask	Shows the mask value for the source IP address.
DstAddrMask	Shows the mask value for the destination IP address.

Configure an ACE Source Port

Configure ACE source port entries to have the filter look for packets with a specific source port.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.

2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Proto**.
8. Select the **Source Port** tab.
9. Select **Insert**.
10. Specify the operator for the source port.
11. Specify the port number or port list to match.
12. Select **Insert**.

Source Port *Field Descriptions*

Use the data in the following table to use the **Source Port** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Port	Specifies the source port (1–65535).
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
OperMask	Specifies the mask parameter, {0-0xFFFF}.

Configure an ACE Destination Port

Configure ACE destination port entries to have the filter look for packets with a specific destination port.

Before You Begin

- The ACE exists.
- The ACL exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Proto**.
8. Select the **Destination Port** tab.
9. Select **Insert**.

10. Specify the operator for the destination port.
11. Specify the port number or port list to match.
12. Select **Insert**.

Destination Port Field Descriptions

Use the data in the following table to use the **Destination Port** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
Port	Specifies the port number. As noted at the bottom of the tab, potential entries include 0–65535, echo, ftpdata, ftpcontrol, ssh, telnet, dns, http, h.323, and undefined.
OperMask	Specifies the mask parameter, {0-0xFFFF}.

Configure an ACE ICMP Message Type

Configure ACE Internet Control Message Protocol (ICMP) message type entries to have the filter look for packets of a specific ICMP message type.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Proto**.
8. Select the **Icmp Msg Type** tab.
9. Select **Insert**.
10. Specify the operator for the ICMP message type.
11. In the **List** box, specify the ICMP messages to match.
12. Select **Insert**.

Icmp Msg Type Field Descriptions

Use the data in the following table to use the **Icmp Msg Type** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	Specifies the operator for the ACE protocol ICMP message type. Equal (eq) is the only option.
List	Specifies the ICMP message type (0–255), or echoreply, destunreach, sourcequench, redirect, echo-request, routeradv, routerselect, time-exceeded, param-problem, timestamp-request, timestamp-reply, addressmask-request, addressmask-reply, or traceroute.

Configure an ACE ICMPv6 Message Type

Before You Begin

- The ACL exists.
- The ACE exists.

About This Task

Configure ACE Internet Control Message Protocol v6 (ICMPv6) message type entries to have the filter look for packets of a specific ICMPv6 message type.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Proto**.
8. Select the **Icmpv6 Msg Type** tab.
9. Select **Insert**.
10. Specify the operator for the ICMPv6 message type.
11. In **List**, specify the ICMPv6 messages to match.
12. In **Count**, specify 1 through 100.
13. Select **Insert**.

Icmpv6 Msg Type *Field Descriptions*

Use the data in the following table to use the **Icmpv6 Msg Type** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	Specifies the operator for the ACE protocol ICMPv6 message type. Equal (eq) is the only option.
List	Specifies the ICMPv6 message type (0–255), or echoreply, destunreach, sourcequench, redirect, echo-request, routeradv, routerselect, time-exceeded, param-problem, timestamp-request, timestamp-reply, addressmask-request, addressmask-reply, or traceroute.
Count	Specifies 1–100. Enables the ability to count matching packets. Use this parameter with either a security or QoS ACE. The default is disabled.

Configure an ACE TCP Flag

Configure ACE TCP flag entries to have the filter look for packets with a specific TCP flag.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **ACE**.
6. Select the appropriate ACE.
7. Select **Proto**.
8. Select the **TCP Flags** tab.
9. Select **Insert**.
10. Specify the operator for the TCP flags entry.
11. In the **List** box, specify the TCP flags to match.
12. Select **Insert**.

TCP Flags Field Descriptions

Use the data in the following table to use the **TCP Flags** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Oper	The eq and mask parameters specify an operator for a field match condition: equal to or mask. The mask operator is an implied eq on the mask bits.
List	Specifies one or more TCP flags—none, fin (finish connection), syn (synchronize), rst (reset connection), push, ack (acknowledge), urg (urgent), and undefined.
OperMask	Specifies the mask value.

Configure an ACE to Filter IPv4 Routed Packets

Configure an ACE to filter IPv4 routed packets.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select **IpRouted** tab.
4. Select **Insert**.
5. In **AcId**, type the appropriate ACL ID.
6. In **AcId**, type the appropriate ACE ID.
7. **RoutedOnly**
8. Select **Insert**.

IpRouted Field Descriptions

The following table describes values on the **IpRouted** tab.

Name	Description
AcId	Specifies the ACL ID. Value range of 1 to 2048.
AcId	Specifies the ACE ID. Value range of 1 to 2000.

Configure an ACE to Filter IPv6 Routed Packets

Configure an ACE to filter IPv6 routed packets only.

Before You Begin

- The ACL exists.
- The ACE exists.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **Ipv6Routed** tab.
4. Select **Insert**.
5. In **AcId**, type the appropriate ACL ID.
6. In **AceId**, type the appropriate ACE ID.
7. Select **RoutedOnly**.
8. Select **Insert**.

Ipv6Routed *Field Descriptions*

The following table describes values on the **Ipv6Routed** tab.

Name	Description
AcId	Specifies the ACL ID.
AceId	Specifies the ACE ID.
RoutedOnly	Specifies the match for IPv6 routed packets only.

Configure Ingress Policer and Port Rate Limiter

Perform this procedure to limit the number of packets that can be attached to an ACL ACE. If you want to modify the ingress bandwidth rate limiter values, you must first disable the ACE.

**Note**

You can attach one policer to ACL and ACE. Every policer is independent.

Before You Begin

Ensure that both the ACL and the ACE exist.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the appropriate ACL.
4. Select the appropriate ACE.
5. Select **Policer**.
6. Select **Insert**.
7. For **SvcRateValue**, type the service rate value.

8. For **PeakRateValue**, type the peak rate value.
9. Select **Insert**.
10. Select **Apply**.

Policer Field Descriptions

Use the data in the following table to use the **Policer** tab.

Name	Description
AcId	Specifies the ACL identifier numeric value for the filter.
AcId	Specifies the ACE identifier numeric value for the filter.
SvcRateValue	Specifies the service rate in kilobits for the flow of packets from the filter. The range is 8 to 4000000000 and the granularity is 8 kilobits.
PeakRateValue	Specifies peak rate in kilobits for the flow of packets from the filter. The range is 8 to 4000000000 and the granularity is 8 kilobits.
Oper	Specifies the operational status of the policer.

View ACE Port Statistics

About This Task

Use port statistics to ensure that the ACE is operating correctly.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select a field on the **ACL** tab.
5. Select **ACE**.
6. Select the **Statistics** tab.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.

Name	Description
AcId	Specifies the associated ACL index.
AcId	Specifies the ACE index.
MatchCountPkts	Specifies a packet count of the matching packets.
MatchCountOctets	Specifies the number of octets of the matching packets.
GreenBytes	Specifies the number of bytes sent with a rate lower than the committed-rate of the matching ACE.

Name	Description
GreenPackets	Specifies the number of packets sent with a rate lower than the committed-rate of the matching ACE.
YellowBytes	Specifies the number of bytes sent with a rate lower than peak-rate and higher than the committed-rate of the matching ACE.
YellowPackets	Specifies the number of packets sent with a rate lower than the peak-rate and higher than the committed-rate of the matching ACE.
RedBytes	Specifies the number of bytes sent with a rate higher than the peak-rate of the matching ACE.
RedPackets	Specifies the number of packets sent with a rate higher than the peak-rate of the matching ACE.

View All ACE Protocol Entries for an ACL

View all of the ACE protocol entries associated with an ACL.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the appropriate ACL.
5. Select **Proto**.

Protocol Field Descriptions

Use the data in the following table to use the **Protocol** tab.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
SrcPort	Specifies the port number or port list to match.
SrcPortOper	Specifies the operator for the ACE protocol source port.
SrcPortOperMaskRange	The value is displayed in hexadecimal format when SrcPortOper is set to mask. When SrcPortOper is set to range, this field is used as the high range value. in this case, the value is displayed in decimal format. When SrcPortOper is set to eq, this field is set to 0.
DstPort	Specifies port number or port list to match.
DstPortOper	Specifies the operator for the ACE protocol destination port.

Name	Description
DstPortOperMaskRange	The value is displayed in hexadecimal format when DstPortOper is set to mask. When DstPortOper is set to range, this field is used as the high range value. In this case, the value is displayed in decimal format. When SrcPortOper is set to eq, this field is set to 0.
IcmpMsgTypeList	Specifies one or a list of ICMP messages to match. The valid range is 0–255 (reserved).
IcmpMsgTypeOper	Specifies the operator for the ACE protocol ICMP message types.
TcpFlagsList	Specifies one or a list of TCP flags to match. The valid range is 0–63.
TcpFlagsOper	Specifies the operator for the ACE protocol TCP flags.
TcpFlagsOperMask	Displays the mask value in hexadecimal format when TcpFlagsOper is set to mask. When TcpFlagsOper is set to eq, this field displays 0x0.



Troubleshooting

[Data Collection Required for Technical Support Cases](#) on page 3789

[Troubleshooting Fundamentals](#) on page 3792

[Connectivity Fault Management](#) on page 3799

[Software Troubleshooting Tool Configuration](#) on page 3865

[Hardware Troubleshooting](#) on page 3933

[Upper Layer Troubleshooting](#) on page 3943

The Troubleshooting section describes common problems and error messages with the techniques to resolve them, as well as information about troubleshooting tools.

The topics in this section provide necessary concepts and procedures to troubleshoot issues with the switch and connectivity issues in the network.

Data Collection Required for Technical Support Cases

If you need to call Technical Support for assistance, you might need to provide specific switch information or files, like core files, to help the team analyze the issue.

The topics in this section provide details about how to gather information before you contact Technical Support.



Important

The following list provides best practices for working with core files:

- Use Secure Copy (SCP) to transfer the files from the switch. If you use FTP, ensure you transfer the files in binary mode. Do not transfer the files in ASCII mode.
- After you download a core file from the switch, extract it to ensure it downloaded correctly and completely. If you receive errors, download the file again.
- Do not delete the core files from the switch until Technical Support confirms that they can read the files.

Data Collection for an Outage

Perform the following data collection procedures when the switch is in an outage condition and you require Technical Support to perform a root cause analysis.

Collect Data Before You Restart

Perform this procedure before you restart the chassis.

Procedure

1. Capture the current state of the chassis:

```
terminal more disable
```

```
show tech
```

**Note**

`show tech` is an alias for `show fulltech`. Both commands provide the same output.

2. Capture Flight Recorder trace information.

```
flight-recorder all {slot[-slot][,...]}
```

The **all** command executes three separate commands: `flight-recorder snapshot`, `flight-recorder trace`, and `flight-recorder archive`.

3. Reset the chassis:

- a. Reset the chassis without creating a core file:

```
reset -y
```

- b. Create an ssio core file, a cbcx-main.x core file, skip the confirmation question, and then reset the chassis:

```
reset -coredump -y
```

- c. Create an ssio core file, a cbcx-main.x core file, prompt for the confirmation question, and then reset the chassis:

```
reset -coredump
```

**Note**

Create a core file only when there is a need to analyze a problem. If you reset the switch for any other reason the command is `reset -y`.

4. Continue with [Collecting data after you restart](#).

Example

The following example shows output of the `flight-recorder all 1` command.

```
Switch:1#flight-recorder all 1
Processing Flight-recorder snapshot for 1 ....

Flight-recorder snapshot for slot 1 complete, filename is /intflash/PMEM/1/pmem.
20111019114431.1.bin.gz.

Processing Flight-recorder trace for 1 ....

Flight-recorder trace for slot 1 complete, filename is /intflash/flrec/1/trace.2
0111019114434.1.txt.

Processing Flight-recorder archive for slot 1 ....
```

```
Flight-recorder archive for slot 1 complete, filename is /intflash/archive/1/archive.20111019114446.1.tar.
```

Collect Data After You Restart

About This Task

Perform this procedure after you restart the affected chassis.

Procedure

1. Use FTP to transfer the following information:
 - Configuration files from each chassis: Stored on the internal flash at `/intflash/`.
 - Log files from each chassis: Stored on the internal flash at `/intflash/`.
 - Generated archive files for slot: Stored on the internal flash. For example: `/intflash/archive/<slot>`



Note

For , if the core file is not on the primary control processor (CP), you can use FTP to connect to the internal flash of the primary CP. For example, `cp mnt/intflash/archive/<slot>/<file name> /intflash/<file name>`

To copy the file to internal flash or usb device in the primary CP, use the following command: `cp mnt/intflash/archive/<slot>/<file name> {/intflash/<file name> | /usb/<file name>}`

2. Show core information:

```
show core-files
```

If the timestamp for an entry in the command output matches the time the outage first occurred, or is later than that time, transfer that core file to an FTP server. Core files are stored on the internal flash at: `/intflash/coreFiles/`

3. Obtain the network diagram of the relevant nodes, down to the port level.

Data Collection for Non Outage Problems

Use the information in this section to collect data for problems that are less service-impacting than an outage.

Gathering Critical Information

This section identifies the critical information that you must gather before you contact Technical Support.

You must attempt to resolve the problem using this document. Contact Technical Support as a final step taken only after you are unable to resolve the issue using the information and steps provided in this document.

Gather the following information before you contact Technical Support:

- a detailed description of the problem
- the date and time when the problem started
- the frequency of the problem
- if this is a new installation
- if there is relevant information recorded on the support portal — Were related problem solutions found? Is there currently a work around for this issue?
For more information, see support on the Extreme Networks Portal at <https://extremeportal.force.com/ExtrSupportHome>.
- if the system was recently upgraded — Have you recently changed or upgraded the system, the network, or a custom application? (For example, has configuration or code been changed?) When were these changes made? Provide the date and time. Who made these changes? Were the changes made by a partner or customer? Provide the names of the individuals who made the changes.

Troubleshooting Fundamentals

You can better troubleshoot the problems on the network by planning for these events in advance. To do this, you must know the following:

- that the system is properly installed and routinely maintained
- the configuration of the network
- the normal behavior of the network

Network Configuration

To keep track of the network configuration, gather the information described in the following sections. This information, when kept up-to-date, is extremely helpful for locating information if you experience network or device problems.

Site network map

A site network map identifies where each device is physically located on site, which helps locate the users and applications that a problem affects. You can use the map to systematically search each part of the network for problems.

Logical connections

The switch supports virtual LANs (VLAN). With VLANs, you must know how the devices connect logically as well as physically.

Device configuration information

Maintain online and paper copies of the device configuration information. Store all online data with the regular data backup for the site. If the site does not use a backup system, copy the information onto an external storage device, and store the backup at an offsite location.

You can use the File Transfer Protocol (FTP) and Trivial FTP (TFTP) to store configuration files on a remote server.

Other important data about the network

For a complete picture of the network, have the following information available:

- all passwords

Store passwords in a safe place. A good practice is to keep records of previous passwords in case you must restore a device to a previous software version and need to use the old password that was valid for that version.

- device inventory

Maintain a device inventory, which lists all devices and relevant information for the network. The inventory allows you to easily see the device type, IP address, ports, MAC addresses, and attached devices.

- MAC address-to-port number list

If you do not manage the hubs or switches, you must keep a list of the MAC addresses that correlate to the ports on the hubs and switches.

- change control

Maintain a change control system for all critical systems. Permanently store change control records.

- contact details

Store the details of all support contracts, support numbers, engineer details, and telephone and fax numbers.

Normal Behavior on the Network

If you are familiar with the network when it is fully operational, you can be more effective at troubleshooting problems that arise. To understand the normal behavior of the network, monitor the network over a long period of time. During this time you can see a pattern in the traffic flow, such as which devices users access most or when peak usage times occur.

To identify problems, you can use a baseline analysis, which is an important indicator of overall network status. A baseline serves as a useful reference of network traffic during normal operation, which you can then compare to captured network traffic while you troubleshoot network problems. A baseline analysis speeds the process of isolating network problems. By running tests on a properly functioning network, you compile normal data for your network. You can compare this normal data against the results that you get when the network experiences trouble.

For example, ping each node to discover how long it typically takes to receive a response from devices on your network. Capture and save each response time and you can use these baseline response times to help you troubleshoot. You can also use the **show tech/show fulltech** and **show khi performance {buffer-pool|cpu|memory|process|pthread|slabinfo}** commands to obtain baseline output for normal system behavior.

In the following example, the **show khi performance memory** command shows the average memory utilization at various time intervals. The **show khi performance memory history** command shows the VMSize utilization values in kilobytes for each process at various time intervals. After 1 hour elapses, the system stores this information in `/intflash/coreFiles/slot/khi_mem_log`.

```
Switch:1>show khi performance memory 1
Slot:1
  Used: 6259174 (KB)
  Free: 1933970 (KB)
  Current utilization: 76 %
  5-minute average utilization: 76 %
  5-minute high water mark: 80 (%)
  10-minute average utilization: 76 %
  10-minute high water mark: 80 (%)
  1-Hour average utilization: 76 %
  1-Hour high water mark: 79 (%)
  1-Day average utilization: 76 %
  1-Month average utilization: 0 %
  1-Year average utilization: 0 %
```

SMLT socket buffer memory usage 0 KB on 0 blocks

```
Switch:1>show khi performance memory history
```

Slot:1

Values indicate VMSize(VMRSS) in KB

Pid	Pname	5-Min	10-Min	1-Hour	1-Day	1-Month	1-Year
2289	logger	2 (1)	2 (1)	2 (1)	2 (0)	--	--
(--)	--(--)						
2305	namServer	851 (7)	851 (7)	851 (7)	851 (0)	--	--
(--)	--(--)						
2306	sockserv	5 (2)	5 (2)	5 (2)	5 (0)	--	--
(--)	--(--)						
2307	oom95	1178 (403)	1178 (403)	1178 (403)	1178 (0)	--	--
(--)	--(--)						
2308	oom90	1178 (403)	1178 (403)	1178 (403)	1178 (0)	--	--
(--)	--(--)						
2309	imgsync.x	788 (7)	788 (7)	788 (7)	788 (0)	--	--
(--)	--(--)						
2380	logServer	794 (12)	794 (12)	794 (12)	794 (0)	--	--
(--)	--(--)						
2381	trcServer	787 (7)	787 (7)	787 (7)	787 (0)	--	--
(--)	--(--)						
2382	oobServer	787 (7)	787 (7)	787 (7)	787 (0)	--	--
(--)	--(--)						
2383	nickServer	724 (8)	724 (8)	724 (8)	724 (0)	--	--
(--)	--(--)						
2384	nickClient	724 (7)	724 (7)	724 (7)	724 (0)	--	--
(--)	--(--)						
3363	dhclient-ipv4-oob	81 (9)	81 (9)	81 (9)	80 (0)	--	--
(--)	--(--)						
2385	hwsServer	1688 (8)	1688 (8)	1688 (8)	1688 (0)	--	--
(--)	--(--)						
2386	redis-server	20 (5)	20 (5)	20 (5)	20 (0)	--	--
(--)	--(--)						
3405	hiveagent	13 (5)	13 (5)	13 (5)	12 (0)	--	--
(--)	--(--)						

2387	bundle-recovery.py	--(--)	--(--)	--(--)	--(--)	--
(--)	--(--)					
2388	cbcp-main.x	2570 (482)	2570 (482)	2570 (482)	2569 (0)	--
(--)	--(--)					
2392	rssServer	787 (6)	787 (6)	787 (6)	787 (0)	--
(--)	--(--)					
2393	dbgServer	787 (7)	787 (7)	787 (7)	787 (0)	--
(--)	--(--)					
2394	dbgShell	717 (6)	717 (6)	717 (6)	717 (0)	--
(--)	--(--)					
2396	khiCollection	1110 (7)	1110 (7)	1110 (7)	1110 (0)	--
(--)	--(--)					
2398	coreManager.x	788 (6)	788 (6)	788 (6)	788 (0)	--
(--)	--(--)					
2399	ovsdbServerCtl	--(--)	--(--)	--(--)	--(--)	--
(--)	--(--)					
2411	ssio	2916 (345)	2916 (345)	2916 (345)	2915 (0)	

Connectivity Problems

Use the following general tasks to isolate connectivity problems:

- Check physical connectivity. Verify if an alarm for link or port down exists.
- Check the link state by viewing the **show interfaces {gigabitEthernet|loopback|vlan}** command output.
- Use tools like ping or trace to verify if the connectivity issue is localized to an individual port or VLAN.
- Try to localize the affected range of ports and slot.

If you contact technical support staff to help troubleshoot connectivity problems, always provide source and destination IP pairs to facilitate in troubleshooting. Be sure to provide both working and non-working pairs for comparison.

```
Switch:1#show interfaces vlan
```

Vlan Basic									
VLAN ID	NAME	TYPE	MSTP INST_ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRFLD	ORIGIN	ACCESS
1	Default	byPort	0	none	N/A	N/A	0	CONFIG	GlobalRouter
2	VLAN-2	byPort	0	none	N/A	N/A	0	CONFIG	GlobalRouter
50	mcast_smlt_3	byPort	1	none	N/A	N/A	0	CONFIG	GlobalRouter
60	mcast_smlt_4	byPort	1	none	N/A	N/A	0	CONFIG	GlobalRouter

All 4 out of 4 Total Num of Vlans displayed

Vlan Port				
VLAN ID	PORT MEMBER	ACTIVE MEMBER	STATIC MEMBER	NOT_ALLOW MEMBER
1	1/5-1/8,1/10-1/48, 2/2-2/4	1/5-1/8,1/10-1/48, 2/2-2/4		

```

2

50  1/1,2/1/3-2/1/4  1/1,2/1/3-2/1/4

60  1/3-1/4          1/3-1/4

All 4 out of 4 Total Num of Port Entries displayed

=====
                        VLAN VRF Association
=====
VLAN  VRF
ID    NAME
-----
1     GlobalRouter
2     GlobalRouter
50    GlobalRouter
60    GlobalRouter
4086  GlobalRouter
4087  GlobalRouter

--More-- (q = quit)

```

Routing Table Problems

Routing table problems include but are not limited to:

- inactive routes
- unnecessary routes
- discard routes



Note

Only discard routes that belong to the static type protocol are supported. An inter-VRF discard route is not installed in the routing table of the destination VRF on the same switch.

- flapping links (links that go up and come down) that cause the routes to flap
- incorrect route tables
- invalid Address Resolution Protocol (ARP) cache that causes incorrect IP to MAC assignment
- problems with administrative distance or other parameters



Important

Do not restart a device to clear a problem. In restarting the device, you also clear the logs. Logs are vital and can help determine many problems.

Cable Connection Problems

You can usually trace port connection problems to a poor cable connection or to an improper connection of the port cables at either end of the link. To remedy such problems, make sure that the cable connections are secure and that the cables

connect to the correct ports at both ends of the link. If you use homemade cables, ensure that the cables are wired correctly.

1000BASE-T cables

1 Gb/s ports operate using Category 5 UTP cabling only. Category 5 UTP cable is a two-pair cable. To minimize crosstalk noise, maintain the twist ratio of the cable up to the point of termination; untwist at termination cannot exceed 0.5 in. (1.27 cm).

Pluggable optic cables

Cables for the optical transceivers vary depending on the specific device type.

For more information about the cable requirements for optical transceivers, see [Extreme Optics](#) website.

Alarm Database

The switch contains a local alarms mechanism. Local alarms are raised and cleared by applications running on the switch. View active alarms by using the **show alarm database** command in the CLI. Local alarms are an automatic mechanism run by the system that do not require additional configuration.

Check local alarms regularly to ensure no alarms require additional attention. The raising and clearing of local alarms also creates a log entry for each event. For more information about viewing logs, see [View Logs](#) on page 2557.

View the alarm database regularly to monitor alarm conditions, even if you do not observe a performance problem. Review the alarm messages to determine if the system performs as expected.

Not all alarm conditions indicate a problem so you must be familiar with expected behavior.

The alarm database shows the following alarm text:

```
CP1 [01/01/70 00:03:06.796] 0x00010844 00000000 Global Router HW WARNING
USB found in slot 1 has VendorId 05dc ProductId a01a and Manufacturer
Lexar and did not match supported devices
```

This alarm means that you have tried to insert an unsupported USB device into the USB slot. Only the USB device provided with your system can be inserted into the USB slot.

LED Indications of Problems

For information on LEDs on the chassis, see the hardware installation documentation.

Failure to Read Failed Configuration File

The device can fail to read and load a saved configuration file after it starts. This situation occurs if you enable the `factorydefaults boot configuration` flag. Configure the flag to `false`: `no boot config flags factorydefaults`.

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

```
Switch:1(config)# no boot config flags factorydefaults
```

No Web Management Interface Access to a Device

If the device and the PC that runs the web browser are in the same network, you can find that even though other applications, for example, Telnet, can access a particular switch, the web management interface cannot. This situation can occur if the web browser has a proxy server that resolves the `www` path and returns the reachable IP address to the browser. If no route exists from the proxy server to the device, the HTTP query does not reach the device, and does not receive a response.

To prevent this problem, ensure that if the web browser uses a proxy server, you specify a route from the proxy server to the device.

Configuration in EDM is not applied using CLI

In EDM, if you apply configuration that disrupts the connection to the web server, and you do not close the EDM session, you cannot apply the same configuration using CLI. This situation occurs if you, for example, delete the Segmented Management Instance, or the management IP address or the route used to access EDM, and you do not close the EDM session afterwards.

To prevent this issue, you must close the EDM session before you can apply the same configuration using CLI.

Connectivity Fault Management

Table 325: Connectivity Fault Management product support

Feature	Product	Release introduced
IEEE 802.1ag Connectivity Fault Management (CFM): <ul style="list-style-type: none"> • Layer 2 Ping • TraceRoute • TraceTree 	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

CFM Fundamentals

The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute. To support troubleshooting of the SPBM cloud, the switch supports a subset of CFM functionality. Configure CFM on all SPBM VLANs.

CFM is based on the IEEE 802.1ag standard.

IEEE 802.1ag Connectivity Fault Management (CFM) provides OAM tools for the service layer, which allows you to monitor and troubleshoot an end-to-end Ethernet service instance. CFM is the standard for Layer 2 ping, Layer 2 traceroute, and the end-to-end connectivity check of the Ethernet network.

The 802.1ag feature divides or separates a network into administrative domains called Maintenance Domains (MD). Each MD is further subdivided into logical groupings called Maintenance Associations (MA). A single MD can contain several MAs.

Each MA is defined by a set of Maintenance Points (MP). An MP is a demarcation point on an interface that participates in CFM within an MD. Two types of MP exist:

- Maintenance End Point (MEP)
- Maintenance Intermediate Point (MIP)

CFM supports three kinds of standard CFM messages: Continuity Check Message (CCM), Loopback Message (LBM), and Linktrace Message (LTM). Messages are sent between Maintenance Points (MP) in the system.

On the switch, CFM is implemented using the LBM and LTM features only to debug SPBM. CCM messages are not required or supported.

You can assign maintenance levels for each CFM SPBM MEP and MIP to each SPBM B-VLAN individually or you can assign maintenance levels and global MEPs for all SPBM VLANs.

Autogenerated CFM and Explicitly Configured CFM

The switch simplifies CFM configuration with autogenerated CFM. With autogenerated CFM, you use the **cfm spbm enable** command and the switch creates default MD, MA, MEPs, and MIPs for SPBM B-VLANs.

If you choose to configure CFM explicitly, you must configure an MD, MA, MEPs, and MIPs.

Autogenerated CFM

You can use autogenerated CFM at a global level to create a MEP and a MIP at a specified level for every SPBM B-VLAN on the chassis. If you use autogenerated CFM commands, you do not have to configure explicit MDs, MAs, MEPs, or MIPs, and associate them with multiple VLANs.

If you do not want to use autogenerated CFM commands, you can choose to configure explicit MDs, MAs, MEPs, and MIPs for SPBM B-VLANs. However, you cannot use both an autogenerated CFM configuration and an explicit CFM configuration together.



Note

Previous explicit CFM configurations of MDs, MAs, and MEPs on SPBM B-VLANs continue to be supported. However, if you want to enable the autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLANs. The switch only supports one type of MEP or MIP for each SPBM B-VLAN.

For information on autogenerated CFM configuration using CLI see [Configuring Autogenerated CFM on SPBM B-VLANs](#) on page 3808. For information on autogenerated CFM configuration using EDM see [Configure Autogenerated CFM on SPBM B-VLANs](#) on page 3830.

Explicitly configured CFM

If you choose to explicitly configure CFM, you must configure an MD, MA, MEPs, and MIPs. You can configure explicit CFM only on SPBM B-VLANs.

For explicit configuration information for CLI see [Configure Connectivity Fault Management Explicit Mode](#) on page 3810.

For explicit configuration information for EDM see [Configure Explicit CFM](#) on page 3831.

Using CFM

For SPBM B-VLANs, the autogenerated MEPs and MIPs respond to **12 ping**, **12 traceroute**, and **12 tracetest** in the same manner as the MEPs and MIPs created explicitly. The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. After you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **12 ping** and **12 traceroute** requests.

Maintenance Domain (MD)

A Maintenance Domain (MD) is the part of a network that is controlled by a single administrator. For example, a customer can engage the services of a service provider, who, in turn, can engage the services of several operators. In this scenario, there can be one MD associated with the customer, one MD associated with the service provider, and one MD associated with each of the operators.

You assign one of the following eight levels to the MD:

- 0–2 (operator levels)
- 3–4 (provider levels)
- 5–7 (customer levels)

The levels separate MDs from each other and provide different areas of functionality to different devices using the network. An MD is characterized by a level and an MD name (optional).

A single MD can contain several Maintenance Associations (MA).

Maintenance Association (MA)

An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance association End Points (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

The following figure shows MD level assignment in accordance with the 802.1ag standard. As shown in the figure, MIPs can be associated with MEPs. However, MIPs can also function independently of MEPs.

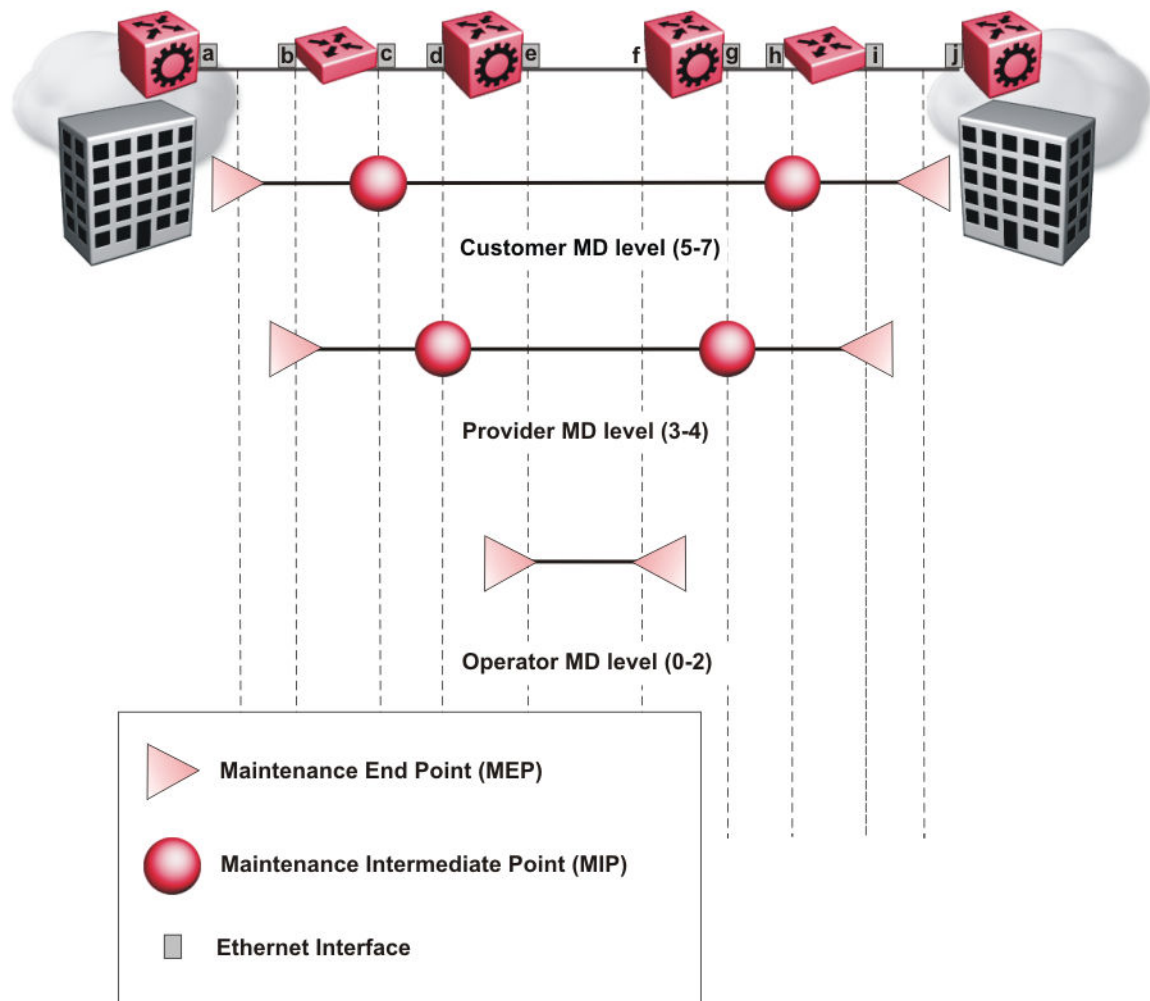


Figure 241: MD level assignment

Maintenance Association Endpoint (MEP)

A Maintenance Endpoint (MEP) represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA. MEP functionality can be divided into the following functions:

- Fault Detection
- Fault Verification
- Fault Isolation
- Fault Notification

Fault detection and notification are achieved through the use of Continuity Check Messages (CCM). CCM messages are not supported.

Fault Verification

Fault verification is achieved through the use of Loopback Messages (LBM). An LBM is a unicast message triggered by the operator issuing an operational command. LBM can be addressed to either a MEP or Maintenance Intermediate Point (MIP) but only a MEP can initiate an LBM. The destination MP can be addressed by its MAC address. The receiving MP responds with a Loopback Response (LBR). LBM can contain an arbitrary amount of data that can be used to diagnose faults as well as performance measurements. The receiving MP copies the data to the LBR.

LBM Message

The LBM packet is often compared to a ping. A MEP transmits the LBM packet. This packet can be addressed to another MEP or to the MAC address of the MP; in the case of SPBM, this is the SPBM system ID or its virtual SMLT MAC. Only the MP for which the packet is addressed responds with an LBR message.

- Provides “ICMP ping like” functionality natively at Layer-2.
- DA is the MAC address of the target.
- Includes a transaction identifier that allows the corresponding LBR to be identified when more than one LBM request is waiting for a response.
- Bridges forward the frame using the normal FDB rules.
- Only the target (MIP or MEP) responds.
- Initiator can choose the size and contents data portion of the LBM frame.
- Can be used to check the ability of the network to forward different sized frames.

L2 Ping

The **12 ping** command is a proprietary command that allows a user to trigger an LBM.

For B-VLANs, specify either the destination MAC address or node name.

You can specify an IP address as the destination address for the **12 ping** command.

You can test connectivity with Anycast IP Gateway I-SID.

Fault Isolation

Fault isolation is achieved through the use of Linktrace Messages (LTM). LTM is intercepted by all the MPs on the way to the destination MP. The switch supports two types of LTM.

The first type, the unicast LTM, can be addressed to either MEP or MIP MAC address. Each MP on the way decrements the TTL field in the LTM frame, sends Linktrace Reply (LTR), and forwards the original LTM to the destination. The LTM is forwarded until it reaches its destination or the TTL value is decremented to zero. LTR is a unicast message addressed to the originating MEP.

The second type, the proprietary LTM, is used to map the MAC addresses of the SPBM network; in this case the target MAC is not an MP, but rather a service instance identifier (I-SID).

Linktrace Message

Connectivity Fault Management offers linktrace messaging for fast fault detection. Use linktrace messages to verify the connectivity that you provide or use, and to debug systems.

Linktrace Message — Unicast

The linktrace message (LTM) is often compared to traceroute. A Maintenance Endpoint (MEP) transmits the LTM packet. This packet specifies the target MAC address of a Maintenance Point (MP), which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with a linktrace reply (LTR).

- Trace the path to any given MAC address.
- DA is unicast
- The LTM contains:
 - Time to live (TTL)
 - Transaction Identifier
 - Originator MAC address
 - Target MAC address
- CFM unaware entities forward the frame as is like any other data frame.
- MIP or MEP that is not on the path to the target discards the LTM and does not reply.
- MIP that is on the path to the target
 - Forwards the LTM after decrementing the TTL and replacing the SA with its own address.
 - Sends a reply (LTR) to the originator.
 - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- If the MIP or MEP is a target
 - Sends an LTR to the originator.
 - Identifies itself in the forwarded LTM and LTR by modifying TLV information.
- A MEP that is not the target but is on the path to the target performs the following actions:
 - Generates a reply as described in previous bullets.
 - Configures one of the flag fields in the reply to indicate that it is the terminal MEP.

Linktrace Message — Multicast

Use the multicast LTM to trace the multicast tree from any node on any I- SID using the nickname MAC address and the I-SID multicast address.

Specifying a multicast target address for an LTM allows for the tracing of the multicast tree corresponding to that destination address (DA). With a multicast target every node that is in the active topology for that multicast address responds with a Linktrace reply and also forwards the LTM frame along the multicast path. Missing Linktrace replies (LTRs) from the nodes in the path indicate the point of first failure.

Use this functionality to better troubleshoot I-SID multicast paths in a SPBM network.

L2 Traceroute

The **12 traceroute** command is a proprietary command that allows you to trigger an LTM message.

For B-VLANs, specify either the destination MAC address or node name.

This command provides a simpler command syntax than the standard LTM commands, which require you to specify the MD, MA, and MEP ID information. The **12 traceroute** command provides a trace equivalent at Layer 2 for use with nodes on the SPBM B-VLAN in the customer domain.

You can use CFM to troubleshoot networks and hosts that support the CFM protocol. After you configure CFM, CFM works in the network whether or not SPBM is in use.

You cannot use CFM to troubleshoot networks and hosts that do not support the CFM protocol, such as a customer domain that does not support CFM. Only devices that support the CFM protocol respond to **12 ping** and **12 traceroute** requests.



Note

While performing CFM Layer 2 traceroute between two BEBs using a transit BCB, the transit BCB hop is not seen, if the transit BCB has IS-IS adjacencies over FE I3core with both source BEB and destination BEB.

12 traceroute with IP address

You can specify an IP address as the destination address for the **12 traceroute** command.

The **12 traceroute** command converts Layer 3 IP information to an appropriate Layer 2 VLAN and MAC combination. The system can also target IP addresses that are not SPBM derived routes.

If you enable ECMP, **12 traceroute** runs internally for each of the VLAN paths returned, and displays a summary of the results. If you disable ECMP, the results display only one path.



Note

If you use the **12 traceroute ip-address** command on a DvR Leaf node, the output only shows DvR Controller IP addresses if the IP address or host route specified is unknown in the DvR domain.

L2 traceroute with Anycast IP Gateway

You can test connectivity with Anycast IP Gateway I-SIDs.

L2 Tracetree

The **12 tracetree** command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

L2 Tracetree-fan

The **12 tracetree-fan** command allows a user to trigger an LTM on the internal Fabric Area Network I-SID. This command allows the user to trace the FAN tree.

Maintenance Domain Intermediate Point (MIP)

MIPs do not initialize any CFM messages. MIPs passively receive CFM messages, process the messages received and respond back to the originating MEP. By responding to received CFM messages, MIPs can support discovery of hop-by-hop path among MEPs, allow connection failures to be isolated to smaller segments of the network to help discover location of faults along the paths. MIPs can be created independent of MEPs. MIP functionality can be summarized as:

- Respond to Loopback (ping) messages at the same level as itself and addressed to it.
- Respond to Linktrace (traceroute) messages.
- Forward Linktrace messages after decrementing the TTL.

Layer 2 Tracemroute

The **12tracemroute** command is a proprietary command that allows the user to trace the multicast tree for a certain multicast flow. The user specifies source, group, and service context (either VLAN or VRF) for the multicast flow to trace.

CFM sends a multicast LTM using an internal calculation to map the source, group, and context to the corresponding target address. The LTR comes from all leaves of the multicast tree for that flow, as well as transit nodes. The target MAC used in the LTM is a combination of the data I-SID and the nickname and the packet is sent on the appropriate SPBM B-VLAN. The user can see the generated multicast tree for that flow, which includes the data I-SID and nickname.

Nodal MPs

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. Each node (chassis) has a given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP. The Nodal B-VLAN MPs supports eight levels of CFM and you configure the Nodal B-VLAN MPs on a per B-VLAN basis. Virtual SMLT 10 MAC addresses are also able to respond for LTM and LBM.

Nodal B-VLAN MEPs

The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs. To support this behavior a MAC Entry is added to the FDB and a new CFM data-path table containing the B-VLAN and MP level are added to direct CFM frames to the CP as required.

Nodal B-VLAN MIPs

The Nodal MIP is associated with a B-VLAN. VLAN and level are sufficient to specify the Nodal MIP entity. The Nodal MIP MAC address is the SPBM system ID for the node on

which it resides. If the fastpath sends a message to the CP, the MIP responds if it is not the target and the MEP responds if it is the target.

Nodal B-VLAN MIPs with SMLT

When Nodal MEPs or MIPs are on SPBM B-VLANs the LTM code uses a unicast MAC DA. The LTM DA is the same as the target MAC address, which is the SPBM MAC address or the SMLT MAC address of the target node.

The switch supports SMLT interaction with SPBM. This is accomplished by using two B-VIDs into the core from each pair of SMLT terminating nodes. Both nodes advertise the Nodal BMAC into the core on both B-VIDS. In addition each node advertises the SMLT virtual BMAC on one of the two B-VLANs.

The Nodal MEP and MIP are expanded to respond to both the Nodal MAC address as well as the Virtual SMLT MAC address if both MACs are being advertised on its B-VLAN. In addition a source mode is added to the LTM and LBM command to use either the Nodal MAC or the SMLT virtual MAC address as the source MAC in the packet.

Configuration Considerations

When you configure CFM, be aware of the following configuration considerations.

General CFM

- A single switch has a limit of one MEP and one MIP on B-VLAN.
- The maintenance level for MEPs and MIPs on a given B-VID (in a network) must be configured to the same level for them to respond to a given CFM command.
- You can configure global CFM at only one MD level for each switch.
- All nodal MEPs and MIPs are restricted to SPBM B-VIDs.

Autogenerated CFM

- Autogenerated MEPs are not unique across the entire network unless you configure the global MEP ID on each switch to a different value. You must configure a unique MEP ID at a global level, for CFM.
- A single switch can have only one autogenerated MEP or MIP for each B-VLAN.

Explicit CFM

- Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to be supported. However, if you want to enable autogenerated CFM you must first remove the existing MEP and MIP on the SPBM B-VLAN.
- You can assign maintenance levels for each CFM SPBM MEP and MIP to each SPBM B-VLAN individually or you can assign maintenance levels and global MEPs for all SPBM VLANs by following the appropriate procedure:
 - [Assign a MEP/MIP Level to an SPBM B-VLAN](#) on page 3814
 - [Assign MEP/MIP Levels to SPBM B-VLANs Globally](#) on page 3815
 - [Configure CFM Nodal MEP](#) on page 3834

CFM Configuration Using CLI

This section provides procedures to configure and use Connectivity Fault Management (CFM) using Command Line Interface (CLI). The Shortest Path Bridging MAC (SPBM) network needs a mechanism to debug connectivity issues and to isolate faults. This is performed at Layer 2, not Layer 3. To support troubleshooting of the SPBM cloud, the switch supports a subset of CFM functionality.



Note

When you enable CFM in an SBPM network, as a best practice, enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

You can configure CFM using one of two modes: simplified or explicit. Both modes are described in the following sections, but the simplified mode should be used.



Note

If you enable the **cfm spbm enable** command, you cannot assign a MEP/MIP level to an individual SPBM B-VLAN or configure CFM MD maintenance levels individually.

Regardless of whether you have chosen to configure individually or globally, there is one MEP per SPBM B-VLAN and one MIP level per SPBM B-VLAN.

Autogenerated CFM

CFM provides two methods for configuration; autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure a MD, MA, and MEP ID to create a MEP.

For SPBM B-VLANs, you can use either autogenerated or explicitly configured CFM MEPs.

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly configured CFM MEPs.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function. However, if you want to enable the autogenerated commands you must first remove the existing MEP and MIP on the SPBM B-VLAN.

The switch only supports one MEP and one MIP, either autogenerated or explicitly configured, on the SPBM B-VLAN. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN.

Configuring Autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This eliminates the need to explicitly configure an MD, MA, and MEP ID, and to associate the MEP and MIP level to the SPBM B-VLAN.

When you enable this feature, the device creates a global MD (named spbm) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the level attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1.

The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.

**Important**

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs:

```
cfm spbm level <0-7>
```

You can change this level from the default of 4 either before or after the feature is enabled.

Only configure global CFM at one MD level for each chassis for each VLAN type.

3. Assign a global CFM MEP ID for all CFM SPBM MEPs:

```
cfm spbm mepid <1-8191>
```

4. Enable the autogenerated CFM for SPBM B-VLANs globally:

```
cfm spbm enable
```

5. (Optional) Configure the maintenance level for every CFM SPBM MEP and MP level on all the SPBM B-VLANs to the default:

```
default cfm spbm level
```

6. (Optional) Assign a global CFM MEP ID for all CFM SPBM MEPs to the default:

```
default cfm spbm mepid
```

7. (Optional) Disable the global CFM MEPs and MIPs:

```
no cfm spbm enable
```

8. Display the global CFM MEP configuration:

```
show cfm spbm
```

Example

Configure autogenerated CFM MEPs and MIPs:

```
Switch>enable  
Switch#configure terminal
```

```
Switch(config)#cfm spbm level 6
Switch(config)#cfm spbm mepid 4
Switch(config)#cfm spbm enable
Switch(config)#show cfm spbm
```

LEVEL	ADMIN	MEPID	MAC
6	enable	4	00:15:e8:b8:a3:df

Variable Definitions

The following table defines parameters for the **cfm spbm** command.

Variable	Value
<i>level</i> <0-7>	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
<i>mepid</i> <1-8191>	Specifies the global MEP ID within the range of 1-8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.
<i>enable</i>	Enables autogenerated CFM on all SPBM B-VLANs.

Configure Connectivity Fault Management Explicit Mode

In the explicit mode of configuring Connectivity Fault Management (CFM), you can manually configure a Maintenance Domain (MD), a Maintenance Association (MA), and a Maintenance Endpoint (MEP) and then associate the MEP to a B-VLAN and assign a MIP level to a B-VLAN.



Note

If you use autogenerated CFM, these steps are unnecessary.

Configure Connectivity Fault Management Maintenance Domain

Use this procedure to configure the Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Create the CFM MD:
cfm maintenance-domain WORD<0-22> [index <1-2147483647>] [maintenance-level <0-7>] [level <0-7>]
3. Display the CFM MD configuration:
show cfm maintenance-domain
4. Delete the CFM MD:
no cfm maintenance-domain WORD<0-22>

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#cfm maintenance-domain mdl index 99 maintenance-level 3
Switch:1(config)#show cfm maintenance-domain
=====
Maintenance Domain
=====
Domain Name          Domain Index    Level Domain Type
-----
mdl                  99              3      NODAL

Total number of Maintenance Domain entries: 1.
Switch:1(config)#no cfm maintenance-domain mdl
Switch:1(config)# show cfm maintenance-domain
=====
Maintenance Domain
=====
Domain Name          Domain Index    Level Domain Type
-----

Total number of Maintenance Domain entries: 0.
```

Variable Definitions

The following table defines parameters for the **cfm maintenance-domain** command.

Variable	Value
<i>WORD</i> <0-22>	Specifies the maintenance domain name.
<i>index</i> <1-2147483647>	Specifies a maintenance domain entry index.
<i>maintenance-level</i> <0-7>	Specifies the MD maintenance level when creating the MD. The default is 4.
<i>level</i> <0-7>	Modifies the MD maintenance level for an existing MD. The default is 4.

Configure Connectivity Fault Management Maintenance Association

Use this procedure to configure the Connectivity Fault Management (CFM) Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its domain. It can therefore represent a set of Maintenance Association Endpoints (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Create the CFM MA:


```
cfm maintenance-association WORD<0-22> WORD<0-22> [index <1-2147483647>]
```

3. Display the CFM MA configuration:

```
show cfm maintenance-association
```

4. Use the following command, if you want to delete the CFM MA:

```
no cfm maintenance-association WORD<0-22> WORD<0-22>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#cfm maintenance-association mdl mal index 98
Switch:1(config)#show cfm maintenance-association
=====
Maintenance Association Status
=====
Domain Name      Assn Name      Domain Idx  Assn Idx
-----
mdl              mal            1           98

Total number of Maintenance Association entries: 1.

=====
Maintenance Association config
=====
Domain Name      Assn Name
-----
mdl              mal

Total number of MA entries: 1.
```

Variable Definitions

The following table defines parameters for the **cfm maintenance-association** command.

Variable	Value
<i>WORD<0-22> WORD<0-22></i>	Creates the CFM MA. The first parameter, specifies the MD name. The second parameter, specifies the MA short name.
<i>index <1-2147483647></i>	Specifies a maintenance association entry index.

Configure Connectivity Fault Management Maintenance Endpoint

Use this procedure to configure the Connectivity Fault Management (CFM) Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create the CFM MEP:

```
cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable [state
<enable>]
```

3. Enable an existing CFM MEP:

```
cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
```

4. Disable an existing CFM MEP:

```
no cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191> enable
```

5. Display the CFM MEP configuration:

```
show cfm maintenance-endpoint
```

6. Delete an existing CFM MEP:

```
no cfm maintenance-endpoint WORD<0-22> WORD<0-22> <1-8191>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#cfm maintenance-endpoint md1 mal 1 state enable
Switch:1(config)#show cfm maintenance-endpoint
=====
Maintenance Endpoint Config
=====
DOMAIN          ASSOCIATION      MEP  ADMIN
NAME            NAME            ID
-----
md1             mal             1    enable

Total number of MEP entries: 1.

=====
Maintenance Endpoint Service
=====
DOMAIN_NAME      ASSN_NAME        MEP_ID TYPE    SERVICE_DESCRIPTION
-----
md1             mal             1      nodal    Vlan 1, Level 4

Total number of MEP entries: 1.
```

Variable Definitions

The following table defines parameters for the **cfm maintenance-endpoint** command.

Variable	Value
<i>WORD<0-22></i>	The first parameter, specifies the MD name.
<i>WORD<0-22></i>	The second parameter, specifies the MA short name.
<i><1-8191></i>	Specifies the MEP ID.
<i>enable</i>	Enables an existing MEP. Use this parameter with the no option to disable an existing MEP.
<i>state {enable disable}</i>	Enables or disables the MEP when creating the MEP. The default is disabled.

Assign a MEP/MIP Level to an SPBM B-VLAN

Use this procedure to assign a nodal MEP to an SPBM B-VLAN. The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and MIP functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Before You Begin

- You must configure a CFM MD, MA, and MEP.

Procedure

- Add nodal MEPs to the B-VLAN:

```
vlan nodal-mep <1-4059> WORD<0-22> WORD<0-22> <1-8191>
```
- Display the nodal MEP configuration:

```
show vlan nodal-mep <1-4059>
```
- Remove the nodal MEPs from the B-VLAN:

```
no vlan nodal-mep <1-4059> WORD<0-22> WORD<0-22> <1-8191>
```
- Add nodal MIP level to the B-VLAN:

```
vlan nodal-mip-level <1-4059> WORD<0-15>
```
- Display the nodal MIP level configuration:

```
show vlan nodal-mip-level [<1-4059>]
```
- Remove the nodal MIP level from the B-VLAN:

```
no vlan nodal-mip-level <1-4059> WORD<0-15>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#vlan nodal-mep 100 md1 ma1 2
Switch:1(config)#show vlan nodal-mep
=====
Vlan Nodal Mep
=====
VLAN_ID    DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
-----
100        spbm.100.6
200        spbm.200.6

Switch:1(config)#vlan nodal-mip 100 6
Switch:1(config)#show vlan nodal-mip
=====
Vlan Nodal Mip Level
=====
VLAN_ID    NODAL_MIP_LEVEL_LIST
-----
1
100        6
```

```
216
304
41000
1001
```

Variable Definitions

The following table defines parameters for the **vlan nodal-mep** command.

Variable	Value
<1-4059>	Specifies the VLAN ID.
WORD<0-22>	The first parameter, specifies the Maintenance Domain name.
WORD<0-22>	The second parameter, specifies the Maintenance Association name.
<1-8191>	Specifies the nodal MEPs to add to the VLAN.

The following table defines parameters for the **vlan nodal-mip-level** command.

Variable	Value
<1-4059>	Adds the nodal MIP level. Specifies the VLAN ID.
WORD<0-15>	Adds the nodal MIP level, which has up to eight levels, ranging from 0 to 7.

Assign MEP/MIP Levels to SPBM B-VLANs Globally



Note

If you enable the **cfm spbm enable** command, you cannot assign a MEP/MIP level to an individual SPBM B-VLAN or configure CFM MD maintenance levels individually.

About This Task

Enables the global CFM MEP and MIPs for all SPBM B-VLANs.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable simplified CFM configuration for SPBM VLANs:

```
cfm spbm enable
```
3. Enter the CFM SPBM level:

```
cfm spbm level <0-7>
```
4. Enter the CFM SPBM MEPID level:

```
cfm spbm mepid <1-8191>
```

Example

```
Switch:1(config)# cfm spbm level 7
Switch:1(config)# cfm spbm mepid 12
Switch:1(config)# cfm spbm enable
```

Variable Definitions

The following table defines parameters for the **simplified CFM** commands.

Variable	Value
<i>spbm level <0-7></i>	Configures the maintenance level for every CFM SPBM MEP and MIP level on all SPBM B-VLANs. The default is 4.
<i>mepid <1-8191></i>	Assigns a global MEP ID for all CFM SPBM MEPs. The default is 1.
<i>no cfm spbm enable</i>	Disables global configuration of CFM SPBM MEP and MIP levels on all SPBM B-VLANs.
<i>default cfm spbm level</i>	Returns maintenance level to default for all CFM SPBM MEP and MIP level on all SPBM B-VLANs.
<i>default cfm spbm mepid</i>	Returns MEP ID for all CFM SPBM MEPs to default.
<i>show cfm spbm</i>	Displays the global CFM MEP configuration for SPBM B-VLANs.

Trigger a Loopback Test (LBM)

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. This can be used to check the ability of the network to forward different sized frames.

Before You Begin

- You must have a MEP that is associated with a B-VLAN.

Procedure

- To enter User EXEC mode, log on to the switch.
- Trigger the loopback test:

```
loopback WORD<0-22> WORD<0-22> <1-8191>
<0x00:0x00:0x00:0x00:0x00:0x00> [burst-count <1-200>] [data-tlv-
size <0-400>] [frame-size <64-1500>] [interframe-interval
<msecs>] [priority <0-7>] [source-mode {nodal| smltVirtual}]
[testfill-pattern <all-zero|all-zero-crc|pseudo-random-bit-sequence|
pseudo-random-bit-sequence-crc>] [time-out <1-10>]
```

Example

```
Switch:1# loopback mdl 4001 13 00:14:0D:A2:B3:DF burst-count 10 priority
3 time-out 5
```

```
Result of LBM from mep: spbm.bvlan1000.8 to MAC address: 00:66:00:66:00:66 :
Sequence number of the first LBM is 150404162
The total number of LBMs sent out is 1
The number of LBRs received is 1
The number of LBRs lost is 0
The percentage of LBMs lost is 0.00%
The RTT Min is 15071 microseconds, Max is 15071 microseconds, Average is 15071.00 microseconds
The Standard Deviation of RTT is 0.00 microseconds
```

Variable Definitions

The following table defines parameters for the **loopback** command.

Variable	Value
<i>WORD</i> <0-22>	The first parameter, specifies the MD name.
<i>WORD</i> <0-22>	The second parameter, specifies the MA name.
<1-8191>	Specifies the MEP ID.
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the remote MAC address to reach the MEP/MIP.
<i>burst-count</i> <1-200>	Specifies the burst-count.
<i>data-tlv-size</i> <0-400>	Specifies the data TLV size.
<i>frame-size</i> <64-1500>	Specifies the frame-size. The default is 0.
<i>priority</i> <0-7>	Specifies the priority. The default is 7.
<i>source-mode</i> { <i>nodal</i> <i>smltVirtual</i> }} Note: Exception: <i>smltVirtual</i> is not supported on 4220 Series or 5320 Series.	Specifies the source mode: <ul style="list-style-type: none"> • <i>nodal</i> • <i>smltVirtual</i>—Use this value with B-VLANs only. The default is <i>nodal</i> .

Variable	Value
<code>testfill-pattern {all-zero all-zero-crc pseudo-random-bit-sequence pseudo-random-bit-sequence-crc}</code>	<p>Specifies the testfill pattern:</p> <ul style="list-style-type: none"> all-zero — null signal without cyclic redundancy check all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors.</p> <p>The default is 1: all-zero.</p>
<code>time-out <1-10></code>	Specifies the time-out interval in seconds. The default is 3.

Trigger Linktrace (LTM)

Use the following procedure to trigger a linktrace.

The Linktrace Message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Before You Begin

- You must have a MEP that is associated with a VLAN.

Procedure

- To enter User EXEC mode, log on to the switch.
- Trigger the linktrace:

```
linktrace WORD<0-22> WORD<0-22> <1-8191>
<0x00:0x00:0x00:0x00:0x00:0x00> [detail] [priority <0-7>] [source-mode
<nodal|smltVirtual>] [ttl-value <1-255>]
```

Example

```
Switch:1# linktrace mdl 4001 13 00:bb:00:00:14:00 priority 7
```

```
Please wait for LTM to complete or press any key to abort
```

```
Received LTRs:
```

```
SeqNum: 10575 MD: mdl MA:4001 MepId: 13 Priority: 7
```

```
-----
TTL SRC MAC FWDYES TERMMEP RELAY ACTION
---
```


63	00:bb:00:00:10:00	true	false	Fdb
62	00:bb:00:00:14:00	false	true	Hit

Variable Definitions

The following table defines parameters for the **linktrace** command.

Variable	Value
<i>WORD</i> <0-22>	The first parameter, specifies the MD name.
<i>WORD</i> <0-22>	The second parameter, specifies the MA name.
<1-8191>	Specifies the MEP ID.
<0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the target MAC address to reach the MEP.
<i>detail</i>	Displays linktrace result details.
<i>priority</i> <0-7>	Specifies the priority. The default is 7.
<i>source-mode</i> { <i>nodal</i> <i>smltVirtual</i> }} Note: Exception: <i>smltVirtual</i> is not supported on 4220 Series or 5320 Series.	Specifies the source mode: <ul style="list-style-type: none"> • <i>nodal</i> • <i>smltVirtual</i>—Use this value with B-VLANs only. The default is <i>nodal</i> .
<i>ttl-value</i> <1-255>	Specifies the Time-to-Live value. The default is 64.

Trigger a Layer 2 Ping

Before You Begin

- You must have a MEP that is associated with a VLAN.

About This Task

Use this procedure to trigger a Layer 2 ping, inside an SPBM cloud or network, which acts like native **ping**. This feature enables CFM to debug Layer 2. It can also help you debug IP shortcuts and the record for the shortcuts' ARP.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Trigger a Layer 2 ping:

- `l2 ping {vlan <1-4059> routernodename WORD<0-255> | vlan <1-4059> mac <0x00:0x00:0x00:0x00:0x00:0x00>} [burst-count <1-200>] [data-tlv-size <0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode <nodal|smltVirtual>] [testfill-pattern <all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>] [time-out <1-10>]`
- `l2 ping {ip-address WORD<0-255>} [burst-count <1-200>] [data-tlv-size <0-400>] [frame-size <64-1500>] [priority <0-7>] [source-mode <nodal|smltVirtual>] [testfill-pattern <all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>] [time-out <1-10>] [vrf WORD<1-16>]`
- `l2 ping {anycast-gw <1-15999999>} [burst-count<1-200> | data-tlv-size <0-400>| frame-size <64-1500> | priority <0-7> | source-mode <nodal|smltVirtual| testfill-pattern <all-zero|all-zero-crc|pseudo-random-bit-sequence|pseudo-random-bit-sequence-crc>] [time-out <1-10>] [vrf WORD<1-16>]`

Example

```
Switch:1#l2 ping vlan 2 mac 00.14.0d.bf.a3.df
Please wait for l2ping to complete or press any key to abort
----00:14:0d:bf:a3:df  L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 0 packets received,   100.00% packet loss

Switch:1#l2 ping vlan 2 routernodename MONTIO
Please wait for l2ping to complete or press any key to abort

----00:14:0d:a2:b3:df  L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 1 packets received,   0.00% packet loss
  round-trip (us)          min/max/ave/stdv =  26895/26895/26895.00/  0.00

Switch:1#l2 ping ip-address 192.0.2.10
Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 192.0.2.10, paths found 1, paths attempted 1
=====
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                                PKTS  PKTS  LOSS      MIN/MAX/AVE
(us)
=====
2  SHAMIM          (00:1a:8f:08:53:df)  1      0    100.00%  0/0/0.00
```

Variable Definitions

The following table defines parameters to configure the Layer 2 ping parameters.

Variable	Value
<pre>{vlan <1-4059> routernodename WORD<0-255> } (vlan <1-4059> mac <0x00:0x00:0x00:0x00:0x00:0x00> } {ip-address WORD<0-255> } {anycast-gw<1-15999999>}</pre>	<p>Specifies the destination for the Layer 2 ping:</p> <ul style="list-style-type: none"> • <1-4059> — Specifies the VLAN ID. • <1-15999999> — Specifies the Anycast IP Gateway service instance identifier (I-SID). • WORD<0-255> — Specifies the Router node name. • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address. • <A.B.C.D> — Specifies the IP address.
<i>burst-count</i> <1-200>	Specifies the burst count.
<i>data-tlv-size</i> <0-400>	Specifies the data TLV size. The default is 0.
<i>frame-size</i> <64-1500>]	Specifies the frame size. The default is 0.
<pre>testfill-pattern <all-zero all-zero-crc pseudo-random- bit-sequence pseudo-random- bit-sequence-crc></pre>	<p>Specifies the testfill pattern:</p> <ul style="list-style-type: none"> • all-zero — null signal without cyclic redundancy check • all-zero-crc — null signal with cyclic redundancy check with 32-bit polynomial • pseudo-random-bit-sequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudo-random-bit-sequence-crc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors. The default is all-zero.</p>
<i>priority</i> <0-7>	Specifies the priority. The default is 7.
<i>time-out</i> <1-10>	Specifies the interval in seconds. The default is 3.
<pre>source-mode{nodal smltVirtual}} </pre> <p>Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.</p>	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • 1: nodal • 2: smltVirtual—Use this value with B-VLANs only. <p>The default is 1: nodal.</p>
<i>vrf</i> WORD<1-16>	Specifies the VRF name.

Trigger a Layer 2 Traceroute

Use this procedure to trigger a Layer 2 traceroute, which acts like native **traceroute**. This feature enables CFM to debug Layer 2 in an SPBM cloud or network. It can determine the path used by IS—IS to get from one MEP to another, by showing all

the hops between. Therefore, it can show where connectivity is lost. It can also work for IP shortcuts.



Important

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

- For B-VLANs, you do not have to trigger an **l2ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.

linktrace traces the path up to the closest device to that MAC address that supports CFM.

Before You Begin

- You must have a MEP that is associated with a VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Trigger a Layer 2 traceroute:
 - `l2 traceroute {ip-address WORD<0-255>} [ttl <1-255>] [vrf WORD<1-16>]`
 - `l2 traceroute {<vlan <1-4059> routernodename WORD<0-255> | <vlan <1-4059> mac <0x00:0x00:0x00:0x00:0x00:0x00>} [priority <0-7>] [source-mode <nodal|smltVirtual>] [ttl <1-255>]`
 - `l2 traceroute {anycast-gw <1-15999999> [priority <0-7>] [source-mode <nodal|smltVirtual>] [ttl <1-255>]`

Examples

```
Switch:1#l2 traceroute vlan 2 routernodename Switch-MONTIO
Please wait for l2traceroute to complete or press any key to abort

l2traceroute to Switch-MONTIO (00:14:0d:a2:b3:df),  vlan 2
0   Switch-PETER4          (00:15:9b:11:33:df)
1   Switch-MONTIO          (00:14:0d:a2:b3:df)

Switch:1#l2 traceroute ip-address 192.0.2.10
Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 192.0.2.10, paths found 1
=====
Switch-SHAMIM (00:1a:8f:08:53:df),  vlan 2
0   Switch-PETER4          (00:15:9b:11:33:df)
1   Switch-MONTIO          (00:14:0d:a2:b3:df)
```

Variable Definitions

The following table defines parameters for the **l2 traceroute** command.

Variable	Value
<pre>{vlan <1-4059> routernodename WORD<0-255> } (vlan <1-4059> mac <0x00:0x00:0x00:0x00:0x00:0x00> } {ip-address WORD<0-255> } {anycast-gw<1-15999999>}</pre>	<p>Specifies the destination for the Layer 2 traceroute:</p> <ul style="list-style-type: none"> • <1-4059> — Specifies the VLAN ID • <1-15999999> — Specifies the Anycast IP Gateway service instance identifier (I-SID). • WORD<0-255> — Specifies the Router Node Name • <XX:XX:XX:XX:XX:XX> — Specifies the MAC address • WORD<0-255> — Specifies the IP address
<code>ttn-value <1-255></code>	Specifies the TTL value. The default is 64.
<code>priority <0-7></code>	Specifies the priority. The default is 7.
<pre>source-mode{nodal smltVirtual}} </pre> <p>Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.</p>	<p>Specifies the source mode:</p> <ul style="list-style-type: none"> • nodal • smltVirtual—Use this value with B-VLANs only. <p>The default is nodal.</p>
<code>vrf WORD<1-16></code>	Specifies the VRF name.

Trigger a Layer 2 Tracetree

Use this procedure to trigger a Layer 2 tracetree. Layer 2 tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Trigger a Layer 2 tracetree:

```
12 tracetree {<1-4059> <1-16777215> [routernodename WORD<0-255>
| <1-4059> <1-16777215>] [mac <0x00:0x00:0x00:0x00:0x00:0x00>]}
[priority <0-7>] [source-mode <nodal|smltVirtual>] [ttn-value <1-255>]
```

Example

```
Switch:1#12 tracetree 500 1
Switch:1# 12 tracetree 500 1

Please wait for 12tracetree to complete or press any key to abort

12tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10
```

```

hops 64
1   Switch-PETER4      00:15:9b:11:33:df -> Switch-MONTI0      00:14:0d:a2:b3:df
2   Switch-MONTI0      00:14:0d:a2:b3:df -> Switch-LEE2        00:15:e8:b8:a3:df

```

Variable Definitions

The following table defines parameters for the **12 tracetree** command.

Variable	Value
<pre>{ <1-4059><1-16777215> routernodename WORD<0-255> <1-4059><1-16777215> mac <0x00:0x00:0x00:0x00:0x00:0x00>}</pre>	<ul style="list-style-type: none"> • <1-4059> — Specifies the VLAN ID. • <1-16777215> — Specifies the I-SID. • WORD<0-255> — Specifies the Router Node Name. • <0x00:0x00:0x00:0x00:0x00:0x00> — Specifies the MAC address.
<code>ttl-value <1-255></code>	Specifies the TTL value. The default is 64.
<code>priority <0-7></code>	Specifies the priority value. The default is 7.
<pre>source-mode<nodal smltVirtual></pre> <p>Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.</p>	Specifies the source mode. The default is nodal.

Trigger a Layer 2 Tracetree-fan

About This Task

Use this procedure to trigger a Layer 2 tracetree-fan from the nickname server to make sure that all the nodes towards the nickname client support the FAN protocol. Layer 2 tracetree-fan allows a user to trigger an LTM on the internal Fabric Area Network (FAN) I-SID. This command allows the user to trace the FAN tree.

Procedure

Trigger a Layer 2 tracetree-fan:

```
12 tracetree-fan [mac <0x00:0x00:0x00:0x00:0x00:0x00>] [priority <0-7>]
[routernodename WORD<0-255>] [ttl-value <1-255>]
```

Example

```

Switch:1# 12 tracetree-fan

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to b1:ad:aa:41:b0:84, vlan 4051 i-sid 16777001 nickname 0.00.00 hops 64

1   Switch-PETER4      b0:ad:aa:41:b0:84 -> Switch-MONTI0      b0:ad:aa:41:48:84
1   Switch-PETER4      b0:ad:aa:41:b0:84 -> Switch-MONTI1      b0:ad:aa:42:88:84
2   Switch-MONTI0      b0:ad:aa:41:48:84 -> Switch-LEE2        b0:ad:aa:43:3c:84

```

Trigger a Layer 2 Tracemroute

Use this procedure to debug the IP Multicast over Fabric Connect stream path using **12 tracemroute** on the VLAN (Layer 2) or the VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.



Note

The VLAN option is only valid for a VLAN that has an I-SID configured and IGMP snooping enabled.

Before You Begin

- On the source and destination nodes, you must configure an autogenerated or an explicit CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Trigger a Layer 2 tracemroute on the VLAN:

```
12 tracemroute source <A.B.C.D> group <A.B.C.D> vlan <1-4059>[priority
<0-7>] [ttl-value <1-255>]
```



Note

For the preceding command, if you do not specify a VLAN, **12 tracemroute** uses the global default VRF.

Wait for the l2 tracemroute to complete or press any key to abort.

3. Trigger a Layer 2 tracemroute on the VRF:

```
12 tracemroute source <A.B.C.D> group <A.B.C.D> vrf WORD<1-16>
[priority <0-7>] [ttl-value <1-255>]
```



Note

For the preceding command, if you do not specify a VRF, **12 tracemroute** uses the global default VRF.

Wait for the l2 tracemroute to complete or press any key to abort.

Examples

The following is a sample output for a Layer 2 tracemroute on a VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#12 tracemroute source 192.0.2.81 233.252.0.1 vlan 201

Please wait for 12 tracemroute to complete or press any key to abort.

Source 192.0.2.81
```

```

Group: 233.252.0.1

VLAN:201

BMAC: 03:00:03:f4:24:01

B-VLAN: 10

I-SID: 16000001

=====
1 PETER4 00:03:00:00:00:00 -> LEE1 00:14:0d:bf:a3:df
2 LEE1 00:14:0d:bf:a3:df -> LEE2 00:15:e8:b8:a3:df

```

The following is a sample output for a Layer 2 tracemroute on a VRF:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#12 tracemroute source 192.0.2.10 group 233.252.0.1 vrf red

Please wait for l2 tracemroute to complete or press any key to abort.

Source 192.0.2.10

Group: 233.252.0.1

VRF: redID 1

BMAC: 03:00:04:f4:24:01

B-VLAN: 20

I-SID: 16000001

=====
1 PETER4 00:03:00:00:00:00 -> LEE1 00:14:0d:bf:a3:df
2 LEE1 00:14:0d:bf:a3:df -> LEE2 00:15:e8:b8:a3:df

```

Variable Definitions

The following table defines parameters for the **12 tracemroute** command.

Variable	Value
<i>source</i> <A.B.C.D>	Specifies the source IP address.
<i>group</i> <A.B.C.D>	Specifies the IP address of the multicast group.
<i>vlan</i> <1-4084>	Specifies the VLAN value.
<i>vrf</i> WORD<1-16>	Specifies the VRF name. If you do not specify a VRF name, then the results are shown for the flow in the Global Router (default) context.
<i>priority</i> <0-7>	Specifies the priority value.
<i>ttl</i> <1-255>	Specifies the time-to-live (TTL) for the trace packet, which is how many hops the trace packet takes before it is dropped.

Use Trace CFM to Diagnose Problems

Use the following procedure to display trace information for CFM.

About This Task

Use trace to observe the status of a software module at a certain time.

For example, if you notice a CPU utilization issue (generally a sustained spike above 90%) perform a trace of the control plane activity.

Use the **trace level 120 <0-4>** command to trace CFM module information, including CLI, instrumentation, show config, and platform dependent code. The CFM module ID is 120.

Use the **trace cfm level <0-4>** command to trace platform independent code and CFM protocol code.



Caution

Risk of traffic loss

Using the trace tool inappropriately can cause CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Clear the trace:

```
clear trace
```

3. Begin the trace operation:

```
trace cfm level <0-4>
```

Wait approximately 30 seconds, and then stop trace.

4. Stop tracing:

```
trace shutdown
```

5. View the trace results:

```
show trace cfm
```

6. Begin the trace operation for the CFM module:

```
trace level 120 <0-4>
```

Wait approximately 30 seconds, and then stop trace.

7. View trace results:

```
trace screen enable
```



Important

If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

8. Save the trace file to the Compact Flash card for retrieval.

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

9. Search trace results for a specific string value, for example, the word error:

```
trace grep [WORD<0-128>]
```

If you use this command and do not specify a string value, you clear the results of a previous search.

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)# clear trace
Switch:1(config)# trace cfm level 3
Switch:1(config)# trace shutdown
Switch:1(config)# show trace cfm
=====
                        CFM Tracing Info
=====
Status      : Enabled
Level       : VERBOSE
Switch:1(config)#trace level 120 3
Switch:1(config)# save trace
Switch:1(config)# trace grep error
Switch:1(config)#trace grep 00-1A-4B-8A-FB-6B
```

Variable Definitions

The following table defines parameters for the **trace** command.

Variable	Value
<i>cfm level</i> [<0-4>]	Starts the trace by specifying the level. <ul style="list-style-type: none"> <0-4> specifies the trace level from 0-4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
<i>filter</i>	Configures a filter trace for a file or module.
<i>flags</i>	Configures trace flags for IS-IS or OSPF.
<i>grep</i> [WORD<0-128>]	Searches trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.

Variable	Value
<code>level [<Module_ID>] [<0-4>]</code>	<p>Starts the trace by specifying the module ID and level.</p> <ul style="list-style-type: none"> • <code><Module_ID></code> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see <i>Fabric Engine and VOSS Feature Matrix</i>. To see the list of module IDs, use the show trace modid-list command or CLI command completion Help. • <code><0-4></code> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
<code>route-map</code>	Enables or disables the trace route-map. The values are on and off.
<code>screen {disable enable}</code>	Enables the display of trace output to the screen.
<code>shutdown</code>	Stops the trace operation.
<code>spbm isis level [<0-4>]</code>	<p>Starts the trace by specifying the level.</p> <ul style="list-style-type: none"> • <code><0-4></code> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose. <p>The default is 1, very terse.</p>

The following table defines parameters for the **save trace** command.

Variable	Value
<code>file WORD<1-99></code>	<p>Specifies the file name in one of the following formats:</p> <ul style="list-style-type: none"> • <code>a.b.c.d:<file></code> • <code>x:x:x:x:x:x:x <file></code> • <code>/intflash/<file></code> • <code>/extflash/<file></code> • <code>/usb/<file></code> • <code>/mnt/intflash/ <file></code> • <code>/mnt/extflash/ <file></code> <p>/mnt/intflash is the internal flash of the CPU. /mnt/extflash is the external flash of the CPU.</p>

CFM Configuration Using EDM

This section provides procedures to configure Connectivity Fault Management (CFM) using Enterprise Device Manager (EDM).



Note

When you enable CFM in an SPBM network, as a best practice, enable CFM on the Backbone Edge Bridges (BEB) and on all Backbone Core Bridges (BCB). If you do not enable CFM on a particular node, you cannot obtain CFM debug information from that node.

Autogenerated CFM

CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. You must choose one or the other. Use the procedures in this section to configure autogenerated MEPs that eliminate the need to configure an MD, MA, and MEP ID to create a MEP.

Previous explicit CFM configurations of MDs, MAs and MEPs on SPBM B-VLANs continue to function. However, if you want to enable the autogenerated commands, you must first remove the existing MEP and MIP on the SPBM B-VLAN. The switch only supports one MEP or MIP on the SPBM B-VLAN, either explicitly configured or autogenerated.

Configure Autogenerated CFM on SPBM B-VLANs

Use this procedure to configure the autogenerated CFM MEP and MIP level for every SPBM B-VLAN on the chassis. This configuration eliminates the need to explicitly configure an MD, MA, and MEP ID and to associate the MEP and MIP level to the SPBM B-VLAN.

About This Task

When you enable this feature, the device creates a global MD (named `spbm`) for all the SPBM Nodal MEPs. This MD has a default maintenance level of 4, which you can change with the `level` attribute. All the MEPs that the device creates use the MEP ID configured under the global context, which has a default value of 1. The nodal MEPs are automatically associated with the SPBM B-VLANs configured. The MIP level maps to the global level. When you enable the feature, the device automatically associates the MIP level with the SPBM B-VLANs configured. The feature is disabled by default.



Important

CFM supports one MEP or MIP for each SPBM B-VLAN only. This means that if you want to use these autogenerated MEPs, you cannot use your existing CFM configuration. You must first remove the existing MEP or MIP on the SPBM B-VLAN. If you want to continue configuring MEPs manually, skip this procedure.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **CFM**.

3. Select the **Global** tab.
4. Select **enable** next to **SpbmAdminState**.
5. Select **Apply**.

Global Field Descriptions

Use the data in the following table to configure the global MEP and MIP parameters.

Name	Description
SpbmAdminState	Enables or disables autogenerated CFM for B-VLANs. The default is disable.
SpbmLevel	Specifies the global SPBM CFM maintenance level for the chassis within the range of 0 to 7. The default is 4. Only configure global CFM at one MD level for each chassis for each VLAN type.
SpbmMepld	Specifies the global MEP ID within the range of 1 to 8191. Select a unique ID for each switch to ensure that the MEPs are unique across the network. The default is 1.

Configure Explicit CFM

For SPBM B-VLANs, CFM provides two methods for creating MEPs: autogenerated and explicit. You cannot use both. Use the procedures in this section to configure MEPs explicitly.

If you want to create autogenerated CFM MEPs that eliminate the need to configure an MD, MA, and MEP ID, see the procedures in [Autogenerated CFM](#) on page 3830.



Note

The CFM show commands that display MD, MA, and MEP information work for both autogenerated and explicitly-configured CFM MEPs.

Configure CFM MD

About This Task

Use this procedure to configure a Connectivity Fault Management (CFM) Maintenance Domain (MD). An MD is the part of a network that is controlled by a single administrator. A single MD can contain several Maintenance Associations (MA).

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **CFM**.
3. Select the **MD** tab.
4. Select **Insert**.
5. In the fields provided, specify an index value, name, and level for the MD.
6. Select **Insert**.

MD Field Descriptions

Use the data in the following table to use the **MD** tab.

Name	Description
Index	Specifies a maintenance domain entry index.
Name	Specifies the MD name.
NumOfMa	Indicates the number of MAs that belong to this maintenance domain.
Level	Specifies the MD maintenance level. The default is 4.
NumOfMip	Indicates the number of MIPs that belong to this maintenance domain
Type	Indicates the type of domain.

Configure CFM MA

Before You Begin

- You must configure a CFM MD.

About This Task

Use this procedure to configure a CFM Maintenance Association (MA). An MA represents a logical grouping of monitored entities within its Domain. It can therefore represent a set of Maintenance Endpoints (MEPs), each configured with the same Maintenance Association ID (MAID) and MD Level, established to verify the integrity of a single service instance.

Procedure

- In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
- Select **CFM**.
- Select the **MD** tab.
- Highlight an existing MD, and then select **MaintenanceAssociation**.
- In the **MA** tab, select **Insert**.
- In the fields provided, specify an index value and name for the MA.
- Select **Insert**.

MA Field Descriptions

Use the data in the following table to use the **MA** tab.

Name	Description
DomainIndex	Specifies the maintenance domain entry index.
AssociationIndex	Specifies a maintenance association entry index.

Name	Description
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
NumOfMep	Indicates the number of MEPs that belong to this maintenance association.

Configure CFM MEP

Use this procedure to configure the CFM Maintenance Endpoint (MEP). A MEP represents a managed CFM entity, associated with a specific Domain Service Access Point (DoSAP) of a service instance, which can generate and receive CFM Protocol Data Units (PDU) and track any responses. A MEP is created by MEP ID under the context of an MA.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **CFM**.
3. Select the **MD** tab.
4. Highlight an existing MD, and then select **MaintenanceAssociation**.
5. In the **MA** tab, highlight an existing MA, and then select **MaintenanceEndpoint**.
6. Select **Insert**.
7. In the fields provided, specify the ID and the administrative state of the MEP.
8. Select **Insert**.

MEP Field Descriptions

Use the data in the following table to use the **MEP** tab.

Name	Description
DomainIndex	Specifies the MD index.
AssociationIndex	Specifies the MA index.
Id	Specifies the MEP ID.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
AdminState	Specifies the administrative state of the MEP. The default is disable.

Name	Description
MepType	Specifies the MEP type: <ul style="list-style-type: none"> • trunk • sg • endpt • vlan • port • endptClient • nodal • remotetrunk • remotesg • remoteendpt • remoteVlan • remotePort • remoteEndptClient
ServiceDescription	Specifies the service to which this MEP is assigned.

Configure CFM Nodal MEP

Before You Begin

- You must configure a CFM MD, MA, and MEP.

About This Task

Use this procedure to configure the CFM nodal Maintenance Endpoint (MEP). The Nodal MEP provides traceability and troubleshooting at the system level for a given B-VLAN. The Nodal B-VLAN MEPs created on the CP and function as if they are connected to the virtual interface of the given B-VLAN. Because of this they are supported for both port and MLT based B-VLANs.

Nodal MPs provide both MEP and Maintenance Intermediate Point (MIP) functionality for SPBM deployments. Nodal MPs are associated with a B-VLAN and are VLAN encapsulated packets. Each node (chassis) has a given MAC address and communicates with other nodes. The SPBM instance MAC address is used as the MAC address of the Nodal MP.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. Select a VLAN with a type of spbm-bvlan.
5. Select **Nodal**.
6. In **NodalMepList**, specify the nodal MEPs to add to the VLAN.
7. Select **Apply**.

Nodal MEP/MIP Field Descriptions

Use the data in the following table to use the **Nodal MEP/MIP** tab.

Name	Description
NodalMepList	Specifies the nodal MEPs to add to the VLAN, in the format <mdName.maName.mepId>, for example md10.ma20.30.
NumOfNodalMep	Indicates the number of nodal MEPs assigned to this VLAN.
NodalMipLevelList	Specifies a MIP level list.
NumOfNodalMipLevel	Indicates the number of nodal MIP levels assigned to this VLAN that allows MIP functionality to be enabled on a per level per VLAN basis.

Configure Layer 2 Ping

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

About This Task

Use this procedure to configure a Layer 2 ping inside an SPBM cloud or network. This feature enables CFM to debug Layer 2. It can also help you debug IP shortcuts and the record for the shortcuts' ARP.

Procedure

- In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
- Select **L2Ping/L2Trace Route**.
- From the **L2Ping** tab, configure the Layer 2 ping properties.
- To initiate a Layer 2 ping, highlight an entry, and then select **Start**.
- To update a Layer 2 ping, select **Refresh**.
- To stop the Layer 2 ping, select **Stop**.

L2Ping Field Descriptions

Use the data in the following table to use the **L2Ping** tab.

Name	Description
VlanId	Identifies the backbone VLAN.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Indicates whether the host name is (true) or is not (false) used for L2Ping transmission.

Name	Description
Messages	Specifies the number of L2Ping messages to be transmitted. The default is 1.
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Ping messages. • abort: the service aborted or is about to abort the L2Ping messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the L2Ping Messages will be (or have been) sent. • false: the L2Ping Messages will not be sent. <p>The default is true.</p>
Priority	<p>Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame.</p> <p>The default is 7.</p>
TimeoutInt	Specifies the interval to wait for an L2Ping time-out. The default value is 3 seconds.
TestPattern	<p>Specifies the test pattern to use in the L2Ping PDU:</p> <ul style="list-style-type: none"> • allZero: null signal without cyclic redundancy check • allZeroCrc: null signal with cyclic redundancy check with 32-bit polynomial • pseudoRandomBitSequence: pseudo-random-bit-sequence without cyclic redundancy check • pseudoRandomBitSequenceCrc: pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors. The default value is allZero.</p>
DataSize	Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The default is 0.

Name	Description
FrameSize	Specifies the frame size. If the frame size is specified then the data size is internally calculated and the calculated data size is included in the data TLV. The default is 0.
SourceMode Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.	Specifies the source mode of the transmit loopback service: <ul style="list-style-type: none"> • nodal • smltVirtual — Use the smltVirtual option with B-VLANs only. The default is nodal.
SeqNumber	The transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Result	Displays the Layer 2 Ping result.

Initiate a Layer 2 Traceroute

Use this procedure to trigger a Layer 2 traceroute. This feature enables CFM to debug Layer 2 in an SPBM cloud or network. It can determine the path used by IS—IS to get from one MEP to another, by showing all the hops between. Therefore, it can show where connectivity is lost. It can also work for IP shortcuts.

If you configure **IsTraceTree** to false then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true then EDM performs TraceTree on the multicast tree.

For more information on configuring tracetree, see [Configure a Layer 2 Tracetree](#) on page 3853.



Important

To trace a route to a MAC address, the MAC address must be in the VLAN FDB table.

For B-VLANs, you do not have to trigger an **L2Ping** to learn the MAC address because IS-IS populates the MAC addresses in the FDB table.

Linktrace traces the path up to the closest device to that MAC address that supports CFM.

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 Traceroute/TraceTree** tab.
4. To start the traceroute, highlight an entry, and then select **Start**.
5. To update the traceroute, select **Refresh**.

6. To stop the traceroute, select **Stop**.

L2 Traceroute/TraceTree Field Descriptions

Use the data in the following table to use the **L2 Traceroute/TraceTree** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the Backbone VLAN (B-VLAN).
Priority	Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Specifies whether the host name is (true) or is not (false) used for the L2Trace transmission.
Isid	Specifies the Service Instance Identifier (I-SID).
IsTraceTree	<p>Specifies whether the multicast tree or unicast path is traced:</p> <ul style="list-style-type: none"> If you configure IsTraceTree to false then EDM performs Traceroute on the unicast path. If you configure IsTraceTree to true then EDM performs TraceTree on the multicast tree.
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> ready: the service is available. transmit: the service is transmitting, or about to transmit, the L2Trace messages. abort: the service aborted or is about to abort the L2Trace messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> true: the L2Trace messages will be (or have been) sent. false: the L2Trace messages will not be sent. <p>The default is true.</p>

Name	Description
Ttl	<p>Specifies the number of hops remaining to this L2Trace.</p> <p>This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the time-to-live (TTL) field in the L2Trace is defined by the originating MEP.</p> <p>The default value is 64.</p>
SourceMode Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.	<p>Specifies the source mode:</p> <p>The default is nodal.</p>
SeqNumber	<p>Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.</p>
Flag	<p>L2Trace result flag that indicates L2Trace status or error code:</p> <ul style="list-style-type: none"> • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace • l2TraceLeafNode (9): On a leaf node in the I-SID tree • l2TraceNotInTree (10): Not in the I-SID tree • l2TraceSmItNotPrimary (11): Requested SMLT source from non-primary node

View Layer 2 Traceroute Results

Use this procedure to view Layer 2 traceroute results. This feature enables CFM to debug Layer 2. It can also help you debug ARP problems by providing the ability to troubleshoot next hop ARP records.

About This Task

You can display Layer 2 tracetable results to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID. For more information, see [View Layer 2 Tracetable Results](#) on page 3856.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2Traceroute/TraceTree** tab.
4. Select **Refresh** to update the results.
5. To view the traceroute results, highlight an entry, and then select **Result**.

L2 Traceroute/Tracetable Result Field Descriptions

Use the data in the following table to use the **L2 Traceroute/Tracetable Result** tab.

Name	Description
VlanId	A value that uniquely identifies the Backbone VLAN (B-VLAN).
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which L2Trace's response of the L2Trace is going to be returned. The default is 0.
Hop	The number of hops away from L2Trace initiator.
ReceiveOrder	An index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Time-to-Live (TTL) field value for a returned L2Trace response.
SrcMac	MAC address of the MP that responds to the L2Trace request for this L2TraceReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	The host name of the node that forwarded the L2Trace to the responding node.

Configure Layer 2 IP Ping

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.

- If you want to run a Layer 2 IP Ping for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in [Select and Launch a VRF Context View](#) on page 4100.

About This Task

Use this procedure to configure Layer 2 IP ping

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 IP Ping** tab.
4. To add a new entry:
 - a. Select **Insert**.
 - b. Specify the destination IP address.
 - c. Select **Insert**.
5. To start the Layer 2 IP ping, highlight an entry, and then select **Start**.
6. To stop the Layer 2 IP ping, select **Stop**.
7. To update the Layer 2 IP ping, select **Refresh**.

L2 IP Ping Field Descriptions

Use the data in the following table to use the **L2 IP Ping** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP Address (only IPv4 is supported).
IpAddr	Specifies the destination IP Address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Messages	Specifies the number of L2IpPing messages to be transmitted per MAC/VLAN pair. Range is 1–200. The default is 1.

Name	Description
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2IpPing messages. • abort: the service is aborted or about to abort the L2IpPing messages. <p>This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time. The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: L2IpPing Messages will be or have been sent. • false: L2IpPing Messages will not be sent. <p>The default is true.</p>
TimeoutInt	<p>Specifies the interval to wait for an L2IpPing time-out with a range of 1–10 seconds with a default value of 3 seconds.</p>
TestPattern	<p>Specifies the test pattern to use in the L2IPPing PDU:</p> <ul style="list-style-type: none"> • allZero — null signal without cyclic redundancy check • allZeroCrc — null signal with cyclic redundancy check with 32-bit polynomial • pseudoRandomBitSequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudoRandomBitSequenceCrc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors. The default value is allZero.</p>
DataSize	<p>Specifies an arbitrary amount of data to be included in the data TLV, if the data size is selected to be sent. The range is 0–400. The default is 0.</p>
PathsFound	<p>Specifies the number of paths found to execute the command. The default is 0.</p>

*View Layer 2 IP Ping Results***About This Task**

Use this procedure to view Layer 2 IP ping results.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 IP Ping** tab.
4. To view the Layer 2 IP ping results, highlight an entry, and then select **Result**.

L2 IP Ping Result Field Descriptions

Use the data in the following table to use the **L2 IP Ping Result** tab.

Name	Description
IpAddrType	The address type of the destination IP Address.
IpAddr	Destination IP Address.
SendOrder	Specifies the order that sessions were sent. It is an index to distinguish among multiple L2Ping sessions. This value is assigned sequentially from 1. It correlates to the number of paths found.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN ID found from the Layer 3 lookup and used for transmission.
DestMacAddress	An indication of the target MAC Address transmitted.
PortNum	Either the value '0', or the port number of the port used for the L2 IP ping.
DestHostName	The host name of the responding node.
Size	The number of bytes of data sent.
PktsTx	Number of Packets transmitted for this VLAN/MAC.
PktsRx	Number of Packets received for this VLAN/MAC.
PercentLossWhole	Percentage of packet loss for this VLAN/MAC.
PercentLossFract	Percentage of packet loss for this VLAN/MAC.
MinRoundTrip	Minimum time for round-trip for this VLAN/MAC in us.
MaxRoundTrip	Maximum time for round-trip for this VLAN/MAC in us.

Name	Description
RttAvgWhole	Average time for round-trip for this VLAN/MAC in us.
RttAvgFract	Fractional portion of average time for round-trip.
Flag	Result flag indicating status or error code: <ul style="list-style-type: none"> • 1 - No error • 2 - Internal error • 3 - Invalid IP • 4 - L2Trace completed • 5 - Lookup failure for IP (no VLAN/MAC entries)

Configure Layer 2 IP Traceroute

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN
- If you want to run a Layer 2 IP Traceroute for a specific VRF, you must use EDM in the specific VRF context first. For more information, see the procedure for selecting and launching a VRF context view in [Select and Launch a VRF Context View](#) on page 4100.

About This Task

Use this procedure to configure Layer 2 IP traceroute.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**
3. Select the **L2 IP Traceroute** tab.
4. To add a new entry:
 - a. Select **Insert**.
 - b. Specify the destination IP address.
 - c. (Optional) Specify the TTL value.
 - d. Select **Insert**.
5. To start the Layer 2 IP traceroute, highlight an entry, and then select **Start**.
6. To stop the Layer 2 IP traceroute, select **Stop**.
7. To update the Layer 2 IP traceroute, select **Refresh**.

L2 IP Traceroute Field Descriptions

Use the data in the following table to use the **L2 IP Traceroute** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address (only IPv4 is supported).
IPAddr	Specifies the destination IP Address.
VrfId	Specifies the VRF ID.
VrfName	Specifies the name of the virtual router.
Ttl	Specifies the number of hops remaining to this L2Trace. This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The default value is 64
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Trace messages. • abort: the service is aborted or about to abort the L2Trace messages. <p>This field is also used to avoid concurrency or race condition problems that could occur if two or more management entities try to use the service at the same time. The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the Trace Messages will be or have been sent. • false: the Trace Messages will not be sent <p>The default is true.</p>
PathsFound	Specifies the number of paths found to execute the L2trace. The default is 0.

View Layer 2 IP Traceroute Results

About This Task

Use this procedure to view Layer 2 IP traceroute results.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 IP Traceroute** tab.
4. To view the Layer 2 IP traceroute results, highlight an entry, and then select **Result**.

L2 IP Traceroute Result Field Descriptions

Use the data in the following table to use the **L2 IP Traceoute Result** tab.

Name	Description
IpAddrType	Specifies the address type of destination IP address.
IpAddr	Specifies the destination IP address.
SendOrder	Denotes the order that sessions are sent. It is an index to distinguish among multiple L2Trace sessions. It correlates to the number of paths found. This value is assigned sequentially from 1.
Hop	Specifies the number of Layer 2 hops away from L2Trace initiator.
ReceiveOrder	Specifies the order that sessions are sent. It is an index to distinguish among multiple L2Trace responses with the same Send Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Specifies the time-to-live (TTL) field value for a returned L2Trace response.
VrfId	Specifies the VRF ID.
VlanId	Specifies the VLAN found from Layer 3 lookup and used for transmission.
DestMacAddress	Indicates the target MAC address transmitted.
PortNum	Specifies either the value '0', or the port number of the port used for the l2trace.
SeqNumber	Specifies the transaction identifier/sequence number used in linktrace message packet. The default is 0.
SrcMac	Specifies the MAC address of the MP that responded to L2Trace request for this L2traceReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the L2Trace to the responding node.

Name	Description
LastHostName	Specifies the host name of the node that forwarded the L2Trace to the responding node.
Flag	L2Trace result flag indicating status or error code: <ul style="list-style-type: none"> • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace

Trigger a Loopback Test

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

About This Task

Use this procedure to trigger a loopback test.

The LBM packet is often compared to ping. An MEP transmits the loopback message to an intermediate or endpoint within a domain for the purpose of fault verification. Use this function to check the ability of the network to forward different sized frames.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **CFM**.
3. Select the **LBM** tab.
4. Configure the loopback test properties as required.
5. Select **Apply**.
6. In **Status**, select **transmit**.
7. Select **Apply**.
8. Select **Refresh**.

LBM Field Descriptions

Use the data in the following table to use the **LBM** tab.

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the Maintenance Endpoint index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
DestMacAddress	Specifies the remote MAC address to reach the MEP/MIP.
Messages	Specifies the number of loopback messages to be transmitted. The default is 1.
VlanPriority	Specifies the priority. The default is 7.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. <p>The default is true.</p>
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the Loopback messages. • abort: The service is aborted or about to abort the Loopback messages. <p>The default is ready.</p>
Result	Displays the LBM result.
TimeoutInt	Specifies the timeout interval in seconds. The default value is 3 seconds.
InterFrameInt	Specifies the interval between LBM frames with a range of (0..1000) msec and a default value of 500 msec. The value of 0 msec indicates to send the frames as fast as possible. The default is 500.

Name	Description
TestPattern	<p>Specifies the testfill pattern:</p> <ul style="list-style-type: none"> • allZero — null signal without cyclic redundancy check • allZeroCrc — null signal with cyclic redundancy check with 32-bit polynomial • pseudoRandomBitSequence — pseudo-random-bit-sequence without cyclic redundancy check • pseudoRandomBitSequenceCrc — pseudo-random-bit-sequence with cyclic redundancy check with 32-bit polynomial. <p>A cyclic redundancy check is a code that detects errors. The default value is allZero.</p>
DataSize	Specifies the data type-length-value (TLV) size. The default is 0.
FrameSize	Specifies the frame-size. The default is 0.
Sourcemode Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.	<p>Specifies the source mode of the transmit loopback service:</p> <ul style="list-style-type: none"> • nodal • smltVirtual — Use the smltVirtual option with B-VLANs only. <p>The default is nodal.</p>

Trigger Linktrace

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.
- Assign a nodal MEP to the B-VLAN.

About This Task

Use the following procedure to trigger a linktrace. The linktrace message is often compared to traceroute. An MEP transmits the Linktrace Message packet to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. The packet specifies the target MAC address of an MP, which is the SPBM system ID or the virtual SMLT MAC. MPs on the path to the target address respond with an LTR.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **CFM**.
3. Select the **LTM** tab.
4. Configure the linktrace test properties as required.

5. Select **Apply**.
6. To trigger the linktrace test, perform one of the following actions:
 - In **Status**, select **transmit**, and then select **Apply**
 - Highlight an entry, and then select **Start**.
7. To update the linktrace, select the **Refresh** button.
8. To stop the linktrace, select **Stop**.

LTM Field Descriptions

Use the data in the following table to use the **LTM** tab.

Name	Description
DomainIndex	Specifies the MD index value.
AssociationIndex	Specifies the MA index value.
Index	Specifies the MEP index value.
DomainName	Specifies the MD name.
AssociationName	Specifies the MA name.
VlanPriority	Specifies the VLAN priority, a 3-bit value to be used in the VLAN tag, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the remote MAC address to reach the MEP.
Ttl	Indicates the number of hops remaining to this LTM. This value is decremented by 1 by each bridge that handles the LTM. The decremented value is returned in the LTR. If the value is 0 on output, the LTM is not transmitted to the next hop. The value of the TTL field in the LTM is specified at the originating MEP. The default value is 64.
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
ResultOk	Indicates the result of the operation: <ul style="list-style-type: none"> • true: The Loopback Messages will be (or have been) sent. • false: The Loopback Messages will not be sent. The default is true.

Name	Description
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: The service is available. • transmit: The service is transmitting, or about to transmit, the LTM messages. • abort: The service is aborted, or about to abort, the LTM message. <p>The default is ready.</p>
Flag	<p>Displays the LTM result flag indicating LTM status or error code. Each value represents a status or error case:</p> <ul style="list-style-type: none"> • 1 - No error • 2 - LTM internal error • 3 - Unknown Remote Maintenance Endpoint • 4 - Invalid Remote Maintenance Endpoint MAC Address • 5 - Unset Remote Maintenance Endpoint MAC address • 6 - MEP must be enabled in order to perform LTM • 7 - No LTR response received • 8 - Linktrace to own MEP MAC is not sent • 9 - Endpoint must be enabled in order to perform LTM • 10 - Pbt-trunk must be enabled in order to perform LTM • 11 - LTM completed • 12 - LTM leaf node
SourceMode Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.	<p>Specifies the source mode of the transmit loopback service:</p> <ul style="list-style-type: none"> • nodal • smltVirtual — Use the smltVirtual option with B-VLANs only. <p>The default is nodal.</p>

View Linktrace Results

About This Task

Use this procedure to view linktrace results.

After you trigger linktrace, you must select **Refresh** to update the results.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **CFM**.

3. Select the **LTM** tab.
4. Highlight an entry, and then select **Result**.

Link Trace Replies Field Descriptions

Use the data in the following table to use the **Link Trace Replies** tab.

Name	Description
DomainIndex	Indicates the Maintenance Domain Index.
AssociationIndex	Indicates the Maintenance Association Index.
MepId	Indicates the Maintenance EndPoint ID.
SeqNumber	Indicates the transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which LTM response is going to be returned. The default is 0.
Hop	Indicates the number of hops away from the LTM initiator.
ReceiveOrder	Indicates the index value used to distinguish among multiple LTRs with the same LTR Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the LTRs.
Ttl	Indicates the TTL field value for a returned LTR.
DomainName	Indicates the Maintenance Domain Name.
AssociationName	Indicates the Maintenance Association Name.
Forwarded	Indicates if a LTM was forwarded by the responding MP, as returned in the FwdYes flag of the flags field.
TerminalMep	Displays a boolean value stating whether the forwarded LTM reached a MEP enclosing its MA, as returned in the Terminal MEP flag of the Flags field.
LastEgressIdentifier	Displays an octet field holding the Last Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Last Egress Identifier identifies the MEP Linktrace Indicator that originated, or the Linktrace Responder that forwarded, the LTM to which this LTR is the response. This is the same value as the Egress Identifier TLV of that LTM.

Name	Description
NextEgressIdentifier	Displays an octet field holding the Next Egress Identifier returned in the LTR Egress Identifier TLV of the LTR. The Next Egress Identifier Identifies the Linktrace Responder that transmitted this LTR, and can forward the LTM to the next hop. This is the same value as the Egress Identifier TLV of the forwarded LTM, if any. If the FwdYes bit of the Flags field is false, the contents of this field are undefined, and the field is ignored by the receiver.
RelayAction	Indicates the value returned in the RelayAction field.
SrcMac	Displays the MAC address of the MP that responded to the LTM request for this LTR.
IngressAction	Displays the value returned in the IngressAction Field of the LTM. The value ingNoTlv indicates that no Reply Ingress TLV was returned in the LTM.
IngressMac	Displays the MAC address returned in the ingress MAC address field. If the rcCfmLtrReplyIngress object contains the value ingNoTlv(5), then the contents of this field are meaningless.
EgressAction	Displays the value returned in the Egress Action Field of the LTM. The value egrNoTlv(5) indicates that no Reply Egress TLV was returned in the LTM.
EgressMac	Displays the MAC address returned in the egress MAC address field. If the rcCfmLtrReplyEgress object contains the value egrNoTlv(5), then the contents of this field are meaningless.

Configure a Layer 2 Tracetree

Use this procedure to configure a Layer 2 Tracetree. This feature enables CFM to debug Layer 2. Layer 2 Tracetree allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. The command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP.
- Enable the MEP.

- Assign a nodal MEP to the B-VLAN.

About This Task

If you configure **IsTraceTree** to false, then EDM performs Traceroute on the unicast path. If you configure **IsTraceTree** to true, then EDM performs TraceTree on the multicast tree.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. From the **L2 Traceroute/TraceTree** tab, configure the Layer 2 tracetree properties.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Select **Apply**.
6. Select **Refresh** to update the results.

L2 Traceroute/TraceTree Field Descriptions

Use the data in the following table to use the **L2 Traceroute/TraceTree** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the Backbone VLAN (B-VLAN).
Priority	Specifies a 3-bit value to be used in the VLAN header, if present in the transmitted frame. The default is 7.
DestMacAddress	Specifies the target MAC address.
HostName	Specifies the target host name.
DestIsHostName	Specifies whether the host name is (true) or is not (false) used for the L2Trace transmission.
Isid	Specifies the Service Instance Identifier (I-SID).
IsTraceTree	Specifies whether the multicast tree or unicast path is traced: <ul style="list-style-type: none"> • If you configure IsTraceTree to false then EDM performs Traceroute on the unicast path. • If you configure IsTraceTree to true then EDM performs TraceTree on the multicast tree.

Name	Description
Status	<p>Indicates the status of the transmit loopback service:</p> <ul style="list-style-type: none"> • ready: the service is available. • transmit: the service is transmitting, or about to transmit, the L2Trace messages. • abort: the service aborted or is about to abort the L2Trace messages. <p>This field is also used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p> <p>The default is ready.</p>
ResultOk	<p>Indicates the result of the operation:</p> <ul style="list-style-type: none"> • true: the L2Trace messages will be (or have been) sent. • false: the L2Trace messages will not be sent. <p>The default is true.</p>
Ttl	<p>Specifies the number of hops remaining to this L2Trace.</p> <p>This value is decremented by 1 by each Bridge that handles the L2Trace. The decremented value is returned in the L2Trace. If 0 on output, the L2Trace is not transmitted to the next hop. The value of the time-to-live (TTL) field in the L2Trace is defined by the originating MEP.</p> <p>The default value is 64.</p>
SourceMode Note: Exception: smltVirtual is not supported on 4220 Series or 5320 Series.	<p>Specifies the source mode:</p> <p>The default is nodal.</p>

Name	Description
SeqNumber	Specifies the transaction identifier/sequence number of the first loopback message (to be) sent. The default is 0.
Flag	<p>L2Trace result flag that indicates L2Trace status or error code:</p> <ul style="list-style-type: none"> • none (1): No error • internalError (2): L2Trace internal error • invalidMac (3): Invalid MAC address • mepDisabled (4): MEP must be enabled in order to perform L2Trace • noL2TraceResponse (5): No L2Trace response received • l2TraceToOwnMepMac (6): L2Trace to own MEP MAC is not sent • l2TraceComplete (7): L2Trace completed • l2TraceLookupFailure (8): Lookup failure for L2Trace • l2TraceLeafNode (9): On a leaf node in the I-SID tree • l2TraceNotInTree (10): Not in the I-SID tree • l2TraceSmltNotPrimary (11): Requested SMLT source from non-primary node

View Layer 2 Tracetree Results

Use this procedure to view Layer 2 Tracetree results. The Layer 2 Tracetree command is a proprietary command that allows a user to trigger a multicast LTM message by specifying the B-VLAN and I-SID. This command allows the user to view a multicast tree on the SPBM B-VLAN from the source node to the destination nodes for a particular I-SID.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **L2Ping/L2Trace Route**.
3. Select the **L2 Traceroute/TraceTree** tab.
4. In the **IsTraceTree** field double-click and select **true** for EDM to perform Tracetree on the multicast tree.
5. Select **Apply**.
6. Select **Refresh** to update the results.
7. To view the tracetre results, highlight an entry, and then select **Result**.

L2 Traceroute/Tracetree Result Field Descriptions

Use the data in the following table to use the **L2 Traceroute/Tracetree Result** tab.

Name	Description
VlanId	A value that uniquely identifies the Backbone VLAN (B-VLAN).
SeqNumber	The transaction identifier/sequence number returned by a previous transmit linktrace message command, indicating which L2Trace's response of the L2Trace is going to be returned. The default is 0.
Hop	The number of hops away from L2Trace initiator.
ReceiveOrder	An index to distinguish among multiple L2Trace responses with the same Transaction Identifier field value. This value is assigned sequentially from 1, in the order that the Linktrace Initiator received the responses.
Ttl	Time-to-Live (TTL) field value for a returned L2Trace response.
SrcMac	MAC address of the MP that responds to the L2Trace request for this L2TraceReply.
HostName	The host name of the replying node.
LastSrcMac	The MAC address of the node that forwarded the L2Trace to the responding node.
LastHostName	The host name of the node that forwarded the L2Trace to the responding node.

Configure Layer 2 Trace Multicast Route

Before You Begin

- On the source and destination nodes, you must configure a CFM MD, MA, and MEP, and assign a nodal MEP to the B-VLAN.
- To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

About This Task

Use this procedure to configure the Layer 2 tracemroute on the VLAN (Layer 2) or VRF (Layer 3). This procedure queries the SPBM multicast module to determine the B-VLAN, I-SID, and nickname for the S and G streams. The nickname and I-SID are used to create a multicast MAC address.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > L2Ping/L2Trace Route**.
2. Select the **L2MCAST Traceroute** tab.
3. Select **Insert** to insert the Layer 2 MCAST Traceroute.
4. Type the **SrcIpAddr**.
5. Type the **GroupIpAddr**.
6. Select the **ServiceType**:
 - To perform a tracemroute on a VLAN, select **vlan**.
 - To perform a tracemroute on a VRF, select **vrfid**.
7. In **ServiceId**, type the VLAN ID.
8. Type the **Priority**.
9. Type the **Ttl** value.
10. Select **Insert**.
11. Select **Apply** to save your changes.
12. To start the Layer 2 tracemoute, configure the Status to transmit, and then select **Start**.
13. Update the Layer 2 tracemroute by selecting **Refresh**.
14. To stop the Layer 2 tracemroute, select **Stop**.
15. To see the result, select **Result**.

L2 MCAST Traceroute Field Descriptions

Use the data in the following table to use the **L2 MCAST Traceroute** tab.

Name	Description
SrcIpAddrType	Specifies the source IP address type as IPv4.
SrcIpAddr	Specifies the source IP address of the flow where the multicast trace tree originates.
GroupIpAddrType	Specifies the group IP address type as IPv4.
GroupIpAddr	Specifies the group IP address.
ServiceType	Specifies where you configure the Layer 2 tracemroute. This is either VLAN or VRF.
ServiceId	Specifies the VLAN ID.
VRFName	Specifies the VRF name.
Priority	Specifies the priority value. The value is between 0 and 7.
Ttl	Specifies the returned trace response. The TTL value is between 1 and 255.
SeqNumber	Specifies the transaction identifier/sequence number of the first message to be sent.

Name	Description
Status	<p>Specifies the status of the transmit loopback service:</p> <ul style="list-style-type: none"> ready: Specifies the service is available. transmit: Specifies the service is transmitting, or about to transmit the trace messages. abort: Specifies the services is aborted or about to abort the trace messages. <p>The column will also be used to avoid concurrency or race condition problems that can occur if two or more management entities try to use the service at the same time.</p>
ResultOK	<p>Specifies the result of the operation:</p> <ul style="list-style-type: none"> true: The trace messages will be or have been sent. false: The trace messages will not be sent.
Flag	<p>Specifies the result flag indicating that the Layer 2 trace status or error code. Each value represents a status or error case.</p> <ul style="list-style-type: none"> 1 — No error 2 — Internal Error 3 — Mep must be enabled to perform the trace 4 — No response received 5 — Trace completed 6 — On a leaf node in the I-SID tree 7 — No data I-SID was found for S, G

View Layer 2 Trace Multicast Route Results

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > L2Ping/L2Trace Route**.
2. Select the **L2 MCAST Traceroute** tab.
3. To view the CFMI2 trace multicast route results, highlight an entry and select **Result**.

L2tracemroute Result Field Descriptions

Use the data in the following table to use the **L2tracemroute Result** tab.

Name	Description
VlanId	Specifies a value that uniquely identifies the C-VLAN.
SeqNumber	Specifies the transaction identifier/sequence number returned by a previous transmit linktrace message command. Indicates which I2 tracemroute response is going to be returned.
Hop	Specifies the number of hops away from the I2 tracemroute initiator.
ReceiveOrder	Specifies an index to distinguish among multiple I2 tracemroute responses with the same transaction identifier field value. This value is assigned sequentially from 1, in the order that the linktrace initiator received the responses.
Ttl	Specifies the TTL value for a returned I2 tracemroute response.
SrcMac	Specifies the MAC address of the MP that responds to the I2 tracemroute request for this I2 tracemrouteReply.
HostName	Specifies the host name of the replying node.
LastSrcMac	Specifies the MAC address of the node that forwarded the I2 tracemroute to the responding node.
LastHostName	Specifies the host name of the node that forwarded the I2 tracemroute to the responding node.

CFM Configuration Examples

This section provides a configuration example for Connectivity Fault Management (CFM).

CFM Configuration Example

The following sections show the steps required to configure CFM.



Note

The following commands are not supported on all hardware platforms:

- cfm maintenance-domain
- cfm maintenance-association
- vlan nodal-mip-level

Switch A

```
MAINTENANCE-DOMAIN CONFIGURATION
```

```
cfm maintenance-domain "spbm" index 1 maintenance-level 6
```

```
MAINTENANCE-ASSOCIATION CONFIGURATION
```

```

cfm maintenance-association "spbm" "2" index 1
cfm maintenance-association "spbm" "3" index 2

MAINTENANCE-ENDPOINT CONFIGURATION

cfm maintenance-endpoint "spbm" "2" 1 state enable
cfm maintenance-endpoint "spbm" "3" 1 state enable

VLAN NODAL MEP/MIP CONFIGURATION

vlan nodal-mep 2 spbm 2 1
vlan nodal-mip-level 2 6
vlan nodal-mep 3 spbm 3 1
vlan nodal-mip-level 3 6

```

Switch B

```

MAINTENANCE-DOMAIN CONFIGURATION

cfm maintenance-domain spbm index 1 maintenance-level 6

MAINTENANCE-ASSOCIATION CONFIGURATION

cfm maintenance-association "spbm" "2" index 1
cfm maintenance-association "spbm" "3" index 2

MAINTENANCE-ENDPOINT CONFIGURATION

cfm maintenance-endpoint "spbm" "2" 2 state enable
cfm maintenance-endpoint "spbm" "3" 2 state enable

VLAN NODAL MEP/MIP CONFIGURATION

vlan nodal-mep 2 spbm 2 2
vlan nodal-mip-level 2 6
vlan nodal-mep 3 spbm 3 2
vlan nodal-mip-level 3 6

```

CFM Sample Output

The following sections show sample CFM output.

L2ping can use the system ID or the router name. The following example shows a case where the VLAN and MAC are given.

show isis adjacencies

```

Switch:1# show isis adjacencies
=====
                        ISIS Adjacencies
=====
INTERFACE L STATE  UPTIME      PRI  HOLDDTIME  SYSID          HOST-NAME  STATUS  AREA  AREA-NAME
-----
Port1/3   1  UP    00:37:37   127   22        0014.0dbf.a3df  Switch-Lab1  ACTIVE  HOME  area-9.00.02
Port1/19  1  UP    1d 05:09:16 127   25        0014.0da2.b3df  Switch-Lab2  BACKUP  HOME  area-9.00.02
-----

Home:    2 out of 2 interfaces have formed an adjacency
Remote:  0 out of 0 interfaces have formed an adjacency
-----

```

l2 ping with vlan

```
Switch:1# l2 ping vlan 500 mac 00.14.0d.bf.a3.df

Please wait for l2ping to complete or press any key to abort

----00:14:0d:bf:a3:df    L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 0 packets received,   100.00% packet loss
```

l2 ping with vlan

```
Switch:1# l2 ping vlan 500 routernodename MONTI0

Please wait for l2ping to complete or press any key to abort

----00:14:0d:a2:b3:df    L2 PING Statistics----  0(68) bytes of data
1 packets transmitted, 1 packets received,   0.00% packet loss
  round-trip (us)          min/max/ave/stdv =  26895/26895/26895.00/  0.00
```

l2 traceroute with vlan

```
Switch:1# l2 traceroute vlan 500 routernodename MONTI0

Please wait for l2traceroute to complete or press any key to abort

l2traceroute to MONTI0  (00:14:0d:a2:b3:df),  vlan 500
0   PETER4                (00:15:9b:11:33:df)
1   MONTI0                (00:14:0d:a2:b3:df)
```

l2 tracetree with vlan

```
Switch:1# l2 tracetree 500 1

Please wait for l2tracetree to complete or press any key to abort

l2tracetree to 53:55:10:00:00:01, vlan 500 i-sid 1 nickname 5.55.10 hops 64
1   PETER4                00:15:9b:11:33:df -> MONTI0          00:14:0d:a2:b3:df
2   MONTI0                00:14:0d:a2:b3:df -> LEE2            00:15:e8:b8:a3:df
```

l2 ping with IP address

You can also use L2ping and L2traceroute with an IP address. The following outputs show examples using an IP address.

```
Switch:1# l2 ping ip-address 192.0.2.10

Please wait for l2ping to complete or press any key to abort

L2 PING  Statistics : IP 192.0.2.10, paths found 1, paths attempted 1
=====
TX    RX    PERCENT  ROUND TRIP TIME
VLAN NEXT HOP                                PKTS  PKTS  LOSS    MIN/MAX/AVE (us)
=====
500  SHAMIM      (00:1a:8f:08:53:df)  1      0    100.00%  0/0/0.00
```

l2 ping with IPv6 address

```
Switch:1# l2 ping ip-address 49:0:0:0:0:0:11
```

```
Please wait for l2ping to complete or press any key to abort

L2 PING Statistics : IP 49:0:0:0:0:0:11, paths found 1, paths attempted 1
=====
VLAN NEXT HOP          TX      RX      PERCENT  ROUND TRIP TIME
                        PKTS  PKTS  LOSS      MIN/MAX/AVE (us)
=====
41  SHAMIM              (00:49:00:01:00:11)  1      1      0.00%  11876/11876/11876.00
```

l2 traceroute with IP address

```
Switch:1# l2 traceroute ip-address 192.0.2.10

Please wait for l2trace to complete or press any key to abort

L2 Trace Statistics : IP 192.0.2.10, paths found 1
=====
SHAMIM (00:1a:8f:08:53:df), vlan 500
0  PETER4 (00:15:9b:11:33:df)
1  MONTIO (00:14:0d:a2:b3:df)
```

l2 traceroute with IPv6 address

```
Switch:1# l2 traceroute ip-address 49:0:0:0:0:0:11

Please wait for l2trace to complete or press any key to abort

L2 Trace Statistics : IP 49:0:0:0:0:0:11, paths found 1
=====
SHAMIM (00:49:00:01:00:11), vlan 41
0  4K-DUT7 (00:49:00:01:00:17)
1  9k-2 (00:49:00:01:00:92)
2  8K-1 (00:49:00:08:00:81)
3  4K-DUT1 (00:49:00:01:00:11)
```

show cfm maintenance-domain

```
Switch:1#show cfm maintenance-domain
=====
Maintenance Domain
=====
Domain Name      Domain Index  Level Domain Type
-----
mdl              99           3      NONE

Total number of Maintenance Domain entries: 1.
```

show cfm maintenance-association

```
Switch:1#show cfm maintenance-association
=====
Maintenance Association Status
=====
Domain Name      Assn Name      Domain Idx  Assn Idx
-----
mdl              mal            1           98

Total number of Maintenance Association entries: 1.
```

```

=====
Maintenance Association config
=====
Domain Name      Assn Name
-----
md1              mal

Total number of MA entries: 1.

```

show cfm maintenance-endpoint

```

Switch:1#show cfm maintenance-endpoint
=====
Maintenance Endpoint Config
=====
DOMAIN      ASSOCIATION      MEP  ADMIN
NAME        NAME            ID
-----
md1         mal             1   enable

Total number of MEP entries: 1.

=====
Maintenance Endpoint Service
=====
DOMAIN_NAME  ASSN_NAME      MEP_ID TYPE  SERVICE_DESCRIPTION
-----
md1         mal             1   unused

Total number of MEP entries: 1.

```

show vlan nodal-mep

```

Switch:1#show vlan nodal-mep
=====
Vlan Nodal Mep
=====
VLAN_ID  DOMAIN_NAME.ASSOCIATION_NAME.MEP_ID
-----
1
2
3
4      md1.mal.1
5
6
7
8
9
10
11
12
13
14

```

show vlan nodal-mip-level

```

Switch:1#show vlan nodal-mip-level
=====
Vlan Nodal Mip Level
=====
VLAN_ID  NODAL_MIP_LEVEL_LIST

```

```
-----  
1  
2  
3  
4      6  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14
```

Software Troubleshooting Tool Configuration

This section provides troubleshooting information for common software problems.

Troubleshooting Tool Fundamentals

This section provides conceptual information about the methods and monitoring tools you can use for troubleshooting problems. This section also contains precautionary notices that you must read for the safe operation of the network.

Troubleshooting Overview

The types of problems that typically occur with networks involve connectivity and performance. This section also contains precautionary notices that you must read for the safe operation of the switch. The switch supports a diverse range of network architectures and protocols, some of which maintain and monitor connectivity and isolate connectivity faults.

In addition, the switch supports a wide range of diagnostic tools that you can use to monitor and analyze traffic, capture and analyze data packets, trace data flows, view statistics, and manage event messages.

Certain protocols and tools are tailored for troubleshooting specific switch network topologies. Other tools are more general in their application and you can use them to diagnose and monitor ingress and egress traffic on the switch.

If connectivity problems occur and the source of the problem is unknown, it is usually best to follow the Open Systems Interconnection (OSI) network architecture layers. Confirm that your physical environment, such as the cable and port connections, operates without failures before moving up to the network and application layers.

To gather information about a problem, consider the following information:

- Consider the OSI model when you troubleshoot. Start at Layer 1 and move upwards. The Address Resolution Protocol (ARP) can cause problems; ARP operates at Layer 2 to resolve MAC addresses to IP addresses (Layer 3).
- Device-specific tools and protocols can help you gather information. This document outlines switch-specific tools.

- You can use client- and server-based tools from Microsoft, Linux, and UNIX. For example, you can use Windows tools like `ifconfig`, `ipconfig`, `winipcfg`, and `route print` to obtain IP information and routing tables. Servers also maintain route tables.

The following command output shows example output of the `route print` command.

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.
C:\Documents and Settings\jsmith>route print

=====
Interface List
0x1 ..... MS TCP Loopback interface
0x2 ...00 12 f0 74 2a 87 ..... Broadcom NetLink (TM) Gigabit Ethernet - Packet
Scheduler Miniport
0x3 ...00 14 38 08 19 c6 ..... Broadcom NetXtreme Gigabit Ethernet - Packet
Scheduler Miniport
0x4 ...44 45 53 54 42 00 ..... IPSECshm Adapter - Packet Scheduler
Miniport
=====
Active Routes:

Network          Destination      Netmask          Gateway          Metric
0.0.0.0           0.0.0.0          192.168.0.1       192.168.0.102    26
0.0.0.0           0.0.0.0          207.179.154.100   207.179.154.100  1
127.0.0.0         255.0.0.0        127.0.0.1         127.0.0.1        1
192.168.0.0       255.255.255.0    192.168.0.102     192.168.0.102    25
192.168.0.0       255.255.255.0    207.179.154.100   207.179.154.100  1
192.168.0.102     255.255.255.255  127.0.0.1         127.0.0.1        25
192.168.0.255     255.255.255.255  192.168.0.102     192.168.0.102    25
198.164.27.30     255.255.255.255  192.168.0.1       192.168.0.102    1
207.179.154.0     255.255.255.0    207.179.154.100   207.179.154.100  30
207.179.154.100   255.255.255.255  127.0.0.1         127.0.0.1        30
207.179.154.255   255.255.255.255  207.179.154.100   207.179.154.100  30
224.0.0.0         240.0.0.0        192.168.0.102     192.168.0.102    25
224.0.0.0         240.0.0.0        207.179.154.100   207.179.154.100  1
255.255.255.255   255.255.255.255  192.168.0.102     192.168.0.102    1
255.255.255.255   255.255.255.255  207.179.154.100   3                 1
255.255.255.255   255.255.255.255  207.179.154.100   207.179.154.10  1
Default Gateway: 207.179.154.100
=====
Persistent Routes:  None
```

- Other network problems can give the impression that a device has a problem. For instance, problems with a Domain Name Service (DNS) server, another switch, firewall, or access point can appear to be routing problems.

Debug Files

The switch stores debug files in the `intflash` directory.

The debug file is in a zipped format and contains information to help debug the device, including:

- a memory snapshot
- logs
- traces

This best practice is to delete these files to ensure enough space exists in the internal flash. New files do not overwrite old files. You must remove the files; otherwise, the internal flash may not have enough free space for necessary activities, for example, to store a core dump file if the switch fails, or you may not have the space to transfer a new release to the internal flash to upgrade your switch.

The switch stores a maximum of 32 files for each debug file for each slot, depending on the file size of each debug file. The internal flash provides 2 GB of storage. The system displays a message on the console to inform you when less than 700 MB is available.

The **debug-file remove** command can delete the following types of debug files:

- core
- archive
- PMEM
- dmalloc
- flrec
- wd_stats

If you want to delete a specific file, you must use the **remove** command.

SNMP

The switch does not support SNMP for the **show debug-file** or the **debug-file remove** commands.

Flight Recorder

Table 326: Flight Recorder product support

Feature	Product	Release introduced
Flight Recorder for system monitoring	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

The Flight Recorder is a high level term for the framework in place on the switch to store both history and current state information for various kernel, system, and application data with minimal overhead to execution. This data can later be accessed on-demand when debugging systems issues to give engineers the best

possible troubleshooting information. Functionally, the Flight Recorder consists of two elements; Persistent Memory and Always-on Trace.

The Persistent Memory feature stores information in volatile memory outside of any process. This feature provides information on crashes, errors, and outages that are not the result of a power failure. Persistent Memory data not saved to non-volatile storage before a power failure will be lost. Persistent Memory snapshots are taken when:

- a critical process stops functioning
- a process stops responding
- the hardware watchdog activates
- the user initiates a snapshot in the CLI

The Always-on Trace feature creates an ongoing, circular log of every trace call recently executed regardless of the trace level enabled by the user. The Always-On Trace feature uses circular logging, and therefore stores the most recent traces of the process.

Flight Recorder functionality is provided only through CLI. This feature uses the following commands:

- **flight-recorder all {slot[-slot]}[,...]**

The command creates a flight-recorder snapshot, trace and archive.

- **flight-recorder archive {slot[-slot]}[,...]**

This command creates a tarball of flight-recorder files, log files, and configuration files.

- **flight-recorder snapshot {slot[-slot]}[,...]**.

This command takes a snapshot of PMEM data.

- **flight-recorder trace {slot[-slot]}[,...]**

This command takes a snapshot of always-on-trace data.

Network Service Probe

Table 327: Network Service Probe product support

Feature	Product	Release introduced
VLAN and I-SID network service probe to troubleshoot IP and DHCP connectivity issues	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.2
	5420 Series	Fabric Engine 9.2
	5520 Series	Fabric Engine 9.2
	5720 Series	Fabric Engine 9.2
	7520 Series	Fabric Engine 9.2
	7720 Series	Fabric Engine 9.2
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 9.2
	VSP 7400 Series	VOSS 9.2
Domain resolution test for a specific DNS server	4220 Series	Fabric Engine 9.4
	5320 Series	Fabric Engine 9.4
	5420 Series	Fabric Engine 9.4
	5520 Series	Fabric Engine 9.4
	5720 Series	Fabric Engine 9.4
	7520 Series	Fabric Engine 9.4
	7720 Series	Fabric Engine 9.4
	7830 Series	Fabric Engine 9.4
	VSP 4900	VOSS 9.4
	VSP 7400	VOSS 9.4

Create a Network Service Probe to troubleshoot network connectivity issues. A Network Service Probe interface supports the following configurations:

- Port-based VLAN with or without an attached I-SID
- Static IP address and static default gateway

**Note**

The default gateway must be non-local to the switch.

- Dynamic address resolution through DHCP

**Note**

The DHCP server and relay agent must be non-local to the switch.

- IPv4 only

A Network Service Probe interface supports the following functions:

- Provide feedback about its own IP and MAC address.
- Launch a ping request to any destination.
- Provide feedback about ping results and latency (CLI only).
- Respond to ARP and ping requests.

Dynamic DHCP and DNS Behavior

When you use dynamic IP configuration, the Service Probe interface receives and stores DNS servers from the DHCP server along with the IP address and default gateway.

The interface can store up to three dynamic DNS servers advertised by DHCP, designated as the following:

- primary-dynamic
- secondary-dynamic
- tertiary-dynamic

These DNS settings apply dynamically and they not saved in the configuration file. Additionally, you can query the state of any of the three dynamic DNS servers to determine whether they are functioning successfully or have failed.

Port Mirroring

Table 328: Port Mirroring product support

Feature	Product	Release introduced
Ingress mirroring (port and flow-based)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3 Note: Mirror destination MLT does not apply to 7830 Series.
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 328: Port Mirroring product support (continued)

Feature	Product	Release introduced
Egress mirroring (port-based)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

True egress port-based mirroring that produces an identical copy of an outgoing packet is not supported. The mirrored copy does not reflect changes that occur in the switch to the outgoing packet (for example, packet fields that are updated during IP routing). As a result, the mirrored copy is not identical to the outgoing packet.

Use the port mirroring feature to monitor and analyze network traffic. Port mirroring supports both ingress (incoming traffic) and egress (outgoing traffic) port mirroring. When you enable port mirroring, the system forwards ingress or egress packets normally from the mirrored (source) port, and sends a copy of the packet to the mirroring (destination) port.

Port Mirroring Overview

Port mirroring causes the switch to make a copy of a traffic flow and send the copy to a device for analysis. Use port mirroring in diagnostic sniffing—use the mirror to view the packets in the flow without breaking the physical connection to place a packet sniffer inline. You can also use mirroring for security reasons.

You can use egress mirroring to monitor packets as they leave specified ports. Egress mirroring on the switch is done at the end of the ingress pipeline. Since packet modifications occur in the egress pipeline, some of the changes will not be reflected in the mirrored version of the packet. Changes that occur in the egress pipeline may be reflected in the mirrored packet due to the metadata that is carried with the packet. Metadata notifies the egress pipeline what to change.

Use a network analyzer to observe and analyze packet traffic at the mirroring port. Unlike other methods that analyze packet traffic, the packet traffic is uninterrupted and packets flow normally through the mirrored port.

You can mirror to a port or list of ports or a MultiLink Trunking (MLT) group. The switch supports one-to-many, many-to-one, and many-to-many mirroring configurations.

Ingress and Egress Mirrored Ports

You can use all ports in the system to function as an ingress port for mirroring (mirrored port), an egress port for mirroring (mirrored port), or as a mirroring port (where all the mirrored traffic is redirected). The number of mirroring ports (also called destination ports) that you can configure is limited by the hardware. The hardware limitation is four ports simultaneously (where each mirroring direction counts as one). For example, if two mirroring ports are designated to mirror both ingress and egress traffic then all four mirroring ports are consumed.

The following table describes ingress mirroring functionality. Only one type of mirroring destination is supported at a time. You cannot mirror the same port to multiple classes of destinations, for example, MLT. However, you can mirror to multiple physical destinations.



Important

Mirroring packets from one network-to-network interface (NNI) port to another NNI port is not supported. Mirror to access ports, not NNI ports.

Table 329: Ingress mirroring functionality

Function	Support information
Ingress port mirroring and ingress flow mirroring	Supported. Maximum of four mirror-to-ports per box.
One port to one port	Supported
One to MLT group [for threat protection system (TPS applications)]	Supported
One to many (multicast group ID/VLAN)	Not supported
One to one (remote mirrored destination)	Not supported
Many to one (multiple mirrored ports to one mirroring port)	Supported
Many to MLT group	Supported
Many to many (VLAN/multicast group ID) (multiple ports with several different destinations)	Not supported
Many to one (relation between Remote Mirror Source [RMS] and Remote Mirror Termination [RMT])	Not supported
VLAN and port combination as a mirroring destination	Not supported
Ingress flow mirroring	Supported
Allow filters to specify a separate destination for each access control entry	Supported

The following table describes egress mirroring functionality.

Table 330: Egress mirroring functionality

Function	Support information
Egress port mirroring	Supported
One port to one port	Supported
One to MLT groups (for TPS applications)	Supported
One to many (multicast group ID/VLAN)	Not supported
Many to one (multiple mirrored ports to one mirroring port)	Supported
Many to MLT group	Supported
Many to many (multicast group ID) (multiple ports with several different destinations)	Supported
Many to one (relation between Remote Mirror Source [RMS] and Remote Mirror Termination [RMT])	Not supported
VLAN and port combination as mirroring destination	Not supported
Egress flow mirroring	Supported
Allow filter to specify a separate destination for each access control entry	Supported

Port Configuration

You can specify a destination multilink trunking (MLT) group, a destination port or set of ports.

There are two port mirroring modes: rx (ingress, that is, inPort) and tx (egress, that is, outPort). In rx mode, when you configure the ACE mirror or ACL global options to mirror packets, use the ACE to configure the mirroring destination port.

To modify a port mirroring instance, first disable the instance. Also, to change a port or MLT entry, first remove whichever parameter is attached to the entry, and then add the required entry.

ACLs, ACEs, and Port Mirroring

You can use filters to reduce the amount of mirrored traffic. You can configure the mirroring action globally in an access control list (ACL), or for a specific access control entry (ACE) by using the ACE mirror actions. If you use the global action, mirroring applies to all ACEs that match in an ACL.

To use filters with port mirroring, apply an ACL to the mirrored port in the egress and ingress directions. Traffic patterns that match the ACL or ACE with an action of permit are forwarded to the destination and also to the mirroring port. Traffic patterns that match an ACE with an action of drop (deny) are not forwarded to the destination, but still reach the mirroring port. For example, for an ACL or ACE with a match action of permit, packets are mirrored to the specified mirroring destination on the ACE. If you enable a port or VLAN filter, then that filter is the mirroring filter.

You can specify more than one mirroring destination by using multiple ACEs. Use each ACE to specify a different destination.

You can configure a port-based and a flow-based mirroring filter on the same port. If such a case occurs, then the flow-based mirror takes precedence.

For information about how to configure ACLs and ACEs for port mirroring using CLI, see the following sections:

- [Configure Global Mirroring Actions with an ACL](#) on page 3897
- [Configure ACE Actions to Mirror](#) on page 3897

For information about how to configure ACLs and ACEs for port mirroring using EDM, see the following sections:

- [Configure ACLs for Mirroring](#) on page 3920
- [Configure ACEs for Mirroring](#) on page 3922

Port Mirroring Considerations and Restrictions

Although you can configure the switch to monitor both ingress and egress traffic, some restrictions apply:

- The software does not support true egress mirroring because packets are mirrored prior to the completion of packet processing, so egress mirrored packets can differ from the packets egressing the port.



Note

To mirror the egress traffic, you can use the NEXT-hop device ingress mirroring to capture the egress packets of the switch.

- Mirrored traffic shares ingress queue and fabric bandwidth with normal traffic and therefore can impact normal traffic. Therefore, use these features with this potential consequence in mind and enable them only for troubleshooting, debugging, or for security purposes such as packet sniffing, intrusion detection, or intrusion prevention.
- You can configure as many ingress mirroring flows as you have filters.
- To avoid VLAN members from seeing mirrored traffic, you must remove mirroring (destination) ports from all VLANs.
- The MAC drops an error packet, for example, packets that are too short or too long. Control packets consumed by the MAC (802.3x flow control) are also not mirrored.
- Certain control packets generated by the CP cannot be egress mirrored, such as those in the following list:
 - BPDU
 - EAPoL
 - IP Directed Broadcast
 - LACP
 - LLDP
 - Multicast routed packets
 - NAAP

- NLB
- Nodal CFM
- TDP
- VLACP
- The system displays ingress multicast packets in egress mirroring.
- For switches that support Extreme Integrated Application Hosting:
 - To use an Extreme Integrated Application Hosting port with a connect type as OVS or SR-IOV for Port Mirroring, associate VLAN 4091 to the virtual machine (VM) vport to send the mirrored packets to the VM.
- If incoming traffic from the same source port is simultaneously ingress mirrored on an incoming port into one I-SID and egress mirrored on another outgoing port into a different I-SID, the mirrored packet carries an I-SID associated with ingress mirroring.
- Incoming traffic that does not contain a VLAN tag is not mirrored into an I-SID if the offset ID is in the range 2 to 1000. It is mirrored to an I-SID only if the offset ID is 1.
- The original CVLAN tag on the mirrored packet is preserved for only one mirrored I-SID if the offset ID is 1. The original CVLAN tag is not preserved in a mirrored packet for all other remaining mirrored I-SIDs if the offset ID is in the range 2 to 1000.

Port Mirroring Resources

Port mirroring resources are limited to four ports simultaneously (where each mirroring direction counts as one). For example, if two mirroring ports are designated to mirror both ingress and egress traffic then all four mirroring ports are consumed.

Port mirroring shares these four resources with other applications such as port mirroring RSPAN, Fabric Extend, Application Telemetry, IPFIX, and ACL with mirror action. Each one of these applications consumes at least one port mirroring resource. (port mirroring RSPAN consumes two if you configure both Ingress and Egress modes.)



Important

To enable any one of the preceding applications, you must have at least one free mirroring resource. If all four port mirroring resources are already in use, the switch displays a `Resource not available` error message when you try to enable the application.

If you receive a `Resource not available` error message, you can use the **show mirror-resources** command to view information about mirror resource usage. For more information, see [Display Mirror Resource Usage](#) on page 3896.

General Diagnostic Tools

The switch has diagnostic features available with Enterprise Device Manager (EDM) and Command Line Interface (CLI). You can use these diagnostic tools to help you troubleshoot operational and configuration issues. You can perform such tasks as configuring and displaying log files, viewing and monitoring port statistics, tracing a route, running loopback and ping tests, and viewing the address resolution table.

Traceroute

Traceroute determines the path a packet takes to reach a destination by returning the sequence of hops (IP addresses) the packet traverses.

According to RFC1393, traceroute operates by: "sending out a packet with a time-to-live (TTL) of 1. The first hop then sends back an ICMP error message indicating that the packet could not be forwarded because the TTL expired. The packet is then resent with a TTL of 2, and the second hop returns the TTL expired. This process continues until the destination is reached. The purpose behind this is to record the source of each ICMP TTL exceeded message to provide a trace of the path the packet took to reach the destination."

Ping

Ping is a simple and useful diagnostic tool to determine reachability. When you use ping, the switch sends an ICMP echo request to a destination IP address. If the destination receives the packet, it responds with an ICMP echo response.

If a ping test is successful, the destination is alive and reachable. Even if a router is reachable, it could have improperly working interfaces or corrupted routing tables.

Trace

Use trace commands to provide detailed data collection about software modules on the switch. The trace toolset can be used to trace multiple modules simultaneously and provides options to specify the verbosity level of the output.

You can enable trace logging through the boot config trace-logging flag.



Caution

Risk of traffic loss

Using the trace tool inappropriately can cause a CPU lockup conditions, loss of access to the switch, loss of protocols, and service degradation.

**Tip**

While these occurrences are uncommon, when using the trace level tool, minimize this risk. The following actions are:

- In situations where trace data is required concurrently from multiple modules, consider troubleshooting during a maintenance window if feasible. Consider a maintenance window period if the switch is stable but CPU utilization is high and CPU traces (example trace levels 9 and 11) are required to diagnose the cause.
- Run trace commands from the console port when the CPU utilization is already high. While you can enable or disable tracing when directly connected to the console port.

Activate tracing on one software module at a time.

- Initially activate tracing at lower verbosity settings (that is, 2 rather than 3). Increase to verbosity level 3 or 4 only if required, and after level 2 runs safely.
- Avoid leaving traces active for extended periods of time. For high CPU utilizations, a few seconds (typically less than 5 seconds) is generally sufficient to identify the cause for sustained high CPU utilization.

Fabric RSPAN (Mirror to I-SID)

Table 331: Fabric RSPAN product support

Feature	Product	Release introduced
Fabric RSPAN (Mirror to I-SID)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Remote mirroring is an important functionality that helps in:

- Intrusion Detection or Intrusion Prevention Systems
- Network Port debugging and packet capture
- Mirror and Monitor traffic to central collector or analyzers
- Mirror and Monitor traffic to distributed collectors or analyzers

With the Fabric RSPAN (Mirror to I-SID) feature, mirrored traffic captured from any switch in the network is sent to a remote switch over an SPB cloud for traffic analysis.

With this feature, you can monitor traffic on ports from different switches connected in the network, using just one network analyzer connected to a remote switch which acts as a collector. The source device where the traffic is mirrored to an I-SID, is known as Mirroring BEB (Backbone Edge Bridge), and the remote device where the traffic analyzer is connected for mirrored traffic analysis is known as Monitoring BEB.

The traffic source on the mirroring BEB is supported in the following ways:

- Port based mirroring — Any packet incoming or outgoing through a port is mirrored to a monitoring I-SID configured for that port.
- Flow based mirroring — Any particular packet flow configured in the system using filter based ACLs is mirrored to a monitoring I-SID configured for that flow.

Fabric RSPAN (Mirror to I-SID) Restrictions and Requirements

- Remote mirroring of traffic is not supported on network-to-network interface (NNI) ports or Fabric Extend Layer 2 core ports.
- Mirroring resources will be shared between Fabric RSPAN and regular port mirroring. Fabric RSPAN uses one out of four resources for mirroring if the mode is configured as Rx (Ingress) mirroring. In case of mode Tx (Egress) mirroring, it uses one more entry with same TX-LB port. Hence if mode Rx and Tx are configured for Fabric RSPAN, then only two unique destination ports can be used for regular port mirroring. For example, if you configure Fabric RSPAN on the switch, the regular port mirroring functionality can use only three unique destination ports. And, if all the four unique ports are used by the port mirroring functionality, you cannot configure Fabric RSPAN functionality on that node.
- When the monitor I-SID used for mirroring Fabric RSPAN traffic ingressing to I-SID is used to mirror regular traffic into SPB network, it will remove the customer tag in the mirrored packets. Hence, as a best practice, use different monitor I-SIDs for mirroring regular traffic and Fabric RSPAN traffic.
- Monitoring egress-ports and egress-mlt will not support regular production network traffic.
- The QoS value must be same for all mirror entries having common monitor I-SID, as the BMAC QoS is mapped to a monitor I-SID. QoS value configured for a specific monitor I-SID offset overrides the existing value for all mirroring entries sharing the same monitor I-SID.
- You cannot configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.
- Egress flow-based I-SID mirroring is not supported.
- **monitor-by-isid** works only for VLAN-tagged traffic.
- 5520 Series does not support the following combinations for I-SID mirroring:
 - Deny option for Filter ACL with monitor-isid-offset for Inport/Invlan/InVSN
 - Redirect-Next-Hop along with monitor-isid-offset when monitor-isid-offset is greater than 1

- Remove Tag option along with monitor-isid-offset
- For 5720 Series, 7520 Series, and 7720 Series:
 - To use an Extreme Integrated Application Hosting port as an analyzer port on a monitoring BEB for Fabric RSPAN (Mirror to I-SID), you must associate outer-tag 4091 to egress port 1/s1 or 1/s2 if the connect type is OVS or SR-IOV. Use the **monitor-by-isid <1-1000> map-to-vid <1-4093>** command to configure VLAN 4091 for Fabric RSPAN.

Software Troubleshooting Tool Configuration Using CLI

Use the tools described in this section to perform troubleshooting procedures using CLI.

Use CLI for Troubleshooting

You can use CLI to provide diagnostic information.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Disable scrolling of the output display:

```
terminal more disable
```

3. View configuration file information:

```
more WORD<1-99>
```

4. Capture the output for the following command after you observe a problem with the device:

```
show running-config [verbose] [module {app-telemetry | boot | cfm |
cli | diag | dvr | eap | endpoint-tracking | energy-saver | fa | fhs
| filter | ike | ip | ipfix | ipsec | ipv6 | iqagent | isis | i-sid |
lacp | license | lldp | lst | macsec | mlt | naap | nls | net-debug |
node-alias | ntp | openapi | ovssdb | port | qos | radius | rate-limit
| restconf | rmon | sflow | security | slpp | smtp | spbm | stg | sys
| tacacs | virtualservice | vlan | web | vxlan | vtp}]
```

5. Capture the output for the following command after you observe a problem with the device:

```
show tech
```



Note

`show tech` is an alias for `show fulltech`. Both commands provide the same output.

- Capture the output for the following commands after you observe a problem with the device:



Note

The **show interfaces gigabitEthernet statistics rmon** command displays information only if you previously configured **rmon stats** or **rmon history**.

- `show interfaces gigabitEthernet statistics <dhcp-relay [vrf WORD<1-16>][vrfids WORD<0-512>] [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}`
- `show interfaces gigabitEthernet statistics lacp [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}`
- `show interfaces gigabitEthernet statistics rate-limiting [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}`
- `show interfaces gigabitEthernet statistics rmon [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}`
- `show interfaces gigabitEthernet statistics verbose [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}`
- `show interfaces gigabitEthernet statistics vlacp [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}`

- Capture the output for the following command after you observe a problem with the device:

```
show interfaces gigabitEthernet error [collision|ospf|verbose] [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]}
```

Examples

Due to the length of command output, the following examples are truncated.

Capture the output for the following commands after you observe a problem with the device:

```
Switch:1>enable
Switch:1#show running-config module cli
Preparing to Display Configuration...
#
# Fri Feb 25 12:34:46 2022 UTC
# box type           : <hardware dependent>
# software version   : 8.6.0.0.GA
# cli mode           : ECLI
#
#Card Info :
# Slot 1 :
#   CardType           : <hardware dependent>
#   CardDescription    : <hardware dependent>
#   CardSerial#        : TB012132K-H0057
#   CardPart#          : 801104-00-02
#   CardAssemblyDate   : 20210817
#   CardHWRevision     : 02
#   CardHWConfig       :
#   AdminStatus        : up
```

```
#      OperStatus      : up

--More-- (q = quit)
Switch:1#show tech

Sys Info:
-----

General Info :

      SysDescr      : <hardware dependent> (8.6.0.0.GA)
      SysName       : <hardware dependent>
      SysUpTime     : 1 day(s), 02:39:22
      SysContact    : http://www.extremenetworks.com/contact/
      SysLocation   :

Chassis Info:

      Chassis       : <hardware dependent>
      ModelName     : <hardware dependent>
      BrandName     : Extreme Networks.
      Serial#       : TB012132K-H0057
      H/W Revision  : 02

--More-- (q = quit)
Switch:1#show interfaces gigabitethernet statistics

=====
                        Port Stats Interface
=====
```

PORT NUM	IN OCTETS	OUT OCTETS	IN PACKET	OUT PACKET
1/1	1215232	1852156	18988	25083
1/2	11866260	3650340	128847	51849
1/3	0	0	0	0
1/4	0	0	0	0
1/5	0	0	0	0
1/6	2606433776	2605569408	40718802	40712022
1/7	2383797478	2368788480	37189478	37012320
1/8	2639779622	2624836140	41201664	40945760
1/9	0	0	0	0
1/10	0	0	0	0
1/11	0	0	0	0
1/12	0	6776546	0	62572
1/13	1215232	997632	18988	15588
1/14	7459408	1396224	69625	18702

```
--More-- (q = quit)

Switch:1>enable
Switch:1#show interfaces gigabitethernet error 1/15

=====
                        Port Ethernet Error
=====
```

PORT NUM	ERROR ALIGN	ERROR FCS	FRAMES LONG	TOO SHORT	LINK FAILURE	CARRIER SENSE	CARRIER ERROR	SQETEST ERRORS	IN DISCARD	RATE-LIMIT DROP-COUNTER
1/15	0	0	0	0	2	0	0	0	7800	780

Variable Definitions

The following table defines parameters for the **more** command.

Variable	Value
<code>WORD<1-99></code>	Specifies the file name to view. Provide the file name in one of the following formats: a.b.c.d:<file> or /intflash/<file>.

The following table defines parameters for the **show running-config** command.

Variable	Value
<code>[module {app-telemetry boot cfm cli diag dvr eap endpoint-tracking energy-saver fa fhs filter ike ip ipfix ipsec ipv6 iqagent isis i-sid lacp license lldp lst macsec mlt naap nls net-debug node-alias ntp openapi ovsdb port qos radius rate-limit restconf rmon sflow security slpp smtp spbm stg sys tacacs virtualservice vlan web vxlan ptp}]</code>	Specifies the command group for which you request configuration settings. Note: All configuration modules are not supported on all hardware platforms. For information about feature support, see <i>Fabric Engine and VOSS Feature Matrix</i> .
<code>verbose</code>	Specifies a complete list of all configuration information about the switch.

The following table defines parameters for the **show interfaces gigabitEthernet statistics** command.

Variable	Value
<code>dhcp-relay [vrf WORD<1-16>] [vrfids WORD<0-512> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]]}</code>	Displays port Dynamic Host Configuration Protocol (DHCP) statistics.
<code>lacp {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	Displays Link Aggregation Control Protocol (LACP) statistics.
<code>rate-limiting {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	Displays port ingress rate-limiting statistics.
<code>rmon {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]} [history]</code>	Displays Remote Network Monitoring (RMON) statistics.

Variable	Value
<i>verbose</i>	Displays a complete list of all statistics.
<i>vlacp</i> [<i>history</i>] { <i>slot/port[/sub-port]</i> [<i>-slot/port[/sub-port]</i>] [,...] [<i>slot/all</i>] [<i>all</i>] }	<p>Displays port Virtual Link Aggregation Control Protocol (VLACP) statistics.</p> <ul style="list-style-type: none"> <i>history</i>—Displays the VLACP port counter profile. {<i>slot/port[/sub-port]</i> [<i>-slot/port[/sub-port]</i>] [,...] [<i>slot/all</i>] [<i>all</i>] } <p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> ◦ a single slot and port (<i>slot/port</i>) ◦ a range of slots and ports (<i>slot/port-slot/port</i>) ◦ a series of slots and ports (<i>slot/port,slot/port,slot/port</i>) ◦ all ports on the same slot (<i>slot/all</i>) ◦ all ports on the switch (<i>all</i>) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format <i>slot/port/sub-port</i>.</p>

The following table defines parameters for the **show interfaces gigabitEthernet error** command.

Variable	Value
<i>collision</i>	Displays port collision error information.
<i>ospf</i>	Displays ports Open Shortest Path First (OSPF) error information.
{ <i>slot/port[/sub-port]</i> [<i>-slot/port[/sub-port]</i>] [,...] [<i>slot/all</i>] [<i>all</i>] }	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> ◦ a single slot and port (<i>slot/port</i>) ◦ a range of slots and ports (<i>slot/port-slot/port</i>) ◦ a series of slots and ports (<i>slot/port,slot/port,slot/port</i>) ◦ all ports on the same slot (<i>slot/all</i>) ◦ all ports on the switch (<i>all</i>) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format <i>slot/port/sub-port</i>.</p>
<i>verbose</i>	Displays all port error information.

Use Software Record Dumps

About This Task

Capture a dump of the software records from ingress traffic to help troubleshoot performance problems. Generally, a verbosity level of 1 suffices.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Dump software record information:
`dump ar <1-12> WORD<1-1536> <0-3>`

Example

```
Switch:1>enable
Switch:1#dump ar 1 vlan 1
```

Variable Definitions

The following table defines parameters for the **dump ar** command.

Variable	Value
<1>	Specifies the slot number.
WORD<1-1536>	Specifies a record type in the AR table. Options include vlan, ip_subnet, mac_vlan, mac, arp, ip, ipmc, protocol, all.
<0-3>	Specifies the verbosity from 0–3. Higher numbers specify more verbosity.

Use Trace to Diagnose Problems

Use trace to observe the status of a software module at a given time.

About This Task

For example, if you notice a CPU utilization issue (generally a sustained spike above 90%) perform a trace of the control plane activity.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Clear the trace:
`clear trace`
3. Identify the module ID for which you want to use the trace tool:
`show trace modid-list`
4. Begin the trace operation:
`trace level [<Module_ID>] [<0-4>]`

5. Wait approximately 30 seconds.

The default trace settings for CPU utilization are:

- High CPU Utilization: 90%
- High Track Duration: 5 seconds
- Low CPU Utilization: 75%
- Low Track Duration: 5 seconds

6. Stop tracing:

```
trace shutdown
```

7. View the trace results:

```
show trace file [tail]
```

8. Search trace results for a specific string value, for example, the word error:

```
trace grep [WORD<0-128>]
```

If you use this command and do not specify a string value, you clear the results of a previous search.

9. Stop tracing:

```
trace shutdown
```

Example

```
Switch:1> enable
```

Clear the trace:

```
Switch:1# clear trace
```

Identify the module ID for which you want to use the trace tool:

```
Switch:1# show trace modid-list
```

```
0  - COMMON
1  - SNMP
2  - RMON
3  - PORT_MGR
4  - CHAS_MGR
5  - BRIDGE
6  - HWIF
7  - SIM
8  - CPP
9  - NETDRV
10 - VLAN_MGR
11 - CLI
12 - MAIN
12 - P2IP
12 - RCIP
15 - WEBSRV
16 - ACIF
17 - GBIF
18 - WDT
19 - TDP
20 - MAN_DIAG
21 - MAN_TEST

--More-- (q = quit)
```

Begin the trace operation:

```
Switch:1# trace level 2 3
```

Stop tracing:

```
Switch:1# trace shutdown
```

Save the trace file to the internal flash card for retrieval:

```
Switch:1# save trace
```

Search trace results for a specific string value, for example, the word error:

```
Switch:1# trace grep error
```

Search trace results for a specific string value, for example, MAC address 00-1A-4B-8A-FB-6B:

```
Switch:1# trace grep 00-1A-4B-8A-FB-6B
```

Variable Definitions

The following table defines parameters for the **trace** command.

Variable	Value
<code>grep [WORD<0-128>]</code>	Search trace results for a specific string value, for example, the word error. Performs a comparison of trace messages.
<code>level [<Module_ID>] [<0-4>]</code>	Starts the trace by specifying the module ID and level. <ul style="list-style-type: none"> <code><Module_ID></code> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see <i>Fabric Engine and VOSS Feature Matrix</i>. To see the list of module IDs, use the show trace modid-list command or CLI command completion Help. <code><0-4></code> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.
<code>shutdown</code>	Stops the trace operation.
<code>screen {disable enable}</code>	Enables the display of trace output to the screen.

The following table defines parameters for the **save trace** command.

Variable	Value
<code>file WORD<1-99></code>	Specifies the file name in one of the following formats: <ul style="list-style-type: none"> a.b.c.d:<file>

Use Trace to Diagnose IPv6 Problems

Use trace to observe the status of IPv6 at a certain time.

Before You Begin

- Confirm that trace level 99 is set to a value of 1 before you use trace to diagnose IPv6 problems. Trace level 1 is very terse.

**Caution****Risk of traffic loss**

Using the trace tool inappropriately can cause CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Activate or deactivate the trace for the IPv6 base:

```
trace ipv6 base <disable|enable> <all|debug|error|icmp|info|ipclient|
nbr|pkt|warn> [vrf WORD <1-16>]
```

3. Activate or deactivate the trace for IPv6 forwarding:

```
trace ipv6 forwarding <disable|enable> <all|debug|error|info|pkt|warn>
[vrf WORD <1-16>]
```

4. Activate or deactivate the trace for IPv6 neighbor discovery:

```
trace ipv6 nd <disable|enable> <all|debug|error|info|nbr|pkt|redirect|
warn> [vrf WORD <1-16>]
```

5. Activate or deactivate the trace for IPv6 OSPF:

```
trace ipv6 ospf <disable|enable> <adj|all|config|error|import|info|
lsa|pkt|spf|warn> [vrf WORD <1-16>]
```

6. Activate or deactivate the trace for the IPv6 routing table manager:

```
trace ipv6 rtm <disable|enable> <all|change-list|debug|error|fib|info|
redist|update|warn> [vrf WORD <1-16>]
```

7. Activate or deactivate the trace for IPv6 transport:

```
trace ipv6 transport <disable|enable> <all|common|tcp|udp> [vrf WORD
<1-16>]
```

8. Deactivate the trace to prevent service degradation:

```
trace shutdown
```

```
clear trace
```

Examples

Activate the trace for all the IPv6 forwarding categories:

```
Switch:1>enable
Switch:1#trace ipv6 forwarding enable all
```

Activate the trace for all the IPv6 neighbor discovery categories:

```
Switch:1#trace ipv6 nd enable all
```

Activate the trace for the all IPv6 routing table manager categories:

```
Switch:1#trace ipv6 rtm enable all
```

Activate the trace for all the IPv6 transport categories:

```
Switch:1#trace ipv6 transport enable all
```

Deactivate the trace:

```
Switch:1#trace shutdown
Switch:1#clear trace
Removed 5 files.
```

Variable Definitions

The following table defines parameters for the **trace ipv6** command.

Variable	Value
<i>base</i> <disable enable> <all debug error icmp info ipclient nbr pkt warn>	Enables or disables a specific trace category for IPv6 base.
<i>forwarding</i> <disable enable> <all debug error info pkt warn>	Enables or disables a specific trace category for IPv6 forwarding.
<i>nd</i> <disable enable> <all debug error info nbr pkt redirect warn>	Enables or disables a specific trace category for IPv6 neighbor discovery.
<i>ospf</i> <disable enable> <adj all config error import info lsa pkt spf warn>	Enables or disables a specific trace category for IPv6 OSPF.
<i>rtm</i> <disable enable> <all change-list debug error fib info redist update warn>	Enables or disables a specific trace category for IPv6 routing table manager.
<i>transport</i> <disable enable> <all common tcp udp>	Enables or disables a specific trace category for IPv6 transport.
<i>vrf</i> WORD<1-16>	Specifies the VRF name.

View and Delete Debug Files

Use this procedure to view and delete debug files.

Delete debug files to free space in the internal flash. As a best practice, delete these files to ensure enough space exists in intflash. New debug files do not overwrite old debug files. You must remove the file; otherwise, enough free space might not exist in the intflash to store the core dump if the switch fails or enough space might not exist for you to transfer a new release to the intflash of the switch to upgrade your switch.

The **debug-file remove** command can delete the following types of files:

- core
- archive
- PMEM
- dmalloc
- flrec
- wd_stats

If you want to delete a specific file, you must use the **remove** command. For more information, see *Fabric Engine Command References*.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. View debug files:

```
show debug-file [all][{slot[-slot]}[,...]]
```

3. Delete debug files:

```
debug-file remove [all]
```

4. Enter Privileged EXEC mode:

```
enable
```

5. View core files:

```
show core-files {slot[-slot]}[,...]]
```

Example

The following example shows how you view all debug files for all slots.

```
Switch>show debug-file
```

```
=====
                                Core Files
=====
Directory: /intflash/coreFiles/1
1. File:   core.logServer.20120611084204.1.tar
   Size:   60928 bytes
   Created: Mon Jun 11 08:42:04 2012
2. File:   core.trcServer.20120611084213.1.tar
   Size:   60928 bytes
   Created: Mon Jun 11 08:42:13 2012
3. File:   core.logServer.20120611164647.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:46:48 2012
4. File:   core.trcServer.20120611164652.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:46:52 2012
5. File:   core.dbgServer.20120611164700.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:47:01 2012
6. File:   core.logServer.20120611164740.1.tar
   Size:   64000 bytes
   Created: Mon Jun 11 16:47:41 2012

Remote CP Directory: /intflash/coreFiles/2
1. File:   core.coreManager.x.20120612085548.2.tar
   Size:   1162240 bytes
```

```

Created: Tue Jun 12 08:55:49 2012
2. File: core.coreManager.x.20120612085602.2.tar
Size: 478208 bytes
Created: Tue Jun 12 08:56:02 2012
3. File: core.coreManager.x.20120612085553.2.tar
Size: 1170432 bytes
Created: Tue Jun 12 08:55:56 2012
4. File: core.coreManager.x.20120612085558.2.tar
Size: 1883136 bytes
Created: Tue Jun 12 08:56:00 2012

```

```

=====
Archive Files
=====

```

```
Directory: /intflash/archive/1
```

```

1. File: archive.20120611083021.1.tar
Size: 34296320 bytes
Created: Mon Jun 11 08:30:22 2012
2. File: archive.20120611163454.1.tar
Size: 31108096 bytes
Created: Mon Jun 11 16:34:54 2012
3. File: archive.20120611164354.1.tar
Size: 31792128 bytes
Created: Mon Jun 11 16:43:55 2012
4. File: archive.20120611164507.1.tar
Size: 31881216 bytes
Created: Mon Jun 11 16:45:08 2012

```

```
Remote CP Directory: /intflash//archive/2
```

```

1. File: archive.20120611163507.2.tar
Size: 30903296 bytes
Created: Mon Jun 11 16:35:08 2012
2. File: archive.20120611164408.2.tar
Size: 31314432 bytes
Created: Mon Jun 11 16:44:09 2012
3. File: archive.20120611164521.2.tar
Size: 31367168 bytes
Created: Mon Jun 11 16:45:21 2012

```

```
Directory: /intflash/archive/4
```

```

1. File: archive.20120611163515.4.tar
Size: 4725760 bytes
Created: Mon Jun 11 16:35:18 2012
2. File: archive.20120611164416.4.tar
Size: 5639168 bytes
Created: Mon Jun 11 16:44:20 2012
3. File: archive.20120611164529.4.tar
Size: 5760000 bytes
Created: Mon Jun 11 16:45:33 2012

```

```
Directory: /intflash/archive/SF4
```

```

1. File: archive.20120611163536.SF4.tar
Size: 1550336 bytes
Created: Mon Jun 11 16:35:40 2012
2. File: archive.20120611164436.SF4.tar
Size: 1781248 bytes
Created: Mon Jun 11 16:44:39 2012
3. File: archive.20120611164549.SF4.tar
Size: 1811968 bytes
Created: Mon Jun 11 16:45:53 2012

```



```

=====
PMEM Files
=====
Directory: /intflash/PMEM/4
1. File:    pmem.20120607194023.4.bin.gz
   Size:    571048 bytes
   Created: Thu Jun  7 19:40:23 2012

=====
DMalloc Files
=====

=====
Flrec Files
=====

=====
WdStats Files
=====
Directory: /intflash/wd_stats/4
1. File:    wd_stats.log.backup
   Size:    2311 bytes
   Created: Mon Jun 11 09:25:07 2012

```

The following example shows how to view only core files on the switch.

```

Switch#show core-files
=====
Core Files
=====
Directory: /intflash/coreFiles/1
1. File:    core.1353113115.lifecycle.CP.1.gz
   Size:    139406 bytes
   Created: Fri Nov 16 19:45:15 2012
2. File:    core.cbc-main.x.20121114043335.1.tar
   Size:    14059520 bytes
   Created: Wed Nov 14 04:35:36 2012
3. File:    core.cbc-main.x.20121114045202.1.tar
   Size:    12809728 bytes
   Created: Wed Nov 14 04:54:03 2012
4. File:    core.cbc-main.x.20121114050825.1.tar
   Size:    12638720 bytes
   Created: Wed Nov 14 05:10:26 2012
5. File:    core.cbc-main.x.20121114122506.1.tar
   Size:    13020160 bytes
   Created: Wed Nov 14 12:27:07 2012
6. File:    core.1353336274.lifecycle.CP.1.gz
   Size:    139390 bytes
   Created: Mon Nov 19 09:44:34 2012
7. File:    core.1353319337.lifecycle.CP.1.gz
   Size:    139404 bytes
   Created: Mon Nov 19 05:02:17 2012
8. File:    core.cbc-main.x.20130122182946.1.tar
   Size:    13683712 bytes
   Created: Tue Jan 22 18:32:08 2013
9. File:    core.cbc-main.x.20130220143809.1.tar
   Size:    13969920 bytes
   Created: Wed Feb 20 14:38:10 2013
10. File:   core.cbc-main.x.20130225155025.1.tar
   Size:    13526016 bytes
   Created: Mon Feb 25 15:50:25 2013
11. File:   core.cbc-main.x.20130225155407.1.tar
   Size:    12674560 bytes
   Created: Mon Feb 25 15:54:07 2013

```

Variable Definitions

The following table defines parameters for the **show core-files** command.

Variable	Value
<code>{slot[-slot][,...]}</code>	Displays the core files for the slot that you select.

The following table defines parameters for the **show debug-file** command.

Variable	Value
<code>all</code>	Displays all types of debug files
<code>{slot[-slot][,...]}</code>	Displays debug files for the slot that you select. If you do not select a slot number, the device displays all types of the archived debug files present in a slot by slot basis. If you select a slot number, the device only displays archived files for the slot you select.

The following table defines parameters for the **debug-file remove** command.

Variable	Value
<code>all</code>	Removes debug files in all slots. If you use the <code>all</code> parameter with the remove debug-file command, the device deletes all types of debug files, including the latest debug files. If you do not use the <code>all</code> parameter, the device retains the latest file of each type.

Configure Port Mirroring

Use port mirroring to aid in diagnostic and security operations.

About This Task

Use port mirroring to make a copy of a traffic flow and send that copy to a device for analysis, for example, for diagnostic sniffing. Use the mirror to see the packets in the flow without breaking into the physical connection to place a packet onto the sniffer inline. You can also use port mirroring for security. You can send flows to inspection engines for post processing.

Connect the sniffer (or other traffic analyzer) to the output port you specify in this procedure.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create a port mirroring instance:

```
mirror-by-port <1-479> in-port {slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]} {monitor-mlt <1-512>| out-port {slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}}
```

3. Create an I-SID mirroring instance:

```
mirror-by-port <1-479> [in-port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]} monitor-isid-offset <1-1000> [mode <rx|tx|both>][qos <qos-level>]]
```

4. Configure the mode:

```
mirror-by-port <1-479> mode <both|rx|tx>
```

**Note**

- When you configure `tx` mode port mirroring on T-UNI and SPBM network-to-network interface (NNI) ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports and the system displays it on the mirror destination port, although they do not egress the mirror source port. This is because `tx` mode port mirroring happens on the mirror source port before the source port squelching logic drops the packets at the egress port.
- The available four mirroring resources are shared between Fabric RSPAN and regular port mirroring, and are allocated based on the mode configured, Ingress (`rx`) or Egress (`tx`). Each configured mode occupies one mirroring resource, but when you configure the mode as `both`, it occupies two mirroring resources (one for Rx and one for Tx).
- You cannot configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

5. Enable the mirroring instance:

```
mirror-by-port <1-479> enable
```

6. Modify existing mirroring entries as required:

```
mirror-by-port mirror-port <1-479> {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}
```

OR

```
mirror-by-port monitor-mlt <1-479> <1-512>
```

OR

```
mirror-by-port monitor-port <1-479> {slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]}
```

**Note**

Before you can modify an existing entry, you must disable the entry: no
 mirror-by-port <1-479> enable.

7. Modify QoS value for Fabric RSPAN mirroring session:

```
mirror-by-port <1-479> qos <0-5>
```

8. Verify the configuration:

```
show mirror-by-port
```

Example

Port mirroring configuration:

```
Switch:1> enable
Switch:1# configure terminal
```

Create the port mirroring instance:

```
Switch:1(config)# mirror-by-port 8 in-port 1/15 out-port 1/1
```

The analyzer connects to port 1/1.

Disable the entry:

```
Switch:1(config)# no mirror-by-port 8 enable
```

Mirror both ingress and egress traffic passing through port 1/16:

```
Switch:1(config)# mirror-by-port 8 mode both
```

Enable mirroring for the instance:

```
Switch:1(config)# mirror-by-port 8 enable
```

Fabric RSPAN configuration:

```
Switch:1> enable
Switch:1# configure terminal
```

Create the Fabric RSPAN mirroring instance:

```
Switch:1(config)#mirror-by-port 3 in-port 1/3 monitor-isid-offset 3 mode both qos 3
```

Disable the entry:

```
Switch:1(config)# no mirror-by-port 3 enable
```

Mirror the egress traffic passing through port 1/3:

```
Switch:1(config)# mirror-by-port 3 mode tx
```

Enable Fabric RSPAN for the instance:

```
Switch:1(config)# mirror-by-port 3 enable
```

The sample command output in the following example does not necessarily reflect the preceding examples.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#show mirror-by-port
```

```
=====
Diag Mirror-By-Port
```

ID	MIRRORED_PORT	MIRRORING_DEST	ENABLE	MODE	REMOTE-MIRROR VLAN-ID	DSCP	TTL
1	1/1	2/1	true	both	0 0		64
2	1/2	2/2	true	rx	0 0		64
3	1/3	2/3	true	tx	0 0		64
4	1/4	2/4	true	both	0 0		64

Variable Definitions

The following table defines parameters for the **mirror-by-port** command.

Variable	Value
<1-479>	Specifies the entry ID.
<i>enable</i>	Enables or disables a mirroring instance already created in the mirror-by-port table.
<i>in-port</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]} {monitor-mlt <1-512> <i>out-port</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]}	<p>Creates a new mirror-by-port table entry.</p> <ul style="list-style-type: none"> <i>in-port</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]} specifies the mirrored port. <i>monitor-mlt</i> <1-512> specifies the mirroring MLT ID from 1-512. <p>Note: Exception: This parameter does not apply to 7830 Series.</p> <ul style="list-style-type: none"> <i>out-port</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]} specifies the mirroring port.
<i>mirror-port</i> <1-479> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]}	<p>Modifies the mirrored port.</p> <p>Before you can modify an existing entry, you must disable the entry: no mirror-by-port <1-479> enable.</p>
<i>monitor-ip</i> <1-479> {A.B.C.D} [dscp <0-63>] [ttl <2-255>]	Creates a mirroring instance for Layer 3 mirroring. The destination must be an IP address {A.B.C.D}. The default DSCP is 0 and the default TTL is 255.
<i>monitor-mlt</i> <1-479> <1-512> Note: Exception: This parameter does not apply to 7830 Series.	<p>Modifies the monitoring MLT. <1-512> specifies the mirroring MLT ID.</p> <p>Before you can modify an existing entry, you must disable the entry: no mirror-by-port <1-479> enable.</p>
<i>monitor-port</i> <1-479> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all][all]}	<p>Modifies the monitoring ports.</p> <p>Before you can modify an existing entry, you must disable the entry: no mirror-by-port <1-479> enable.</p>

Variable	Value
<code>monitor-vlan <1-479> <1-4059></code>	<p>Modifies the monitoring VLAN.</p> <p>Before you can modify an existing entry, you must disable the entry: <code>no mirror-by-port <1-479> enable</code>.</p> <p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p>
<code>mode <both rx tx></code>	<p>Configures the mirroring mode. The default is rx.</p> <ul style="list-style-type: none"> • <i>both</i> mirrors both egress and ingress packets. • <i>rx</i> mirrors ingress packets. • <i>tx</i> mirrors egress packets.
<code>monitor-isid-offset <1-1000></code> Note: Exception: This parameter does not apply to 7830 Series.	<p>Specifies the offset ID that is mapped to the actual monitor I-SID where packets are mirrored.</p> <p>Monitor I-SID = base monitor I-SID + offset ID. The base monitor I-SID is 16776000.</p>
<code>qos <0-5></code>	<p>Specifies the Quality of Service (QoS) profiles for the system. Monitoring I-SID supports six different QoS levels, each QoS level can be configured individually. Default value is 1.</p>

Display Mirror Resource Usage

About This Task

Mirror resources can be consumed by internal processes that are not easily identified, which can lead to unexpected resource shortages. Use the following procedure to display mirror resource usage.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display mirror resource usage:

```
show mirror-resources
```

Example

```
Switch:1#show mirror-resources
<Switch> Mirror Resource Manager stats:
      2 unique mirror destinations allocated, system max is 4
```

```
-----
Mirror      | Mirror      | Total | Mirror
```

Destination	Direction	Refs	Users
Port 1/2	Egress	1	MBP ID 1
Port 1/2	Ingress	1	MBP ID 1

Function execution completed successfully

Configure Global Mirroring Actions with an ACL

Configure the global action to mirror packets that match an access control list (ACL).

Before You Begin

- The ACL exists.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

- Configure the global action for an ACL:

```
filter acl set <1-2048> global-action {monitor-dst-mlt <1-512>|
monitor-dst-ports {slot/port[-slot/port][,...][slot/all][all]}
```

Example

Configure the global action for an ACL:

```
Switch:1> enable
Switch:1#configure terminal
Switch:1(config)#filter acl set 200 global-action monitor-dst-mlt 20
```

Variable Definitions

The following table defines parameters for the **filter acl set** command.

Variable	Value
<1-2048>	Specifies an ACL ID from 1–2048.
<i>default-action</i> <deny permit>	Specifies the global action to take for packets that do not match an ACL.
<i>global-action</i> { <i>monitor-dst-mlt</i> <i>PT MLT</i> <1-512> <i>monitor-dst-ports</i> { <i>slot/port</i> [/ <i>sub-port</i>][- <i>slot/port</i> [/ <i>sub-port</i>]][,...][<i>slot/all</i>][<i>all</i>]}	Specifies the global action to take for matching ACLs: <ul style="list-style-type: none"> monitor destination MLT—Configures mirroring to a destination MultiLink Trunking (MLT) group. monitor destination ports—Configures mirroring to a destination port or ports.

Configure ACE Actions to Mirror

Configure actions to use filters for flow-based mirroring.

Before You Begin

- The access control entry (ACE) exists.

About This Task

If you use the mirror action, ensure that you specify the mirroring destination: MLTs or ports.

Procedure

- Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

- Configure actions for an ACE:

```
filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-mlt  
<1-512>
```

OR

```
filter acl ace action <1-2048> <1-2000> {permit|deny} monitor-dst-  
ports {slot/port[-slot/port][,...][slot/all][all]}
```

- Ensure the configuration is correct:

```
show filter acl action [<1-2048>] [<1-2000>]
```

Example

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#filter acl ace action 901 1 permit monitor-dst-mlt 5
```

Variable Definitions

The following table defines parameters for the **filter acl ace action** command.

Variable	Value
<i>1-2048</i>	Specifies the ACL ID from 1-2048
<i>1-2000</i>	Specifies the ACE ID from 1-2000.
<i>monitor-dst-mlt <1-512></i>	Configures mirroring to a destination MLT group.
<i>monitor-dst-ports {slot/ port[/sub-port][-slot/ port[/sub-port]][,...] [slot/all][all]}</i>	Configures mirroring to a destination port or ports.
<i>{permit deny}</i>	Configures the action mode for security ACEs. The default value is permit.

Clear ARP Information for an Interface

Clear the Address Resolution Protocol (ARP) cache as part of ARP problem resolution procedures.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Clear ARP information:

```
clear ip arp interface gigabitethernet {slot/port[/sub-port] [-slot/
port[/sub-port]][,...][slot/all][all]}
```

OR

```
clear ip arp interface vlan <1-4059>
```

Example

```
Switch:1>enable
Switch:1#clear ip arp interface gigabitethernet 1/1
```

Variable Definitions

The following table defines parameters for the **clear ip arp interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]][,...][slot/all][all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Flush Routing, MAC, and ARP Tables for a Port

Flush or clear the routing tables for administrative and troubleshooting purposes. The clear and flush commands perform the same function; they remove the contents of the table.

Procedure

- 1. Enter GigabitEthernet Interface Configuration mode:
`enable`

`configure terminal`

`interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}`



Note If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- 2. Flush IP routing tables by port:
`action flushIp`
- 3. Flush the MAC address tables by port:
`action flushMacFdb`
- 4. Flush ARP tables by port:
`action flushArp`
- 5. Flush all tables with one command:
`action flushAll`
- 6. Exit to Global Configuration mode:
`exit`
- 7. Clear a routing table for a port:
`clear ip route gigabitethernet {slot/port[sub-port]}`

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitEthernet 1/1
Switch:1(config-if)#action flushAll
Switch:1(config-if)#exit
Switch:1(config)#clear ip route gigabitethernet 1/1
```

Variable Definitions

The following table defines parameters for the `clear ip route gigabitethernet` command.

Variable	Value
<code>{slot/port[/sub-port]}</code>	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Flush Routing, MAC, and ARP Tables for a VLAN

Flush or clear the routing tables for administrative and troubleshooting purposes. The clear and flush commands perform the same function; they remove the contents of the table.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Flush IP routing tables by VLAN:
vlan action <1-4059> flushIp
3. Flush the MAC address tables by VLAN:
vlan action <1-4059> flushMacFdb
4. Flush ARP tables by VLAN:
vlan action <1-4059> flushArp
5. Flush all tables with one command:
vlan action <1-4059> all
6. Clear a routing table for a VLAN:
clear ip route vlan <1-4059>

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config-if)#vlan action 123 all
Switch:1(config)#clear ip route vlan 123
```

Variable Definitions

The following table defines parameters for the **vlan action** and **clear ip route vlan** commands.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

*Ping an IP Device***About This Task**

Ping a device to test the connection between the switch and another network device. After you ping a device, the switch sends an Internet Control Message Protocol (ICMP)

packet to the target device. If the device receives the packet, it sends a ping reply. After the switch receives the reply, a message displays that indicates traffic can reach the specified IP address. If the switch does not receive a reply, the message indicates the address does not respond.

If you originate the **ping** or **traceroute** operation from a device other than the switch to a routing IP interface associated with a VRF, the operation can fail if you use large packet sizes. For these situations, do not use packet sizes larger than the following:

- Ping for VRF routes: 1472 bytes
- Traceroute for VRF routes: 1444 bytes



Note

Exceptions:

- The CLIP Segmented Management Instance can support large packets for VRF routes.
- If you originate the **ping** or **traceroute** operation from the switch, you can configure large datasizes and the operation functions properly.

You can specify a management instance ID to use the correct source for the outgoing ICMP ECHO request packet.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Ping an IP network connection:

```
ping WORD<0-256> [-d] [-I <1-60>] [-s] [-t <1-120>] [count <1-9999>]
[datasize <28-9216|28-51200>] [grt] [interface gigabitEthernet {slot/
port[sub-port]}| tunnel <1-2000> | vlan <1-4059>] [scopeid <1-9999>]
[source WORD<1-256>] [vrf WORD<1-16>]
```

3. Ping an IP network connection using a Global Routing Table (GRT):

```
ping WORD<0-256> [-d] [-I <1-60>] [-s] [-t <1-120>] [count <1-9999>]
[datasize <28-9216|28-51200>] grt [interface gigabitEthernet {slot/
port[sub-port]}| tunnel <1-2000> | vlan <1-4059>] [source WORD<1-246>]
```

4. Ping a network connection using a Segmented Management Instance:

```
ping WORD<0-256> [-s] [-t <1-120>] [count <1-9999>] [datasize
<28-9216|28-51200>] mgmt [clip | oob | vlan]
```



Note

If you do not use the *mgmt* parameter, the **ping** command uses the IP routing stack to initiate the ping request.

If you ping a device using a management CLIP, the ping source IP address is configured as the management CLIP IP address. If you ping a device using a management VLAN, the ping source IP address is configured as the management VLAN IP address.

5. (Optional) Ping a network connection using a Network Service Probe interface:

```
ping WORD<0-256> service-probe
```

Examples

Ping an IP device from a GRT VLAN IP interface:

```
Switch:1#ping 192.0.2.16 grt interface vlan 1
192.0.2.16 is alive
```

Ping a device using the management routing table:

```
Switch:1#ping 192.0.2.12 mgmt
```

Ping a device using a management CLIP:

```
Switch:1#ping 192.0.2.12 mgmt clip
```

Ping an IP device using a management VLAN:

```
Switch:1#ping 192.0.2.12 mgmt vlan
```

Ping an IP device using the configured Network Service Probe interface:

```
Switch:1#ping 192.0.2.6 service-probe
Sending ping in context grt
PING 192.0.2.6 (192.0.2.6) from 192.0.2.1 : 64(92) bytes of data.
72 bytes from 192.0.2.6: icmp_seq=1 ttl=255 time=1.11 ms

--- 192.0.2.6 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.111/1.111/1.111/0.000 ms
```

Variable Definitions

The following table defines parameters for the **ping** command.

Variable	Value
<i>count</i> <1-9999>	Specifies the number of times to ping. The default is 1.
<i>-d</i>	Configures the ping debug mode. This variable detects local software failures (ping related threads creation or write to sending socket) and receiving issues (icmp packet too short or wrong icmp packet type). This parameter does not apply if you use the mgmt [clip oob vlan] parameter.
<i>datasize</i> <28-9216 28-51200>	Specifies the size of ping data sent in bytes. The datasize for IPv4 addresses is 28-9216. The datasize for IPv6 addresses is 28-51200. The default is 64.
<i>grt</i>	Specifies the ping in Global Routing Table context.
<i>-I</i> <1-60>	Specifies the interval between transmissions in seconds.

Variable	Value
<code>interface gigabitEthernet {slot/port[sub-port]} tunnel <1-2000> vlan <1-4059></code>	<p>Specifies the outgoing interface.</p> <p>Additional ping interface parameters:</p> <ul style="list-style-type: none"> • <code>gigabitEthernet {slot/port[sub-port]}</code>: <code>gigabitEthernet port</code> • <code>tunnel</code>: tunnel ID as a value from 1 to 2000 • <code>vlan</code>: Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <code>vrf-scaling</code> and <code>spbm-config-mode</code> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. <p>This parameter does not apply if you use the <code>mgmt [clip oob vlan]</code> parameter.</p>
<code>mgmt [clip oob vlan]</code> Note: Exception: <code>oob</code> only applies to hardware with a dedicated physical management interface.	<p>Specifies the Segmented Management Instance as the source for the outgoing ICMP ECHO packet. The packet goes out this specific interface only.</p> <p>If you do not specify the management interface type, the ping command uses the management routing table to determine the best management interface and selects the source IP based on the egress management interface. This also includes the <code>Mgmt-sdwan</code> interface if you need to use the ping command for the SD-WAN Appliance. For more information, see Fabric Extend (FE) States on page 32.</p>
<code>-s</code>	Configures the continuous ping at the interval rate defined by the <code>[-I]</code> parameter or until you enter a Ctrl + C keystroke.
<code>scopeid <1-9999></code>	<p>Specifies the circuit ID for IPv6.</p> <p>This parameter does not apply if you use the <code>mgmt [clip oob vlan]</code> parameter.</p>
<code>service-probe</code>	Launches a ping request from the configured Network Service Probe interface.
<code>source WORD<1-256></code>	<p>Specifies the source IP address for the ping command.</p> <p>This parameter does not apply if you use the <code>mgmt [clip oob vlan]</code> parameter.</p>

Variable	Value
<code>-t <1-120></code>	Specifies the no-answer timeout value in seconds. The default is 5.
<code>WORD<0-256></code>	Specifies the host name or IPv4 (a.b.c.d) or IPv6 (x:x:x:x:x:x:x) address.
<code>vrf WORD<1-16></code>	Specifies the virtual router and forwarder (VRF) name. This parameter does not apply if you use the <code>mgmt [clip oob vlan]</code> parameter.

Run a Traceroute Test

Use traceroute to determine the route packets take through a network to a destination.

About This Task

If you originate the **ping** or **traceroute** operation from a device other than the switch to a routing IP interface associated with a VRF, the operation can fail if you use large packet sizes. For these situations, do not use packet sizes larger than the following:

- Ping for VRF routes: 1472 bytes
- Traceroute for VRF routes: 1444 bytes



Note

Exceptions:

- The CLIP Segmented Management Instance can support large packets for VRF routes.
- If you originate the **ping** or **traceroute** operation from the switch, you can configure large datasizes and the operation functions properly.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Run a traceroute test:

```
traceroute WORD<0-256> [<1-1176>] [-m <1-255>] [-p <1-65535>] [-q  
<1-255>] [-v] [-w <1-255>] [grt [source]] [mgmt [clip | oob | vlan]]  
[source <WORD 1-256>] [vrf <WORD 1-16>]
```

3. Run a traceroute test using a Segmented Management Instance:

```
traceroute WORD<0-256> [<1-1176>] [-m <1-255>] [-p <1-65535>] [-q  
<1-255>] [-w <1-255>] [grt [source]] mgmt [<clip | oob | vlan>]
```



Note

If you do not use the `mgmt` parameter, the **traceroute** command uses the IP routing stack to initiate the traceroute request.

Examples

Run a traceroute test with a probe packet size of 200 and a max time to live of 60:

```
Switch:1>enable
Switch:1#traceroute 192.0.2.33 200 -m 60
```

Run a traceroute test for an IPv6 address:

```
Switch:1#traceroute 2001:db8::1
```

Run a traceroute test using the management routing table:

```
Switch:1#traceroute 192.0.2.12 mgmt
```

Run a traceroute test using a management CLIP:

```
Switch:1#traceroute 192.0.2.12 mgmt clip
```

Run a traceroute test using a management VLAN:

```
Switch:1#traceroute 192.0.2.12 mgmt vlan
```

Variable Definitions

The following table defines parameters for the **traceroute** command.

Variable	Value
<code>-m <1-255></code>	Specifies the maximum time-to-live (TTL).
<code>-p <1-65535></code>	<p>Specifies the base UDP port number. The default is 33434.</p> <p>Note: If the traceroute action is directed to an IPv6 host address, Linux increments the UDP port on a per-TTL basis. For an IPv4 host address, Linux increments the UDP port on a per-probe basis. Because the traceroute command sends a default of three probes, three incrementing ports will be sent for an IPv4 host address. If you use the <code>-p</code> parameter with a value greater than 65533, the traceroute command fails for an IPv4 host address because the maximum port number, 65535, is exceeded.</p> <p>If you send a traceroute probe into the device through the Segmented Management Instance or any routing interface, you must use the default UDP port range of 33434 to 33534. Using other ports will fail.</p>
<code>-q <1-255></code>	Specifies the number of probes per TTL.
<code>-v</code>	<p>Specifies verbose mode (detailed output).</p> <p>This parameter does not apply if you use the <code>mgmt [clip oob vlan]</code> parameter.</p>
<code>-w <1-255></code>	Specifies the wait time for each probe.

Variable	Value
<1-1176>	Specifies the size of the probe packet.
<i>grt</i>	Specifies the traceroute command is executed in grt context.
<i>mgmt</i> [<clip oob vlan>]	Specifies the Segmented Management Instance as the source for the outgoing packet. The packet goes out this specific interface only. If you do not specify the management interface type, the traceroute command uses the management routing table to determine the best management interface and selects the source IP based on the egress management interface.
<i>source</i> <WORD 1-256>	Specifies the source IP address. This parameter does not apply if you use the mgmt [clip oob vlan] parameter.
<i>WORD</i> <0-256>	Specifies the destination IPv4 or IPv6 address, or hostname.
<i>vrf</i> <WORD 1-16>	Specifies the VRF instance by VRF name. This parameter does not apply if you use the mgmt [clip oob vlan] parameter.

Show SNMP Logs

Show the full SNMP logs. SNMP logs display the alarms and events that have been registered on the device.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Show the logs:
show fulltech file *WORD*<1-99>

Variable Definitions

The following table defines parameters for the **show fulltech** command.

Variable	Value
<i>file</i> <i>WORD</i> <1-99>	This variable represents the log file to be opened and displayed. It is displayed in the following format: • /intflash/<file>

Use Trace to Examine IS-IS Control Packets

Use trace as a debug tool to examine the code flow and Intermediate-System-to-Intermediate-System (IS-IS) control packets.

About This Task

The default trace settings for CPU utilization are:

- High CPU Utilization: 90%
- High Track Duration: 5 seconds
- Low CPU Utilization: 75%
- Low Track Duration: 5 seconds

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Identify the module ID for which you want to use the trace tool:
`show trace modid-list`
3. Clear any previous trace:
`clear trace`
4. Begin the trace operation:
`trace level [<Module_ID>] [<0-4>]`
5. Wait approximately 30 seconds.
6. Stop tracing:
`trace shutdown`
7. View the trace results:
`trace screen enable`



Important

If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

8. (Optional) Save the trace file.
`save trace [file WORD<1-99>]`

If you do not specify a file name, the file name is systrace.txt.

9. (Optional) Display trace lines saved to a file:
`show trace file [tail]`
10. (Optional) Search trace results for a specific string value:
`trace grep [WORD<0-128>]`

If you use this command and do not specify a string value, you clear the results of a previous search.

Example

Clear prior trace information, and then enable the display of trace output to the screen. Begin the trace operation for received IS-IS hello packets:

```
Switch:1#clear trace
Switch:1>enable
```

```
Switch:1#trace screen enable
Switch:1#Screen tracing is on
Switch:1#trace level 267 2
```

Variable Definitions

The following table defines parameters for the **show trace** command.

Variable	Value
<i>auto</i>	Displays the current configuration for the automatic trace function.
<i>file [tail]</i>	Displays the trace results saved to a file.
<i>level</i>	Displays the current trace level for all modules.
<i>modid-list</i>	Specifies the module ID list.

The following table defines parameters for the **trace** command.

Variable	Value
<i>grep [WORD<0-128>]</i>	Specifies the search keyword. You can use a specific MAC address. You can search for errors, using the command, trace grep error .
<i>level [<Module_ID>] [<0-4>]</i>	<p>Starts the trace by specifying the module ID and level.</p> <p><i><Module_ID></i> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see <i>Fabric Engine and VOSS Feature Matrix</i>. To see the list of module IDs, use the show trace modid-list command or CLI command completion Help.</p> <p>0-4 specifies the trace level:</p> <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
<i>shutdown</i>	Stops the trace operation.
<i>screen {disable enable}</i>	<p>Enables or disables the display of trace output to the screen.</p> <p>Important:</p> <p>Avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.</p>

The following table defines parameters for the **save trace** command.

Variable	Value
<i>file</i> WORD<1-99>	<p>Specifies the file name in one of the following formats:</p> <ul style="list-style-type: none"> a.b.c.d: <file> <p>WORD<1-99> is a string of 1-99 characters.</p> <p>Note: If you do not specify a file name, the file name is systrace.txt.</p>

View Detailed TLV Information for IS-IS

About This Task

Use the following procedure to view the detailed information about metric type of IS-IS routes in TLVs in Link State Packets (LSP).

Procedure

1. Display the detail view of TLV 135:
show isis lsdb tlv 135 detail
2. Display the detail view of TLV 184:
show isis lsdb tlv 184 detail

Examples

Viewing the metric type of IS-IS route in TLV 135:

```
Switch:1#show isis lsdb tlv 135 detail

=====
                        ISIS LSDB (DETAIL)
=====
-----
Level-1 LspID: 4072.0000.0000.00-02      SeqNum: 0x00000009      Lifetime: 1110
      Chksum: 0x31ce  PDU Length: 46
      Host_name: evp4k
      Attributes:      IS-Type 1
TLV:135 TE IP Reachability: 2
      Metric: 1 Metric Type:Internal  Prefix Length: 32  RouteTag = 0
(0x0)
      UP/Down Bit: FALSE                Sub TLV Bit: FALSE
      IP Address: 15.15.15.72
      Metric: 1 Metric Type:External  Prefix Length: 24  RouteTag = 0
(0x0)
      UP/Down Bit: FALSE                Sub TLV Bit: FALSE
      IP Address: 192.0.2.5
```

Viewing the metric type of IS-IS route in TLV 184

```
Switch:1#show isis lsdb tlv 184 detail

=====
                        ISIS LSDB (DETAIL)
=====
```

```

-----
Level-1 LspID: 4072.0000.0000.00-03      SeqNum: 0x00000008      Lifetime: 1103
      Chksum: 0x3ce6  PDU Length: 72
      Host_name: evp4k
      Attributes:      IS-Type 1
TLV:184 SPBM IPVPN Reachability:
      Vrf ISID:100
      Metric:1 Metric Type:External      Prefix Length:32 RouteTag = 0
(0x0)
      IP Address: 192.0.2.3
      Metric:1 Metric Type:Internal      Prefix Length:32 RouteTag = 0
(0x0)
      IP Address: 192.0.2.72

```

View Summarized TLV Information for IS-IS

About This Task

Use the following procedure to view the summarized information about the metric type of IS-IS routes in TLVs. You can also view the metric type of the prefix.

Procedure

Display the summarized view of TLVs 135 and 184:

```
show isis lsdb ip-unicast
```

Example

Display the summarized view of TLVs:

```

evp4k:1#show isis lsdb ip-unicast
=====
                        ISIS IP-UNICAST-ROUTE SUMMARY
=====
I-SID      ADDRESS      PREFIX      METRIC      METRIC      TLV      LSP      HOST      AREA      ROUTE-TAG
LENGTH    TYPE          TYPE        FRAG      NAME
-----
-          192.0.2.5      32          1          Internal   135      0x2      evp4k      HOME      0
-          192.0.2.20     24          1          External   135      0x2      evp4k      HOME      0
100        192.0.2.15     32          1          External   184      0x3      evp4k      HOME      0
100        192.0.2.5      32          1          Internal   184      0x3      evp4k      HOME      0
-          192.0.2.4      32          1          Internal   135      0x2      esp1       HOME      0
=====
5 out of 5 Total Num of Entries

```

Configure I-SID Monitoring

Use the following procedure to configure Fabric RSPAN (Mirror to I-SID) on the Backbone Edge Bridge (BEB) connected to the monitor station.

**Note**

You cannot configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

**Note**

If you change the egress port or egress MLT for a particular session using a separate CLI command, it overwrites the existing egress port list or egress MLT.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a monitor by I-SID entry:

**Note**

If you use the Extreme Integrated Application Hosting (IAH) port 1/s1 as the analyzer port on the monitoring BEB for remote mirroring, you must associate it to VLAN ID 4091.

```
monitor-by-isid <1-1000> [monitor-isid-offset <1-1000> {egress-mlt
<1-512> | egress-ports {slot/port[-slot/port][,...][slot/all][all]} }
[map-to-vid <1-4093>]]
```

3. Configure map to VLAN ID:

```
monitor-by-isid <1-1000> map-to-vid <1-4093>
```

4. Configure egress MLT:

```
monitor-by-isid <1-1000> egress-mlt <1-512>
```

5. Configure egress port:

```
monitor-by-isid <1-1000> egress-ports {slot/port[-slot/port][,...]
[slot/all][all]}
```

6. Enable monitoring by I-SID entry:

```
monitor-by-isid <1-1000> enable
```

**Note**

Disable the entries (egress ports, MLT, and VLAN ID) to modify or remove parameters in the existing configuration.

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# monitor-by-isid 1 monitor-isid-offset 1 egress-port 1/6
Switch:1(config)# monitor-by-isid 2 monitor-isid-offset 2 egress-port 1/7 map-to-vid 200
Switch:1(config)# monitor-by-isid 3 monitor-isid-offset 3 egress-port 1/7 map-to-vid 201
Switch:1(config)# monitor-by-isid 2 egress-port 1/8
Switch:1(config)# monitor-by-isid 1 monitor-isid-offset 1000 egress-ports 1/1 egress-mlt
16 map-to-vid 1000
Switch:1(config)# monitor-by-isid 7 monitor-isid-offset 7 egress-mlt 2 map-to-vid 203
Switch:1(config)# monitor-by-isid 2 egress-mlt 3
```

Variable Definitions

The following table defines parameters for the **monitor-by-isid** command.

Variable	Value
<1-1000>	Specifies the monitoring session.
<i>monitor-isid-offset</i> <1-1000> Note: Exception: This parameter does not apply to 7830 Series.	Specifies the offset ID that is mapped to the actual monitor I-SID where packets are mirrored. Monitor I-SID = Base monitor I-SID + Offset ID. The base monitor I-SID is 16776000.
<i>egress-ports</i> {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}	Specifies the port to which the analyzers connect.
<i>egress-mlt</i> <1-512>	Specifies the MLT to which the analyzers connect.
<i>map-to-vid</i> <1-4093>	Maps the mirrored packet to a specified VLAN ID for analysis. This parameter is optional. Note: If you use the Extreme Integrated Application Hosting (IAH) port 1/s1 as the analyzer port on the monitoring BEB for remote mirroring, you must associate it to VLAN ID 4091.

Display I-SID Monitoring Diagnostics

Use the following procedure to display the **monitor-by-isid** entries on the monitoring BEB.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Enter the following command:

```
show monitor-by-isid WORD<1-1024>
```

Example

```
Switch:1>show monitor-by-isid 3
=====
                Diag Monitor-By-ISID
=====
ID MONITOR_ISID ISID_OFFSET  EGRESS_PORTS  EGRESS_MLT  MAP_TO_VLAN  ENABLE
-----
3      16776000      1           1/4, 1/5         1           999         true
-----
```

Variable Definitions

The following table defines parameters for the **show monitor-by-isid** command.

Variable	Value
<i>WORD<1-1024></i>	Specifies the session ID list ranging from 1 to 1000.

Display I-SID Mirroring Statistics

Use the following procedure to display the statistics of the number of packets mirrored into I-SID on the mirroring BEB.

Procedure

1. Enter Global Configuration mode:

```
enable

configure terminal
```
2. Enter the following command:

```
show isid-mirroring stats [monitor-isid-offset WORD<1-1024>]
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# show isid-mirroring stats monitor-isid-offset 1
=====
                Mirror Statistics Info
=====
ISID          ISID_OFFSET  PACKETS
-----
16776000      1            100
-----
```

Variable Definitions

The following table defines parameters for the **show isid-mirroring stats** command.

Variable	Value
<i>monitor-isid-offset WORD<1-1024></i>	Specifies the offset ID mapped to monitor the I-SID.
Note: Exception: This parameter does not apply to 7830 Series.	The offset ID ranges from 1 to 1000.

Clear Fabric RSPAN (Mirror to I-SID) Statistics

Use the following procedure to clear Fabric RSPAN (Mirror to I-SID) statistics of packets mirrored into the specified mirroring I-SID or all mirroring I-SIDs on the BEB.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enter the following command:

```
clear isid-mirroring stats monitor-isid-offset WORD<1-1024>
```



Note

You must use this command on the Mirroring BEB to clear the statistics of packets mirrored into I-SID.

Example

```
Switch:1>enable
Switch:1#configure terminal
```

Clear all Fabric RSPAN statistics:

```
Switch:1(config)# clear isid-mirroring stats
```

Clear all Fabric RSPAN (Mirror to I-SID) statistics of packets mirrored into the specified mirroring I-SID

```
Switch:1(config)# clear isid-mirroring stats monitor-isid-offset 1
```

Variable Definitions

The following table defines parameters for the **clear isid-mirroring stats** command.

Variable	Value
<i>monitor-isid-offset WORD<1-1024></i>	Specifies the offset ID that is mapped to the actual monitor I-SID where packets are mirrored.
Note: Exception: This parameter does not apply to 7830 Series.	Monitor I-SID = base monitor I-SID + offset ID.
	The range of the <i>monitor-isid-offset</i> is 1 to 1000.
	The base monitor I-SID is 16776000.

Configure a Network Service Probe

Configure a Network Service Probe to troubleshoot network connectivity issues.

Before You Begin

- Ensure that the port-based VLAN exists on the switch.

About This Task

You can create only one service probe on the switch. If you delete the VLAN, the Network Service Probe configuration is also removed.

If you do not configure a static IP address for the Network Service Probe or assign a static default gateway, the switch uses DHCP to acquire an IP address or dynamically learn the gateway address. The DHCP Server and relay agent must be non-local to the switch. The gateway must be non-local to the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create the Network Service Probe and associate it with a VLAN ID:

```
service-probe vlan <1-4059>
```
3. (Optional) Create a static IP address and subnet mask for the Network Service Probe:

```
service-probe ip static {<A.B.C.D/X> | (A.B.C.D) (A.B.C.D) }
```
4. (Optional) Create a static IP association with the default gateway:

```
service-probe default-gateway {A.B.C.D}
```
5. Enable the Network Service Probe:

```
service-probe enable
```
6. View the Network Service Probe configuration:

```
show service-probe
```

Example

Create, and enable, a Network Service Probe with a static IP configuration, and then view the configuration:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#service-probe vlan 15
Switch:1(config)#service-probe ip static 15.0.0.1 255.255.255.0
Switch:1(config)#service-probe enable
Switch:1>show service-probe
```

Service Probe Info

```
=====
Vlan: 15
Administrative State: enabled
MAC Address: 4e:bf:2d:0c:1f:bd
IP Address: 192.0.2.11
IP Mask: 255.255.255.0
IP Address Config: dynamic
Default Gateway: 192.0.2.1
Default Gateway State: ACTIVE
Primary dynamic DNS server
=====
IP address: 198.51.100.2
```

```

Query status: Fail

Secondary dynamic DNS server
=====
IP address:    203.0.113.3
Query status: Success

```

NEW! Query Domain Name Server using Service Probe

About This Task

Perform this procedure to query the DNS servers that the switch receives and stores from the DHCP server.



Note

This only applies for dynamic IP configuration on Service Probe.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Query the DNS servers to perform domain resolution tests:

```
service-probe query dns [primary-dynamic | secondary-dynamic | tertiary-dynamic]
```

NEW! Variable Definitions

The following table defines parameters for the **service-probe query dns** command.

<i>primary-dynamic</i>	Specifies the primary DNS server learned through DHCP. If you do not specify a DNS server parameter, the switch uses primary-dynamic by default.
<i>secondary-dynamic</i>	Specifies the secondary DNS server learned through DHCP. This option sends the DNS query to the second DHCP-assigned DNS server.
<i>tertiary-dynamic</i>	Specifies the tertiary DNS server learned through DHCP. This option sends the DNS query to the third DHCP-assigned DNS server.

Software Troubleshooting Tool Configuration Using EDM

Use the tools described in this section to perform troubleshooting procedures using Enterprise Device Manager (EDM).

Flush Routing Tables by VLAN

About This Task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use EDM to flush the routing tables by VLAN or flush them by port. Perform this procedure to flush the IP routing table for a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. In the **Vlan Operation Action** column, select a flush option.
In a VLAN context, all entries associated with the VLAN are flushed. You can also flush the Address Resolution Protocol (ARP) entries and IP routes for the VLAN.
5. Select **Apply**.

Flush Routing Tables by Port

About This Task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use EDM to flush the routing tables by VLAN or flush them by port. Perform this procedure to flush the IP routing table for a port.

Procedure

1. On the Device Physical View, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. In the **Action** section, select **flushAll**.
In a port context, all entries associated with the port are flushed. You can flush the ARP entries and IP routes for a port. After you flush a routing table, it is not automatically repopulated. The repopulation time delay depends on the routing protocols in use.
6. Select **Apply**.

Configure Port Mirroring

Before You Begin

- To change a port mirroring configuration, first disable mirroring.

About This Task

Use port mirroring to aid in diagnostic and security operations.

Use port mirroring to make a copy of a traffic flow and send that copy to a device for analysis, for example, for diagnostic sniffing. Use the mirror to see the packets in the flow without breaking into the physical connection to place a packet onto the sniffer inline. You can also use port mirroring for security. You can send flows to inspection engines for post processing.

Connect the sniffer (or other traffic analyzer) to the output port you specify in this procedure.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **General**.
3. Select the **Port Mirrors** tab.
4. Select **Insert**.
5. To enable port mirroring for the instance, select the **Enable** check box.
6. Configure mirroring as required.



Note

- When you configure `tx` mode port mirroring on T-UNI and SPBM NNI ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports and the system displays on the mirror destination port, although they do not egress the mirror source port. This is because `tx` mode port mirroring happens on the mirror source port before the source port squelching logic drops the packets at the egress port.
- The available four mirroring resources are shared between Fabric RSPAN and regular port mirroring, and are allocated based on the mode configured, Ingress (`rx`) or Egress (`tx`). Each configured mode occupies one mirroring resource, but when you configure the mode as `both`, it occupies two mirroring resources (one for Rx and one for Tx).
- You cannot configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.

7. Select **Insert**.

Port Mirrors Field Descriptions

Use the data in the following table to use the **Port Mirrors** tab.

Name	Description
Id	Specifies an assigned identifier for the configured port mirroring instance.
MirroredPortList	Specifies a port to be mirrored (the source port).
Enable	Enables or disables this port mirroring instance. The default value is Enable.
Mode	Specifies the traffic direction of the packet being mirrored: <ul style="list-style-type: none"> • <code>tx</code> mirrors egress packets. • <code>rx</code> mirrors ingress packets. • <code>both</code> mirrors both egress and ingress packets. The default is <code>rx</code> .

Name	Description
MirroringPortList	Specifies a destination port (the port to which the mirrored packets are forwarded). Configures the mirroring port.
MirroringMltId	Specifies the destination MultiLink trunking ID.
MonitoringIsidOffset	Used to configure the monitoring I-SID offset value. The offset ID is mapped to the actual monitor I-SID value to which the packets are mirrored.
MonitoringIsid	Specifies the actual monitor I-SID value to which the packets are mirrored.
MirroringQos	Used to define the Quality of Service (QoS) profiles for the mirrored packet into monitoring I-SID.

Configure ACLs for Mirroring

Before You Begin

- The ACL exists.

About This Task

Configure the access control list (ACL) to mirror packets for an access control entry (ACE) that matches a particular packet.

You cannot change a parameter that the system displays it dimmed; in this case, delete the ACL, and then configure a new one.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Double-click the parameter **MirrorMltId** to configure mirroring to a destination MLT group.
5. Double-click the parameter **MirrorDstPortList** to configure mirroring to a destination port or ports.
6. Select **Apply**.

ACL Field Descriptions

Use the data in the following table to use the **ACL** tab.

Name	Description
AcId	Specifies a unique identifier for the ACL.
Type	<p>Specifies the ACL type. Valid options are</p> <ul style="list-style-type: none"> • inVlan • inPort • outPort • inVsn <p>Important: The inVlan ACLs drop packets if you add a VLAN after ACE creation.</p> <p>Important: You can insert an inVsn ACL type for a Switched UNI only if the Switched UNI I-SID is associated with a platform VLAN.</p>
Name	Specifies a descriptive user-defined name for the ACL.
VlanList	For inVlan ACL types, specifies all VLANs to associate with the ACL.
PortList	For inPort and outPort ACL types, specifies the ports to associate with the ACL.
DefaultAction	Specifies the action taken when no ACEs in the ACL match. Valid options are deny and permit, with permit as the default. Deny means the system drops the packets; permit means the system forwards packets.
ControlPktAction	Specifies the action taken for control packets. Valid options are deny and permit.
State	Enables or disables all of the ACEs in the ACL. The default value is enable.
PktType	Indicates the packet type to which this ACL applies.
MirrorMltId	Configures mirroring to a destination MLT.
MirrorDstPortList	Configures mirroring to a destination port or ports.
MatchType	<p>For inVsn ACL types, specifies the match type to associate with the ACL. Valid options are:</p> <ul style="list-style-type: none"> • <i>both</i> for traffic ingressing on both UNI ports and network-to-network interface (NNI) ports terminating on this node • <i>terminatingNNIOnly</i> for traffic ingressing on NNI ports only and terminating on this node • <i>uniOnly</i> for traffic ingressing on UNI ports only <p>The default value is <i>both</i></p>

Name	Description
Isid	<p>For inVsn ACL types, specifies the I-SID associated with the customer VLAN (Layer 2 VSN) or the customer VRF (Layer 3 VSN). This I-SID should already be configured on the fabric node.</p> <p>The InVSN Filter supports IP Shortcut traffic if the inVsn ACL match type is <i>both</i>. In this case, the I-SID is zero (0).</p> <p>Important: You can specify a Switched UNI I-SID if the I-SID is associated with a platform VLAN.</p>
Origin	<p>Indicates the origin of the ACL:</p> <ul style="list-style-type: none">• config - ACL created by the user.• eap - ACL created by Extensible Authentication Protocol (EAP) through Remote Authentication Dial-In User Service (RADIUS) response.
DefaultSvcRate	<p>Specifies the service rate limit in kbps {8-4000000000}.The granularity is 8 kbps.</p>
DefaultPeakRate	<p>Specifies the value when exceeded causes packets to drop on ingress. Peak rate limit in kbps {8-4000000000}.The granularity is 8 kbps.</p>

Configure ACEs for Mirroring

Before You Begin

- The ACL exists.
- The ACE exists.

About This Task

Configure actions to use filters for flow mirroring. Use an ACE to define the mirroring actions the filter performs.

If you use the mirror action, ensure that you specify the mirroring destination: IP address, MLTs, ports, or VLANs.

Procedure

1. In the navigation pane, expand **Configuration > Security > Data Path**.
2. Select **Advanced Filters (ACE/ACLs)**.
3. Select the **ACL** tab.
4. Select the ACL for which to modify an ACE.
5. Select **ACE**.
6. Select an ACE, and then select **Action**.
7. Configure one of: **DstPortList**, **DstMltId**, or **DstIp**.
8. Select **Apply**.

Action Field Descriptions

Use the data in the following table to use the **Action** tab.



Note

The table lists the options for both Security ACEs and QoS ACEs. Dependent upon the ACE, the system displays different options on the EDM interface.

Name	Description
AcId	Specifies the ACL ID.
AcId	Specifies the ACE ID.
Mode	Configures the action mode for security ACEs. The default value is deny.
RemarkDscp	Specifies the new Per-Hop Behavior (PHB) for matching packets: phbcs0, phbcs1, phbaf11, phbaf12, phbaf13, phbcs2, phbaf21, phbaf22, phbaf23, phbcs3, phbaf31, phbaf32, phbaf33, phbcs4, phbaf41, phbaf42, phbaf43, phbcs5, phbef, phbcs6, phbcs7. This action is a QoS action.
RemarkDot1Priority	Specifies the new 802.1 priority bit for matching packets: zero, one, two, three, four, five, six, or seven. This action is a QoS action. Note: This does not apply to IPv6 filtering.
InternalQoS	This variable is a QoS action. The default value is 1. Note: This does not apply to IPv6 filtering.
RedirectNextHop	Specifies the next-hop IPv4 address (a.b.c.d) or IPv6 address (aaaa:bbbb:cccc:dddd:eeee:ffff:gggg:hhhh) for redirect mode. Applies to ingress ACLs (routed and Layer 2 packets).
Count	Enables the ability to count matching packets. Use this parameter with either a security or QoS ACE. The default is disabled.
DstPortList	Configures mirroring to a destination port or ports. This action is a security action.
DstMltId	Configures mirroring to a destination MLT. This action is a security action.

Name	Description
MonitoringIsidOffset	Configures the monitoring I-SID offset value. The offset ID is mapped to the actual monitor I-SID value to which the packets are mirrored.
MirroringQoS	Defines the Quality of Service (QoS) profiles for the mirrored packet into monitoring I-SID.

Run a Ping Test

About This Task

Use ping to determine if an entity is reachable.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. From the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Ping/Trace Route**.
3. Select the **Ping Control** tab.
4. Select **Insert**.
5. In **OwnerIndex**, type the owner index.
6. In **TestName**, type the name of the test.
7. In **TargetAddress**, type the host IP address.
8. From **AdminStatus**, select **enabled**.
9. In the remainder of the options, type the desired values.
10. Select **Insert**.
11. Select an entry, and then select **Start**.
Let the test run for several seconds.
12. Select an entry, and then select **Stop**.
13. View the Ping results.

Ping Control Field Descriptions

Use the data in the following table to use the **Ping Control** tab.

Name	Description
OwnerIndex	Provides access control by a security administrator using the View-Based Access Control Model (VACM) for tables in which multiple users need to independently create or modify entries. This is a string of up to 32 characters.
TestName	Specifies the name of the ping test.

Name	Description
TargetAddress	Specifies the host address to use at a remote host to perform a ping operation.
DataSize	Specifies the size of the data portion (in octets) to transmit in a ping operation. The default is 16.
TimeOut	Specifies the timeout value, in seconds, for a remote ping operation. The default is 3 seconds.
ProbeCount	Specifies the number of times to perform a ping operation at a remote host. The default is 1.
AdminStatus	Specifies the state of the ping control entry: enabled or disabled. The default is disabled.
DataFill	Determines the data portion of a probe packet.
Frequency	Specifies the number of seconds to wait before repeating a ping test. The default is 0.
MaxRows	Specifies the maximum number of entries allowed in the PingProbeHistory table. The default is 50.
TrapGeneration	<p>Specifies when to generate a notification. The options are:</p> <ul style="list-style-type: none"> ProbeFailure—Generates a PingProbeFailed notification subject to the value of TrapProbeFailureFilter. The object TrapProbeFailureFilter can specify the number of successive probe failures that are required before a pingProbeFailed notification is generated. TestFailure—Generates a PingTestFailed notification. The object TrapTestFailureFilter can determine the number of probe failures that signal when a test fails. TestCompletion—Generates a PingTestCompleted notification. <p>The value of this object defaults to zero, indicating that none of the options have been selected.</p>
TrapProbeFailureFilter	Specifies the number of successive probe failures that are required before a pingProbeFailed notification is generated. The default is 1.
TrapTestFailureFilter	Determines the number of probe failures that signal when a test fails. The default is 1.
Descr	Describes the remote ping test. The default is 0x00.
SourceAddress	Specifies the IP address (a.b.c.d) as the source address in outgoing probe packets.
ByPassRouteTable	Enables (optionally) the bypassing of the route table. The default is disabled.

View Ping Results

About This Task

View ping results to view performance-related data.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. From the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Ping/Trace Route**.
3. Select the **Ping Control** tab.
4. Select a ping test entry.
5. Select **Ping Result**.

Ping Result Field Descriptions

Use the data in the following table to use the **Ping Result** tab.

Name	Description
OwnerIndex	Specifies the ping test owner.
TestName	Specifies the test name.
OperStatus	Indicates the operational status of the test. The default is disabled.
IpTargetAddressType	Specifies the IP address type of the target. The default is unknown.
IpTargetAddress	Specifies the IP address of the target.
MinRtt	Specifies the minimum ping round-trip-time (RTT) received. A value of 0 means that no RTT is received.
MaxRtt	Specifies the maximum ping RTT received. A value of 0 means that no RTT is received.
AverageRtt	Specifies the current average ping RTT.
ProbeResponses	Specifies the number of responses to probes.
SentProbes	Specifies the number of sent probes.
RttSumOfSquares	Specifies the sum of squares of RTT for all probes received.
LastGoodProbe	Specifies the date and time when the last response is received for a probe.

View Ping Probe History

About This Task

View the ping probe history to view the history of ping tests performed by the switch.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. From the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Ping/Trace Route**.
3. Select a ping entry.
4. Select **Ping Probe History**.

Ping Probe History Field Descriptions

Use the data in the following table to use the **Ping Probe History** tab.

Name	Description
OwnerIndex	Specifies the owner index
TestName	Indicates the name given to the test.
Index	Specifies the index number.
Response	Indicates the amount of time, measured in milliseconds, between request (probe) and response, or when the request timed out. Response is reported as 0 when it is not possible to transmit a probe.
Status	Indicates the status of the response; the result of a particular probe done by a remote host.
LastRC	Indicates the last implementation-method-specific reply code (RC) received. If ICMP Echo is used, then a successful probe ends when an ICMP response is received that contains the code ICMP_ECHOREPLY(0).
Time	Indicates the timestamp for this probe result.

Run a Traceroute Test

About This Task

Run a traceroute test to determine the route packets take through a network to a destination.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. From the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Ping/Trace Route**.
3. Select the **Trace Route Control** tab.
4. Select **Insert**.
5. Configure the instance as required.
6. Select **Insert**.
7. Select an entry, and then select **Start**.
Let the test run for several seconds.
8. Select an entry, and then select **Stop**.
9. View the traceroute test results.

Trace Route Control Field Descriptions

Use the data in the following table to use the **Trace Route Control** tab.

Name	Description
OwnerIndex	Provides access control by a security administrator using the VACM for tables in which multiple users need to independently create or modify entries.
TestName	Specifies the name of the traceroute test.
TargetAddressType	Specifies the type of host address to use on the traceroute request at the remote host. The default is IPv4.
TargetAddress	Specifies the host address used on the traceroute request at the remote host.
ByPassRouteTable	Enables bypassing of the route table. If you enable this variable, the remote host bypasses the normal routing tables and sends directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. You can use this variable to perform the traceroute operation to a local host through an interface that has no route defined. The default is disabled.
DataSize	Specifies the size of the data portion of a traceroute request in octets. The default is 1.
TimeOut	Specifies the timeout value, in seconds, for a traceroute request. The default is 3.
ProbesPerHop	Specifies the number of times to reissue a traceroute request with the same time-to-live (TTL) value. The default is 3.

Name	Description
Port	<p>Specifies the UDP port to which to send the traceroute request. Specify a port that is not in use at the destination (target) host. The default is the IANA assigned port 33434.</p> <p>Note: If the traceroute action is directed to an IPv6 host address, Linux increments the UDP port on a per-TTL basis. For an IPv4 host address, Linux increments the UDP port on a per-probe basis. Because the traceroute command sends a default of three probes, three incrementing ports will be sent for an IPv4 host address. If you use the -p parameter with a value greater than 65533, the traceroute command fails for an IPv4 host address because the maximum port number, 65535, is exceeded.</p>
MaxTtl	Specifies the maximum time-to-live from 1–255. The default is 30.
DSField	Specifies the value to store in the Differentiated Services (DS) field in the IP packet used to encapsulate the traceroute probe. The default is 0.
SourceAddressType	Specifies the type of the source address to use at a remote host.
SourceAddress	Uses the specified IP address (which must be an IP number, not a hostname) as the source address in outgoing probe packets.
IfIndex	Directs the traceroute probes to be transmitted over the specified interface. The default is 0.
MiscOptions	Enables an application to specify implementation-dependent options.
MaxFailures	Indicates the maximum number of consecutive timeouts allowed before terminating a remote traceroute request. The default is 5.
DontFragment	Enables setting of the do not fragment (DF) flag in the IP header for a probe. The default is disabled.
InitialTtl	Specifies the initial time-to-live (TTL) value to use. The default is 1.
Frequency	Specifies the number of seconds to wait before repeating a traceroute test as defined by the value of the various objects in the corresponding row. The default is 0.
StorageType	Specifies the storage type for this row.
AdminStatus	Specifies the desired state for TraceRouteCtlEntry. The options are enabled or disabled. The default is disabled.

Name	Description
MaxRows	Specifies the maximum number of entries allowed in the TraceRouteProbeHistoryTable. The default is 50.
TrapGeneration	Determines when to generate a notification for this entry. The options are <ul style="list-style-type: none">• pathChange —Generates a TraceRoutePathChange notification after the current path varies from a previously determined path.• testFailure —Generates a TraceRouteTestFailed notification after the full path to a target cannot be determined.• testCompletion —Generates a TraceRouteTestCompleted notification after the path to a target has been determined.
Descr	Describes the remote traceroute test.
CreateHopsEntries	Stores the current path for a traceroute test in the TraceRouteHopsTable on an individual hop basis when the value of this object is true. The default is false.
Type	Reports or selects the implementation method to use for performing a traceroute operation.

View Traceroute Results

About This Task

View traceroute results to view performance-related data.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. From the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Ping/Trace Route**.
3. Select the **Trace Route Control** tab.
4. Select a traceroute entry.
5. Select **Trace Route Result**.

Trace Route Result Field Descriptions

Use the data in the following table to use the **Trace Route Result** tab.

Name	Description
OwnerIndex	Specifies the index of the owner.
TestName	Specifies the name of the test.

Name	Description
OperStatus	Specifies the operational status of the test. The default is disabled.
CurHopCount	Specifies the current count of hops.
CurProbeCount	Specifies the current count of probes.
IpTgtAddressType	Specifies the IP target address type
IpTgtAddr	Specifies the IP target address.
TestAttempts	Specifies the number of test attempts.
TestSuccesses	Specifies the number of successful test attempts.
LastGoodPath	Specifies the date and time when the last response is received for a probe.

View the Traceroute History

About This Task

View the traceroute history to view the history of traceroute tests performed by the switch.

The traceroute probe history contains probe information for the hops in the routing path.



Note

Troubleshooting using ping and traceroute is not supported on EDM. For more information, see *Fabric Engine Release Notes*. As an alternative, use CLI.

Procedure

1. From the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **Ping/Trace Route**.
3. Select the **Trace Route Control** tab.
4. Select an entry.
5. Select **Trace Route Probe History**.

Route Probe History Field Descriptions

Use the data in the following table to use the **Trace Route Probe History** tab.

Name	Description
OwnerIndex	Identifies the Trace Route entry to which a probe result belongs.
TestName	Specifies the test name.
Index	Specifies the Index.
HopIndex	Indicates for which hop in a traceroute path the probe results are intended.

Name	Description
ProbeIndex	Specifies the index of a probe for a particular hop in a traceroute path.
HAddrType	Specifies the IP address type of the hop to which this probe belongs.
HAddr	Specifies the IP address of the hop to which this probe belongs.
Response	Specifies the cumulative results at any time.
Status	Specifies the status of the probe.
LastRC	When a new entry is added, the old entry is purged if the total number of entries exceeds the specified maximum number of entries in the Control Table Entry.
Time	Specifies the response time of the probe.

Configure I-SID Monitoring

Use the following procedure to configure Fabric RSPAN on the Backbone Edge Bridge (BEB) connected to the monitor station.

About This Task

You cannot configure the source of mirrored traffic (mirroring to an I-SID) and the analyzer (monitoring an I-SID) on the same local device with the same I-SID offset. If you require mirroring and monitoring on the same local device, use standard port-based mirroring instead of Fabric RSPAN. Fabric RSPAN mirrors traffic into an I-SID of the SPB Fabric network and monitors traffic on the remote device; the network analyzer resides on the remote monitoring device and not on the same local device.



Note

If you change the egress port or egress MLT for a particular session using a CLI command, it overwrites the existing egress port list or egress MLT.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **General**.
3. Select the **Monitor-By-ISID** tab.
4. Select **Insert**.
5. Configure the parameters as required.
6. Select **Insert**.
7. To modify mappings, double-click a parameter to view a list of options.
8. Select **Apply**.

Monitor-By-ISID Field Descriptions

Use the data in the following table to use the **Monitor-By-ISID** tab.

Name	Description
Index	Specifies the entry that contains monitor by I-SID information.
MonitorIsidOffset	Configures the monitoring I-SID offset value. The offset ID is mapped to the actual monitor I-SID value to which the packets are mirrored.
MonitorIsid	Specifies the actual monitor I-SID value to which packets are mirrored.
EgressPortList	Specifies the egress ports to which traffic analyzers connect.
EgressMltId	Specifies the egress MLT ID to which traffic analyzers connect.
MapToVlanId	Specifies the VLAN ID to map with mirrored traffic on the monitoring node.
Enable	Enables or disables monitoring by I-SID.

View Fabric RSPAN (Mirror to I-SID) Statistics

About This Task

Use the following procedure to view or clear statistics of the number of packets mirrored into I-SID on the mirroring BEB.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics**.
2. Select **General**.
3. Select the **Isid-Mirroring Stats** tab.

Isid-Mirroring Stats Field Descriptions

Use the data in the following table to use the **Isid-Mirroring Stats** tab.

Name	Description
Index	Specifies the entry that contains Fabric RSPAN statistics information.
MonitorIsid	Specifies the actual monitor I-SID value to which the packets are mirrored.
MirroredPackets	Specifies the number of packets mirrored into I-SID on the mirroring BEB.
ClearStats	Clears the Fabric RSPAN statistics.

Hardware Troubleshooting

The following sections provide troubleshooting information for common hardware problems.

Time Domain Reflectometer (TDR)

Table 332: Time Domain Reflectometer product support

Feature	Product	Release introduced
Time Domain Reflectometer (Cable Diagnostic)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 9.0.2
	5420 Series	Fabric Engine 8.10 Not supported on 5420F-24S-4XE
	5520 Series	Fabric Engine 9.0.2 Not supported on 5520-24X or 5520-48SE
	5720 Series	Fabric Engine 9.0.2
	7520 Series	Fabric Engine 9.0.2 Supported on 7520-48XT-6C
	7720 Series	Not Supported
	7830 Series	Not Applicable
	VSP 4900 Series	Not Supported
	VSP 7400 Series	Not Supported

With Time Domain Reflectometer (TDR) cable diagnostic, you can test Ethernet cables connected to a switch port, or ports, for defects. Use this functionality to diagnose a link with poor connectivity. The cable diagnostic detects the state of the port or ports, the status of the cable pairs, and the length of the cable.



Note

You can perform TDR cable diagnostic on copper ports.

The following table lists the states of the cables connected to a port or ports.

State	Description
Ok	No cable fault is detected.
Open	Detects lack of continuity between the pins at each end of the cable.
Short	Detects a short circuit in the cable.
Open/Short	Detects either a short circuit on the cable or a lack of continuity between the pins at each end of the cable.
Crosstalk	Detects interference and degradation between two pairs in the cable.
Unknown	Detects an unknown issue.

Troubleshoot Fiber Optic Links

About This Task

You can troubleshoot fiber optic links to ensure that the optical transmitters and receivers operate correctly, and to determine if a receiver is saturated, or does not receive enough power.

To troubleshoot optical links and devices, you can use Digital Diagnostic Monitoring (DDM), as well as published optical specifications.

For more information about transceivers, see [Extreme Optics](#) website.



Important

As a best practice, use transceivers documented in [Extreme Optics](#) website, as they have been through extensive qualification and testing. Extreme Networks is not responsible for issues related to third party transceivers.

Procedure

1. Measure the transmit power.
2. Compare the measured transmit power with the specified launch power.
The values are similar. If the measured power is far below the specified value, a faulty transmitter is a possible cause.
3. Compare the measured transmit power for the near-end optical device to the measured transmit power for the far-end device.
Large differences can mean that the optical devices are mismatched (that is, -SX versus -LX).
4. Measure the receive power at each end of the link.
5. Compare the receive power to the transmit power.
 - For short fiber links, the transmit and received power are similar (after taking into consideration connection losses).
 - For long fiber links, the transmit and received power are similar (after taking into consideration connection losses and fiber attenuation).

Large differences can mean a damaged fiber or dirty or faulty connectors. Large differences can also mean that the link does not use the right type of fiber (single mode or multimode). If the receiver power is measured to be zero, and the link worked previously, it is probable that the far-end transmitter is not operating or the fiber is broken.

6. Compare the measured receive power for the near-end optical device to the measured receive power for the far-end device.

Large differences could mean that the optical devices are mismatched (that is, -SX versus -LX). If optical devices are mismatched, the receiver can be saturated (overdriven).

7. If a receiver is saturated but still operable, install a suitable attenuator.

For long-haul optical devices, the receive power must be significantly less than the transmit power.

8. To help debug the link, loop back the local transmit and receive ports, and use the DDM parameters to help determine the fault.

Hardware Troubleshooting using CLI

This section provides procedures to troubleshoot common hardware problems using the command line interface (CLI).

Use Trace to Diagnose Hardware Problems

Use trace to observe the status of a hardware module at a given time.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Begin the trace operation:
line-card 1 trace level [*<Module_ID>*]{*<0-4>*}
3. Search the trace for a specific string value:
line-card 1 trace grep {*WORD<0-1024>*}

Example

```
Switch:1>enable
```

Begin the trace operation:

```
Switch:1#line-card 1 trace level 67 1
```

Search the trace for a specific string value:

```
Switch:1#line-card 1 trace grep 00-1A-4B-8A-FB-6B
```

Variable Definitions

The following table defines parameters for the **line-card 1** command.

Variable	Value
<i><Module_ID></i> { <i><0-4></i> }	<p>Starts the trace by specifying the module ID and level.</p> <p><i><Module_ID></i> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see <i>Fabric Engine and VOSS Feature Matrix</i>. To see the list of module IDs, use the show trace modid-list command or CLI command completion Help.</p> <p><i><0-4></i> specifies the trace level from 0–4, where 0 is disabled; 1 is very terse; 2 is terse; 3 is verbose, 4 is very verbose.</p>
<i>WORD<0-1024></i>	Performs a string search in the trace.

Troubleshoot USB Viewing Problems

After you insert a USB device in the USB slot, the Linux system automatically detects and mounts the device. If you cannot view files on the device, perform this procedure.



Note

Not all hardware platforms can use the USB device for additional file storage. Some platforms use the USB as part of the system operation and must never be removed. For more information, see your hardware documentation.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Check the file system:

```
ls /usb/
```

3. Remove a USB device:

- a. Unmount the USB device:

```
usb-stop
```

- b. Wait for the response that indicates it is safe to remove the device.

- c. Physically remove the device.

4. Remove and then reinsert the device.

5. Check the device for errors:

```
dos-chkdsk /usb
```

Run the **dos-chkdsk /usb repair** command, if at the end of the **dos-chkdsk /usb** command output you see:

1) Correct

2) Don't correct

6. If errors are detected, then you can reformat the device:

```
dos-format /usb
```



Caution

If you format the device, you erase all data on the device.

Example

Check the file system:

```
Switch:1>enable
Switch:1#ls /usb/
Listing Directory /usb/:
drwxr-xr-x 4 0 0 4096 Jan 1 1970 ./
drwxrwxr-x
22 0 0 0 Sep 9 20:22 ../
drwxr-xr-x 2 0 0 4096 Mar 17 16:03 Photos-of-Flash-
drwxr-xr-x 2 0 0 4096 Jun 13 20:56 intflash/
```

Check the device for errors:

```
Switch:1#usb-stop
It is now safe to remove the USB device.
Switch:1#dos-chkdsk /usb
/usr/sbin/fsck.vfat /dev/usb1 -v >& /dev/console dosfsck 2.11a
(05 Mar 2010)
dosfsck 2.11a, 05 Mar 2010, FAT32, LFN
Checking we can access the last sector of the filesystem
Boot sector contents:
System ID "mkdosfs"
Media byte 0xf8 (hard disk)
512 bytes per logical sector
4096 bytes per cluster
32 reserved sectors
First FAT starts at byte 16384 (sector 32)
2 FATs, 32 bit entries
3897344 bytes per FAT (= 7612 sectors)
Root directory start at cluster 2 (arbitrary size)
Data area starts at byte 7811072 (sector 15256)
974240 data clusters (3990487040 bytes)
62 sectors/track, 124 heads
0 hidden sectors
7809178 sectors total
Checking for unused clusters.
Checking free cluster summary.
/dev/usb1: 17 files, 174804/974240 clusters
```

If errors are detected, reformat the disk:

```
Switch:1#dos-format /usb
```

Reset a Transceiver

Reset a transceiver to simulate removal and reinsertion of the transceiver, which can be helpful in troubleshooting. For example, if authentication of the transceiver fails but you believe the transceiver is a qualified Extreme Networks part, you can reset the transceiver to begin the authentication process again.

Before You Begin

- Before you use the **pluggable-optical-module reset** command, ensure the port is administratively down to avoid link flaps.

About This Task

Resetting the transceiver stops traffic and triggers log messages similar to the removal and insertion of the transceiver.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```


2. Reset the transceiver:

```
pluggable-optical-module reset {slot/port[/sub-port]}
```

**Important**

Not all hardware platforms support these port types. For more information, see your hardware documentation.

Examples

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#pluggable-optical-module reset 1/41
Switch:1(config)#
CP1 [06/25/14 22:15:09.644] 0x0000c5e7 00300001.232 DYNAMIC SET GlobalRouter HW INFO Link
Down(1/41)
CP1 [06/25/14 22:15:10.267] 0x000e0597 00000000 GlobalRouter HAL INFO GBIC removed from
slot 1 Port 41 Type:40GbSR4 Vendor:Extreme Networks
CP1 [06/25/14 22:15:13.015] 0x000e0598 00000000 GlobalRouter HAL INFO GBIC inserted in
slot 1 Port 41 Type:40GbSR4 Vendor:Extreme Networks
CP1 [06/25/14 22:15:14.562] 0x0000c5ec 00300001.232 DYNAMIC CLEAR GlobalRouter HW INFO
Link Up(1/41)

Switch:1(config)#pluggable-optical-module reset 1/1
Switch:1(config)#CP1 [03/31/16 10:48:24.492:UTC] 0x0000c5e7 00300001.384 DYNAMIC SET
GlobalRouter HW INFO Link Down(1/1)
CP1 [03/31/16 10:48:24.601:UTC] 0x000e0597 00000000 GlobalRouter HAL INFO GBIC removed
from slot 1 Port 1 Type:100GbCR4 Vendor:Extreme Networks
CP1 [03/31/16 10:48:24.710:UTC] 0x0000c5e7 00300001.385 DYNAMIC SET GlobalRouter HW INFO
Link Down(1/2)
CP1 [03/31/16 10:48:26.668:UTC] 0x000e0598 00000000 GlobalRouter HAL INFO GBIC inserted
in slot 1 Port 1 Type:100GbCR4 Vendor:Extreme Networks
CP1 [03/31/16 10:48:26.988:UTC] 0x0000c5ec 00300001.385 DYNAMIC CLEAR GlobalRouter HW
INFO Link Up(1/2)
CP1 [03/31/16 10:48:27.099:UTC] 0x0000c5ec 00300001.384 DYNAMIC CLEAR GlobalRouter HW
INFO Link Up(1/1)
```

Variable Definitions

The following table defines parameters for the **pluggable-optical-module reset** command.

Variable	Value
{slot/port[/sub-port]}	Specifies location of the transceiver to reset. Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Initiate a Time Domain Reflectometer Diagnostic Test

About This Task**Note**

This procedure does not apply to 7720 Series. For model support, see *Fabric Engine and VOSS Feature Matrix*.

Perform this procedure to initiate a Time Domain Reflectometer (TDR) diagnostic test to detect defects on Ethernet cables or to determine the status and length of the cable pairs attached to copper ports on the switch. Use this functionality to diagnose a link with poor connectivity.

TDR cable diagnostic operates on administratively enabled ports only. You cannot initiate a TDR test on a port with Auto-Negotiation disabled.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Initiate the diagnostic test on a port:

```
tdr-test [port {slot/port[/sub-port]}[-slot/port[/sub-port]][,...]  
[slot/all][all]]
```



Note

The following warning message displays to inform you that the TDR test impacts the traffic flow, if you continue. %WARNING: TDR test will impact the traffic flow. Do you want to continue (y/n) ?

3. Display the results of the TDR diagnostic test:

```
show tdr-test [port ]{slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]]
```

Example

Initiate the cable diagnostic test on port 1/1:

```
Switch:1#tdr-test port 1/1  
%WARNING: TDR test will impact the traffic flow.  
Do you want to continue (y/n)? y  
Test Started for port 1/1
```

Display the results of the cable diagnostic test on port 1/1:

```
Switch:1>show tdr-test port 1/1  
Port          1/1 (cable length accuracy +/- 10 meters)  
Port Status   Ok  
Pair1 Status   Ok  
Pair1 Length   4M  
Pair2 Status   Ok  
Pair2 Length   9M  
Pair3 Status   Ok  
Pair3 Length   9M  
Pair4 Status   Crosstalk  
Pair4 Length   4M
```

Initiate the cable diagnostic test on ports 1/1-1/2:

```
Switch:1#tdr-test port 1/1-1/2  
%WARNING: TDR test will impact the traffic flow.  
Do you want to continue (y/n)? y  
Test Started for port 1/1  
Test Started for port 1/2
```

Display the results of the cable diagnostic test on ports 1/1-1/3

```
Switch:1#show tdr-test port 1/1-1/3
Port                               1/1 (cable length accuracy +/- 10 meters)
Port Status                        Ok
Pair1 Status                       Ok
Pair1 Length                       4M
Pair2 Status                       Ok
Pair2 Length                       9M
Pair3 Status                       Ok
Pair3 Length                       9M
Pair4 Status                       Crosstalk
Pair4 Length                       4M

Port                               1/2 (cable length accuracy +/- 10 meters)
Port Status                        Ok
Pair1 Status                       Ok
Pair1 Length                       5M
Pair2 Status                       Ok
Pair2 Length                       2M
Pair3 Status                       Ok
Pair3 Length                       124M
Pair4 Status                       Ok
Pair4 Length                       4M

Port                               1/3 Cable diagnosis has not been performed on the specified port 1/3
```

Variable Definitions

The following table defines parameters for the **tdr-test** command.

Variable	Value
<code>port {slot/port[-slot/port][,...] [slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all)

Hardware Troubleshooting using EDM

This section provides procedures to troubleshoot common hardware problems using Enterprise Device Manager (EDM).

Initiate a Time Domain Reflectometer Diagnostic Test

About This Task



Note

This procedure does not apply to 7720 Series. For model support, see *Fabric Engine and VOSS Feature Matrix*.

Perform this procedure to initiate a Time Domain Reflectometer (TDR) test to detect defects on Ethernet cables or to determine the status and length of the cable pairs attached to copper ports on the switch. Use this functionality to diagnose a link with poor connectivity.

TDR cable diagnostic operates on administratively enabled ports only. You cannot initiate a TDR test on a port with Auto-Negotiation disabled.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **TDR** tab.
5. Select **StartTest**.
6. Select **Apply**.
7. Select **Yes** to continue.
8. Select **Refresh** to display the results.

TDR Field Descriptions

Use data from the following table in the **TDR** tab.

Name	Description
StartTest	Enables the cable diagnostic test on the current port.
TestDone	Specifies whether the test is complete.
NormalCableLength	Specifies the cable length in meters.
CableLength	Specifies the length of the cable in meters based on the average electrical length of 4 pairs.
CableStatus	<p>Specifies the status of the cable as a summation of the status of the cable conductor pairs.</p> <ul style="list-style-type: none">• 1—Fail: the cable experiences any combination of open and shorted pairs.• 2—Normal: the cable operates without any faults

Name	Description
Pair#Status where # is the number of the pair. For example, Pair1Status .	Specifies the status of the pair. The following lists the states of the cables connected to a port or ports: <ul style="list-style-type: none"> • pairFail • pairNormal • pairOpen • pairShorted • pairNotApplicable • pairNotTested • pairForce • pinShort • pairUnknown • pairOpenShorted • pairCrossTalk
Pair#ErrLength where # is the number of the pair. For example, Pair1ErrLength .	Specifies the approximate distance to the break, measured in meters on the first pair.

Upper Layer Troubleshooting

Use the information in this section to troubleshoot upper layer problems.

Troubleshooting SNMP

About This Task

Troubleshoot Simple Network Management Protocol (SNMP) if the network management station (NMS) does not receive traps.

Verify the management configurations for the management station. Also verify the management station setup. If the management station can reach a device but not receive traps, verify the trap configurations (that is, the trap destination address and the traps to be sent).

If you enable enhanced secure mode, the switch does not support the default SNMPv1 and default SNMPv2 community strings, and default SNMPv3 user name. The individual in the administrator access level role can configure a non-default value for the community strings, and the switch can continue to support SNMPv1 and SNMPv2. The individual in the administrator access level role can also configure a non-default value for the SNMPv3 user name and the switch can continue to support SNMPv3. If you disable enhanced secure mode, the SNMPv1 and SNMPv2 support for community strings remains the same, and the default SNMPv3 user name remains the same. Enhanced secure mode is disabled by default.

Procedure

1. From the NMS, ping the IP address for the switch. If you can ping successfully, the IP address is valid and you may have a problem with the SNMP setup.
If you cannot ping the switch, you have a problem with either the path or the IP address.
2. Telnet to the switch.
If you can Telnet, the switch IP address is correct.
3. If Telnet does not work, connect to the console port using a serial line connection and ensure that the IP address configuration is correct.
4. If the management station is on a separate subnet, make sure that the gateway address and subnet mask are correct.
5. Using a management application, perform an SNMP Get request and an SNMP Set request (that is, try to poll the device or change a configuration using management software).
6. If you cannot reach the device using SNMP, access the console port, and then ensure that the SNMP community strings and traps are correct.
7. Use sniffer traces to verify that the switch receives the poll.
8. Use sniffer traces to verify that the NMS receives the response.
9. Verify that the data in the response is the data that was requested.

SNMP Trap not Received

Perform the following procedure to troubleshoot issues in which an SNMP trap is not received.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Show the BPDU Guard status for the port:
show spanning-tree bpduguard {slot/port[-slot/port][,...][slot/all]
[all]}
3. Configure the correct SNMP target information:
snmp-server host WORD<1-256> [port <1-65535>] v3 {noAuthNoPriv|
authNoPriv|authPriv WORD<1-32> [inform [timeout <1-2147483647>]
[retries <0-255>]] [filter WORD<1-32>]}

Example

In the following example, BPDU guard is enabled on port 1/8, BPDU packets are received, port 1/8 is disabled, and the TimerCount is incrementing, but no SNMP trap is ever received.

```
Switch:1>enable
Switch:1#show spanning-tree bpduguard 1/8
```

```
=====
Bpdu Guard
```

Port NUM	PORT MLTID	PORT ADMIN_STATE	PORT OPER_STATE	TIMEOUT	TIMER COUNT	BPDUGUARD ADMIN_STATE	BPDUGUARD ORIGIN
1/8		Down	Down	120	0	Disabled	CONFIG

Variable Definitions

The following table defines parameters for the **show spanning-tree** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

The following table defines parameters for the **snmp-server host** command.

Variable	Value
<code>filter WORD<1-32></code>	Specifies a filter profile name.
<code>host WORD<1-256></code>	Specifies the IPv4 or IPv6 host address
<code>inform [timeout <1-2147483647>]</code>	Specifies the notify type. The optional timeout parameter configures the timeout value, which specifies the time to wait for a reply before resending the inform message. Time is specified in centiseconds
<code>noAuthNoPriv authNoPriv authPriv WORD<1-32></code>	Specifies the security level.
<code>port <1-65535></code>	Specifies the port number that will be set as the destination port at the UDP level in the trap packet.
<code>retries <0-255></code>	Specifies the number of packets to be sent if no reply is received.
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Troubleshooting DHCP

About This Task

Perform this procedure to troubleshoot the following Dynamic Host Configuration Protocol (DHCP) scenarios:

- The client cannot obtain a DHCP address when in the same subnet.
- The client cannot obtain a DHCP address when in a different subnet.

When the DHCP server and client are on the different subnets or VLANs, you must configure the device as a DHCP relay agent. The device must forward DHCP requests to the DHCP server. You must perform extra troubleshooting steps to troubleshoot the DHCP relay agent.

Procedure

1. Check the physical connectivity between the DHCP client and server.
2. Verify network connectivity by configuring a static IP address on a client workstation.
If the workstation still cannot reach the network, the problem is not DHCP. Start troubleshooting network connectivity.
3. Attempt to obtain an IP address from the DHCP server by manually forcing the client to send a DHCP request.
If the client obtains an IP address after the PC startup is complete, the issue is not the DHCP server.
4. Obtain an IP address on the same subnet or VLAN as the DHCP server.
If the issue persists, the problem may be with the DHCP server. If DHCP is working on the same subnet or VLAN as the DHCP server, the DHCP issue can be with the DHCP relay agent.
5. Confirm the DHCP relay agent configuration is correct.
6. Obtain sniffer traces where the traffic ingresses and egresses the switch and also on the client side of the network.
7. Check the logs on the switch for errors such as size exceeded or incorrect packet format.

Troubleshooting DHCP Relay

Before You Begin

- Configure the server to reply to the client subnet. Check the server configuration file to verify the configuration.
- Configure a route on the server for the client subnet to create a path on which to send replies.

About This Task

Perform this procedure to troubleshoot the DHCP relay agent.

Procedure

1. Verify that the interfaces that link the client and server are up, and that the ports are in the forwarding state.

- a. To verify client availability, you can configure a temporary static IP address on the client, and then use the ping command.

```
ping WORD<0-256>
```

- b. To verify the port is in the forwarding state, use the following command for the slot and port number:

```
show spanning-tree [rstp|mstp] port role [{slot/port[-slot/port]}
[,...][slot/all][all]]
```

If STP detects loops in the configuration, it blocks ports to avoid flooding in the network. In this situation, the port is not in the forwarding state.

2. Ensure that DHCP is enabled on the client interface and that a valid forwarding path exists and is enabled. Ensure the server is reachable.
3. View the statistics counters for the relay.
4. If request or reply counters do not increase, use a sniffer tool to ensure that the client sends the packets, and that the interface module receives the packets.

You can configure mirroring for the ingress port to verify if the packets reach the module.

- a. If the client sends the packets, check that the packets reach the CPP and search the trace results for the ingress port:

```
trace level 9 3
```

```
trace grep WORD<0-128>
```

- b. If the packets reach the CPP, check that they reach the DHCP protocol; check for errors or packet drop messages:

```
trace level 170 3
```

```
trace grep WORD<0-128>
```

5. If Option 82 is enabled, check the statistic counters for dropped packets, and perform a trace for the DHCP protocol:

```
trace level 170 3
```

Example

```
Switch:1#ping 192.0.2.31
```

```
Switch:1:#show spanning-tree mstp port role
```

CIST Port Roles and States				
Port-Index	Port-Role	Port-State	PortSTPStatus	PortOperStatus
1/1	Disabled	Forwarding	Disabled	Disabled
1/2	Disabled	Forwarding	Disabled	Disabled
1/3	Disabled	Discarding	Enabled	Disabled
1/4	Disabled	Discarding	Enabled	Disabled
1/5	Disabled	Forwarding	Disabled	Disabled
1/6	Disabled	Forwarding	Disabled	Disabled
1/7	Disabled	Forwarding	Disabled	Disabled
1/8	Disabled	Forwarding	Disabled	Disabled

```

1/9      Disabled   Discarding   Enabled     Disabled
1/10     Disabled   Discarding   Enabled     Disabled
1/11     Disabled   Discarding   Enabled     Disabled
1/12     Designated Forwarding   Enabled     Enabled
1/13     Disabled   Forwarding   Disabled    Disabled
1/14     Disabled   Forwarding   Disabled    Disabled

--More-- (q = quit)
Switch:1:#trace level 9 3
Switch:1:#trace grep 00-1A-4B-8A-FB-6B

```

Variable Definitions

Use the data in the following table to use the troubleshooting commands in this procedure.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>WORD<0-128></code>	Specifies the text string to use as the search criterion.
<code>WORD<0-256></code>	Specifies the IP address.

Troubleshoot Client Connection to the DHCP Server

About This Task

Perform this procedure if the client cannot reach the DHCP server.

Procedure

1. Check that the DHCP relay agent in the network switch is correctly configured.
2. Check that the DHCP server configuration is correct.
3. Check for routing issues.

The routing in the network may not be configured so that DHCP request and reply packets are propagated. You can use ping and traceroute.

4. Check that the DHCP pools are correctly configured.
5. If the client cannot reach the server because the link is down, enable auto-negotiation on the link.

Troubleshooting IPv6 DHCP Relay

The following sections provide troubleshooting information for IPv6 DHCP Relay.

IPv6 DHCP Relay Switch Side Troubleshooting

With DHCP Relay, the switch only participates in forwarding the requests and replies to and from the client and the DHCP server. The switch always acts as the relay agent, on which you configure the forward path to the server.

To troubleshoot DHCP Relay issues on the switch, use the following procedure.

Procedure

1. Verify that the DHCP server is reachable using ping. If ping is working and the DHCP server is reachable, DHCP should work.
2. Verify that the relay agents and the forward path configured are reachable. Ping the server and the gateway to the server.
3. Check that the relay agent configurations are correct. Also verify that DHCP is enabled on the switch:

```
show ipv6 dhcp-relay interface {gigabitEthernet {slot/port[/sub-port]}
[-slot/port[/sub-port]][,...][slot/all][all]}|vlan <1-4059>
```

4. Verify that IPv6 forwarding is enabled globally:
5. Verify that the IPv6 based VLAN where the DHCP relay agent is configured is enabled:
6. In a scenario with VRRP and SMLT, configure VRRP IP as the DHCP relay agent.
7. When using the VRRP VRID as the relay agent, make sure the VRRP configurations are proper.
8. To verify that relay forward and relay receive are working, enable trace for DHCP with IPv6, and grep trace for relay:

```
trace level 66 3
```

```
trace grep relay
```

```
trace screen enable
```

9. Display the count of DHCP Relay requests and replies to verify the system received requests and replies:

```
show ipv6 dhcp-relay counters
```

IPv6 DHCP Relay Server Side Troubleshooting

Use the following procedure to troubleshoot IPv6 DHCP Relay on the server side.

Procedure

1. Enable the services on the server side, and then create an IP pool.
The IP pool must contain the range of addresses that you want to assign to the clients.

Configure the IP pool with the same network subnet as that of the relay agent.

2. When the configuration is complete, initiate a DHCP request from a client.
3. Check the log file available on the server to verify the reason for packet drop.

4. Capture the packets on the server side using Ethereal.
5. From the server side, use ping to verify that the relay agent address is reachable.
Ensure that a route to the relay is configured.
6. For more configuration aspects, see the Microsoft webpage for troubleshooting and configuration issues.

**Note**

You can receive some log messages that indicate the system cannot forward packets. However, certain situations are not DHCP failures.

Example 1: if you receive the message `0x00108796 (relayMsgSend)`: cannot find route entry for destination on the console, you must ping the server. If the server is not reachable, the system cannot forward the packet. This is not a DHCP issue.

Example 2: if you receive the message `0x00108705` this indicates a problem at the transmission level. Check the server reachability and ensure that MAC learning is correct before you pursue DHCP issues.

IPv6 DHCP Relay Client Side Troubleshooting

You can collect a client console dump, which can be used to analyze why the received packet cannot be processed and the allocated address cannot be used by the client.

In addition, restarting the client can also fix the issue in some cases.

Make sure the client supports IPv6 requests.

Connect the server directly to the client. If the IP is assigned, then the problem is with the relay.

Enable Trace Messages for IPv6 DHCP Relay

Use this procedure to enable trace for IPv6 DHCP Relay and enable IPv6 forwarding trace.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. To troubleshoot IPv6 DHCP Relay, you can enable rcip6 trace messages using the following command:
`trace level 66 3`
3. You can also enable IPv6 forwarding trace using the following command:
`trace ipv6 forwarding enable <all|debug|error|info|pkt|warn>`

Example

Enable rcip6 trace messages and enable IPv6 forwarding trace:

```
Switch:1>enable
Switch:1#trace level 66 3
Switch:1#trace ipv6 forwarding
```

Troubleshooting TACACS+

The switch supports the Terminal Access Controller Access Control System plus (TACACS+) client. TACACS+ is a remote authentication protocol that provides centralized validation of users who attempt to gain access to a router or network access server (NAS). The TACACS+ feature is disabled by default.

The switch implementation of TACACS+ does not support:

- Earlier versions of TACACS
- Point-to-Point Protocol (PPP) authentication and accounting
- IPv6 addresses

See the following sections to troubleshoot TACACS+.

Unable to Log On Using Telnet

If you cannot log on using Telnet, perform the following steps.

Procedure

1. Check whether the TACACS+ server is available or unreachable.
2. On the TACACS+ server, check whether you configured the privilege level correctly. On successful authorization, the TACACS+ server returns an access level to the switch for the current user, which determines the user access privileges. The switch supports access levels 1 to 6 and access level 15.

The following table maps user accounts to TACACS+ privilege level.

Switch access level	TACACS+ privilege level	Description
NONE	0	If the TACACS+ server returns an access level of 0, the user is denied access. You cannot log into the device if you have an access level of 0.
READ ONLY	1	Permits you to view only configuration and status information.
LAYER 1 READ WRITE	2	Permits you to view most of the switch configuration and status information and change physical port settings.
LAYER 2 READ WRITE	3	Permits you to view and change configuration and status information for Layer 2 (bridging and switching) functions.
LAYER 3 READ WRITE	4	Permits you to view and change configuration and status information for Layer 2 and Layer 3 (routing) functions.

Switch access level	TACACS+ privilege level	Description
READ WRITE	5	Permits you to view and change configuration and status information across the switch. This level does not allow you to change security and password settings.
READ WRITE ALL	6	Permits you to have all the rights of read-write access and the ability to change security settings, including command line interface (CLI) and web-based management user names and passwords, and the SNMP community strings.
NONE	7 to 14	If the TACACS+ server returns an access level of 7 to 14, the user is denied access. You cannot log into the device if you have an access level of 7 to 14.
READ WRITE ALL	15	<p>Permits you to have all the rights of read-write access and the ability to change security settings, including command line interface (CLI) and web-based management user names and passwords, and the SNMP community strings.</p> <p>Note: Access level 15 is internally mapped to access level 6, which ensures consistency with other vendor implementations. The switch does not differentiate between an access level of 6 and an access level of 15.</p>

After you enable TACACS+ authorization, the current privilege-level to command mapping on the switch is no longer relevant because the TACACS+ server has

complete responsibility for command authorization. TACACS+ authorization provides access to the system based on user name, not based on privilege level.



Note

If you want to switch to a privilege level 'X' using **tacacs switch level <1-15>** command, you must create a user "\$enabX\$" on the TACACS+ server. X is the privilege level that you want to change.

3. On the TACACS+ server, check whether you configured the password and user name correctly.
4. On the TACACS+ server, check whether you configured the switch IP address in the trust list.
5. Check whether you configured the encryption key, connection mode (single connection or per-session connection), and TCP port number the same on the TACACS+ server and switch.
6. If you can log on to the switch, check whether the TACACS+ server configured on the platform has the correct IP address:

```
show tacacs
```

7. Use the output from the **show tacacs** command to verify whether you configured the single connection option on the platform, and whether the TACACS+ server supports the single connection.

Example

Check whether the TACACS+ server configured on the platform has the correct IP address:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

    global enable : false

    authentication enabled for : cli

    accounting enabled for : none

    authorization : disabled

    User privilege levels set for command authorization : None

Server:

    create :

Prio   Status  Key      Port  IP address  Timeout  Single Source SourceEnabled
Primary NotConn ***** 3    192.0.2.254   30    true 5.5.5.5   true
Backup  NotConn ***** 47   198.51.100.1  10    false 0.0.0.0   false
```

Unable to Log On Using SSH

If you cannot log on using Secure Shell (SSH), perform the following steps.

Procedure

1. Verify that the network, the switch, and the TACACS+ server is reachable.
2. Verify whether you configured the SSH client correctly.
3. Verify whether you enabled and configured the SSH function correctly on the switch:

```
show ssh global
```

Example

Verify whether you enabled and configured SSH function correctly on the switch:

```
Switch:1(config)#show ssh global
Total Active Sessions      : 1
version                   : v2only
port                      : 22
max-sessions              : 4
timeout                   : 60
action rsa-host key       : rsa-hostkeysize 2048
action dsa-host key       : dsa-hostkeysize 1024
rsa-auth                  : true
dsa-auth                  : true
pass-auth                 : true
keyboard-interactive-auth : false
x509-auth                 : true
x509-auth Trustpoint CA Name :
x509-auth Identity Subject Name : not configured
x509-auth overwrite       : false
x509-auth strip-domain    : false
x509-auth use-domain      : -
x509-auth revocation-check-method : OCSP
sftp enable               : true

client enable             : true

enable                   : true
authentication-type       : aead-aes-128-gcm-ssh aead-aes-256-gcm-ssh hmac-sha1 hmac-sha2-256
encryption-type           : 3des-cbc aead-aes-128-gcm-ssh aead-aes-256-gcm-ssh aes128-cbc aes128-ctr
                           aes192-cbc aes192-ctr aes256-cbc aes256-ctr blowfish-cbc rijndael128-cbc
                           rijndael192-cbc
key-exchange-method       : diffie-hellman-group-exchange-sha256 diffie-hellman-group14-sha1
                           diffie-hellman-group14-sha256 diffie-hellman-group16-sha512
                           diffie-hellman-group18-sha512
```

Unable to Log On by any Means

If you cannot log on by any means, perform the following steps.

Procedure

1. Check whether the TACACS+ server runs properly and try to restart the TACACS+ server.
2. Check whether you enabled both TACACS+ and RADIUS on the switch.

```
show radius
```

```
show tacacs
```

If TACACS+ fails, RADIUS can take over the authentication, authorization, and accounting (AAA) process.

3. Check whether you configured the TACACS+ server to unencrypted mode, as the switch always sends encrypted TACACS+ messages.

4. Check whether you configured the switch properly.



Note

In particular, check the IP address and key.

```
show tacacs
```

5. Check whether you configured the encryption key, connection mode (single connection or per-session connection), and TCP port number the same on the TACACS+ server and switch.
6. If the server connects directly, check whether the administrative and operation status of the port is up:

```
show interfaces gigabitethernet {slot/port[/sub-port]}[-slot/port[/sub-port]][,...]
```

7. If the server is connected in a network, check whether the switch has a route configured to the server network:

```
show ip route
```

8. If the server is connected in a network, check whether the switch has a route configured to the server network:

```
show ip route
```

```
show ipv6 route
```

9. For the Out-of-Band (OOB) or VLAN Segmented Management Instance, check whether the switch has a route configured to the server network:

```
show mgmt ip route
```

```
show mgmt ipv6 route
```

```
show mgmt ip route static
```

10. For Segmented Management Instance troubleshooting, check the management network statistics:

```
show mgmt ip arp
```

```
show khi mgmt statistics
```

```
show mgmt ip ip-statistics
```

```
show mgmt ip icmp-statistics
```

Examples

Check if you enabled both TACACS+ and RADIUS on the switch:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

  global enable : false

  authentication enabled for : cli

  accounting enabled for : none
```

```

authorization : disabled

User privilege levels set for command authorization : None

Server:

                create :

Prio   Status  Key      Port  IP address  Timeout Single Source SourceEnabled
Primary NotConn ***** 3    192.0.2.254    30   true 5.5.5.5   true
Backup  NotConn ***** 47   198.51.100.1   10  false 0.0.0.0   false

Switch:1>show radius
      acct-attribute-value : 193
      acct-enable : false
      acct-include-cli-commands : false
      access-priority-attribute : 192
      auth-info-attr-value : 91
      command-access-attribute : 194
      cli-commands-attribute : 195
      cli-cmd-count : 40
      cli-profile-enable : false
      enable : false
      igap-passwd-attr : standard
      igap-timeout-log-fsize : 512
      maxserver : 10
      mcast-addr-attr-value : 90
      supported-vendor-ids : 1584, 562, 1916
      secure-flag : false

```

Check if the administrative and operation status of the port is up:

```

Switch:1#show interfaces gigabitethernet 1/2

=====
                        Port Interface
=====
PORT      LINK  PORT      PHYSICAL      STATUS
NUM  INDEX DESCRIPTION  TRAP  LOCK    MTU   ADDRESS      ADMIN  OPERATE
-----
1/2    257   1000BaseTX  true  false   1950  00:24:7f:a1:70:61 up    up

=====
                        Port Name
=====
PORT      OPERATE  OPERATE  OPERATE
NUM  NAME      DESCRIPTION  STATUS  DUPLEX  SPEED  VL
AN
-----
1/2    gged      1000BaseTX  up      full    1000   Ta

=====
                        Port Config
=====
PORT      DIFF-SERV  QOS  MLT  VENDOR
--More-- (q = quit)

```

Check if the switch has a route configured to the server network:

```
Switch:1(config)#show ip route
```

```
=====
                        IP Route - GlobalRouter
=====
=====
```

DST	MASK	NEXT	NH VRF/ISID	COST	FACE	INTER PROT	AGE	TYPE	PRF	ACCESS
198.51.100.1	255.255.255.255	192.0.2.65	GlobalRouter	1	100	OSPF	0	IB	125	Trusted
198.51.100.5	255.255.255.255	192.0.2.5	-	1	0	LOC	0	DB	0	Trusted
198.51.100.13	255.255.255.255		GlobalRouter	10	1000	ISIS	0	IBS	7	Trusted
198.51.100.200	255.255.255.255		GlobalRouter	10	1000	ISIS	0	IBS	7	Trusted

```

4 out of 4 Total Num of Route Entries, 4 Total Num of Dest Networks displayed.
=====
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
=====
```

Check if the Segmented Management Instance has a route configured to the server network:

```
Switch:1(config)#show mgmt ip route
```

```
=====
                        Mgmt IPv4 Route Information - Table main
=====
```

DEST/MASK	NEXTHOP	METRIC	INTERFACE	TYPE
198.51.100.0/16	198.51.100.1	300	Mgmt-oob1	STATIC
198.51.100.0/23	0.0.0.0	1	Mgmt-oob1	LOCAL
192.0.2.0/8	192.0.2.1	300	Mgmt-oob1	STATIC

```

3 out of 3 Total Num of mgmt ip route displayed
=====
```

Administrator Unable to Obtain Accounting Information from the TACACS+ Server

If the administrator is unable to obtain accounting information from the TACACS+ server, perform the following steps.

Procedure

1. Check whether you enabled accounting on the switch:

```
show tacacs
```
2. Check whether you enabled accounting on the TACACS+ server.

Example

Check whether accounting is enabled on the switch:

```
Switch:1>enable
Switch:1(config)#show tacacs

Global Status:

global enable : false
```

```

authentication enabled for : cli

accounting enabled for : none

authorization : disabled

User privilege levels set for command authorization : None

Server:

                create :

Prio   Status  Key      Port  IP address  Timeout Single Source SourceEnabled
Primary NotConn ***** 3    192.0.2.254   30    true 5.5.5.5   true
Backup  NotConn ***** 47   198.51.100.1  10    false 0.0.0.0   false

```

Trap Server Cannot Receive Trap Packets from the Switch

If the trap server cannot receive trap packets from the switch, perform the following steps.

Procedure

1. Check whether you configured the trap server correctly on the switch:

```
show snmp-server host
```
2. Check whether a firewall exists between the switch and the trap server.

Example

Check whether you configured the trap server correctly on the switch:

```

Switch:1>enable
Switch:1#show snmp-server host

=====
Notify Configuration
=====
Notify Name          Tag          Type
-----
Inform              informTag    inform
Trap                trapTag      trap

=====
Notify Profile Configuration
=====
Params Name          Profile Name
-----
AuthNoPriv-md5       profile2
AuthPriv-md5         profile3
NoAuthNoPriv-md5     profile1

=====
Target Address Configuration
=====
Target Name          TDomain      TAddress      TMask
-----
4c20cc369925edbd1fe3cf8e2584c498 ipv4         47.17.142.155:162

```

```
55fca382ffba169e986783bbbdedc334 ipv4 47.17.143.57:162
```

Target Address Configuration			
Target Name	Timeout	Retry	TagList
Params	MMS		
4c20cc369925edbd1fe3cf8e2584c498	1500	3	trapTag
4c20cc369925edbd1fe3cf8e2584c498	484		
55fca382ffba169e986783bbbdedc334	1500	3	trapTag
55fca382ffba169e986783bbbdedc334	484		

Target Params Configuration			
Target Name	MP Model	Security Name	Sec
Level			
4c20cc369925edbd1fe3cf8e2584c498	snmpv1	readview	noAu
thNoPriv			
55fca382ffba169e986783bbbdedc334	snmpv2c	secret	noAu
thNoPriv			
TparamV1	snmpv1	readview	noAu
thNoPriv			
TparamV2	snmpv2c	readview	noAu
thNoPriv			

Troubleshoot TACACS+ Problems

Use the **trace level** command to check traps and log files to see any TACACS+ failure. If TACACS+ experiences failure conditions, the TACACS+ module sends SNMP traps to notify the user. The TACACS+ module also logs the failure information into the system log file.

About This Task



Caution

Using the trace tool inappropriately can cause CPU lockup conditions, loss of access to the device, loss of protocols, and service degradation. If you use trace level 3 (verbose) or trace level 4 (very verbose), do not use the screen to view commands due to the volume of information the system generates and the effect on the system.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Configure the trace level for the TACACS+ module:
trace level 109 <1-4>

The TACACS+ module ID is 109.

3. Stop trace:

```
trace shutdown
```

4. View the trace results on screen:

```
trace screen enable
```

5. View trace saved to a file:

```
show trace file [tail]
```

6. Save the trace file to the Compact Flash card for retrieval:

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt. By default, the system saves the file to the external flash.

7. Save the trace file for retrieval:

```
save trace [file WORD<1-99>]
```

If you do not specify a file name, the file name is systrace.txt.

Variable Definitions

The following table defines parameters for the **trace** command.

Variable	Value
<i>level</i> [<Module_ID>] [<1-4>]	Starts the trace by specifying the module ID and level. Module ID 23 represents the IGMP module <Module_ID> specifies the module for the trace. Different hardware platforms support different IDs because of feature support differences. For feature support information, see <i>Fabric Engine and VOSS Feature Matrix</i> . To see the list of module IDs, use the show trace modid-list command or CLI command completion Help. <0-4> specifies the trace level: <ul style="list-style-type: none"> • 0 — Disabled • 1 — Very terse • 2 — Terse • 3 — Verbose • 4 — Very verbose
<i>shutdown</i>	Stops the trace operation.
<i>screen</i> {disable enable}	Enables or disables the display of trace output to the screen. Important: As a best practice, avoid using the screen to view commands if you use trace level 3 (verbose) or trace level 4 (very verbose) due to the volume of information generated and the effect on the system.

The following table defines parameters for the **show trace** command.

Variable	Value
<i>file [tail]</i>	Displays the trace results saved to a file.
<i>level</i>	Displays the current trace level for all modules.
<i>modid-list</i>	Specifies the module ID list.

Use BGP Debugging Commands

Use global and peer debug commands to display specific debug messages for the global and peer Border Gateway Protocol (BGP) configuration, including the BGP neighbors.

You can use these commands to troubleshoot the BGP configuration.

Procedure

1. Enter BGP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router bgp
```
2. Show specific debug messages for the global BGP configuration:

```
global-debug mask WORD<1-100>
```
3. Display specific debug messages for the global BGP neighbors:

```
neighbor-debug-all mask WORD<1-100>
```
4. Display specific debug messages for BGP peers or peer groups:

```
neighbor <nbr_ipaddr|peer-group-name> neighbor-debug-mask WORD<1-100>
```
5. Display debug messages on the console:

```
debug-screen <on|off>
```

Example

Display the global debug messages for error and packet:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router bgp
Switch:1(router-bgp)#global-debug mask error,packet
```

End (disable) the display of global debug messages:

```
Switch:1(router-bgp)#global-debug mask none
```

Display specific debug messages for the global BGP neighbors:

```
Switch:1(router-bgp)#neighbor-debug-all mask packet,event
```

Display specific debug messages for BGP peers or peer groups:

```
Switch:1(router-bgp)#neighbor 192.0.2.10 neighbor-debug-mask event,trace
```

Display debug messages on the console:

```
Switch:1(router-bgp)#debug-screen on
```

Variable Definitions

The following table defines parameters for the **global-debug mask** and **neighbor-debug-all mask** commands.

Variable	Value
<i>WORD<1-100></i>	Specifies one or more mask choices that you enter, separated by commas with no space between choices. For example: [<mask>,<mask>,<mask>...]. Options include: none, all, error, packet, event, trace, warning, state, init, filter, update.

The following table defines parameters for the **neighbor** command.

Variable	Value
<i><nbr_ipaddr peer-group-name></i>	Specifies the IP address or the group name of the peer.
<i>WORD<1-100></i>	Specifies one or more mask choices that you enter, separated by commas with no space between choices. For example: [<mask>,<mask>,<mask>...]. Options include: none, all, error, packet, event, trace, warning, state, init, filter, update.

Job Aid

Use debug command values to control debug messages for global BGP message types, and for message types associated with a specified BGP peer or peer group. The following table identifies mask categories and messages.

Table 333: Mask categories and messages

Mask category	Message
none	None disables the display of all debug messages.
all	All configures the device to show all categories of debug messages.
error	Error configures the device to show error debug messages.
packet	Packet configures the device to show packet debug messages.
event	Event configures the device to show event debug messages.
warning	Warning configures the device to show warning debug messages.

Table 333: Mask categories and messages (continued)

Mask category	Message
init	Init configures the device to show initialization debug messages.
filter	Filter configures the device to show filter-related debug messages.
update	Update configures the device to show update-related debug messages.



Virtual Link Aggregation Control Protocol

[Virtual Link Aggregation Control Protocol](#) on page 3965

[VLACP Configuration using CLI](#) on page 3967

[VLACP Configuration using EDM](#) on page 3976

Table 334: Virtual Link Aggregation Control Protocol (VLACP) product support

Feature	Product	Release introduced
Virtual Link Aggregation Control Protocol (VLACP)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
VLACP Flap Detect and Damping	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.1

Virtual Link Aggregation Control Protocol

Use Virtual Link Aggregation Control Protocol (VLACP) as an extension to LACP for end-to-end failure detection. VLACP is not a link aggregation protocol, it is a mechanism to periodically check the end-to-end status of a point-to-point connection. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure end-to-end communication. After Hello packets are not received, VLACP transitions to a failure state, which indicates a service provider failure and that the port is disabled.

The VLACP only works for port-to-port communications where there is a guarantee for a logical port-to-port match through the service provider. VLACP does not work for port-to-multiport communications where there is no guarantee for a point-to-point match through the service provider. You can configure VLACP on a port.

You can also use VLACP with MLT to complement its capabilities and provide quick failure detection.

VLACP trap messages are sent to the management stations if the VLACP state changes. If the failure is local, the only traps that are generated are port linkdown or port linkup.

The Ethernet cannot detect end-to-end failures. Functions such as remote fault indication or far-end fault indication extend the Ethernet to detect remote link failures. A major limitation of these functions is that they terminate at the next Ethernet hop. They cannot determine failures on an end-to-end basis.

For example, in [Figure 242](#), after the Enterprise networks connect the aggregated Ethernet trunk groups through a service provider network connection (for example, through a VPN), far-end failures cannot be signaled with Ethernet-based functions that operate end-to-end through the service provider network. The multilink trunk (between Enterprise switches S1 and S2) extends through the Service Provider (SP) network.

The following illustration shows an MLT running with VLACP. VLACP can operate end-to-end, but you can also use it as a point-to-point link.




Figure 242: Problem description (1 of 2)

In the following illustration, if the Layer 2 link on S1 (S1/L2) fails, the link-down failure is not propagated over the SP network to S2 and S2 continues to send traffic over the failed S2/L2 link.




Figure 243: Problem description (2 of 2)

Use VLACP to detect far-end failures, which causes MLT to failover if end-to-end connectivity is not guaranteed for links in an aggregation group. VLACP prevents the failure scenario.

The switch software uses the following VLACP timers:

- fast periodic timer—100 to 20 000 ms; default 200 ms
- slow periodic timer—10 000 to 30 000 ms; default 30 000 ms

VLACP Considerations

Use the information in this section to understand the considerations while configuring VLACP into your network.

- If a VLACP-enabled port does not receive a VLACP Data Unit (VLACPDU), it must enter the disabled state. There are occasions where a VLACP-enabled port does not receive a VLACPDU but remains in the forwarding state. To avoid this situation, ensure that the VLACP configuration at the port level is consistent. You must either enable or disable both sides of the point-to-point connection.
- If VLACP is enabled on a MACsec Key Agreement-enabled link, it takes approximately 30 seconds for the VLACP session to begin.

You can configure VLACP on each port. The port can be either an individual port or an MLT member. VLACPDUs can be sent periodically on each port where VLACP is enabled to exchange VLACPDUs from an end-to-end perspective. If VLACPDUs are not received on a particular link, that link is taken down after the expiry timeout occurs (timeout scale x periodic time).

VLACP Flap Detect and Damping

When there is instability or packet loss in a connection, the Virtual Link Aggregation Control Protocol (VLACP) state of the port flaps, causing services such as IP multicast to stop and start rapidly. This behavior causes system-wide instability, including high CPU utilization. In such cases, VLACP Flap Detect and Damping provides link flap detection capability, by configuring a specific time interval and frequency to count the number of VLACP flaps. After the system detects excessive VLACP flaps, it disables the specific VLACP port until the root cause is resolved.

For example, VLACP Flap Detect and Damping detects 3 VLACP flaps within 60 seconds, by default. After the system detects the first VLACP flap, the flap timer starts to count the number of VLACP flaps occurring within 60 seconds. If the flap count reaches 3 before the timer ends, the system disables the specific VLACP port, and generates a Simple Network Management Protocol (SNMP) trap. For more information about SNMP traps, see [Logs and Traps Fundamentals](#) on page 2536.

If you view the port state information, the port displays as operationally down and the reason provided is VLACP_FLAP. The following example shows a port that is disabled because of VLACP flapping; only the relevant sections of the show command output is included.

```
Switch:1#show interfaces gigabitethernet 1/1
```

Port Interface									
PORT NUM	INDEX	DESCRIPTION	LINK TRAP	PORT LOCK	MTU	PHYSICAL ADDRESS	STATUS		
							ADMIN	OPERATE	
1/1	192	1000BaseTX	true	false	1950	e4:5d:52:3c:64:00	up		down

Port State					
PORT NUM	ADMINSTATUS	PORTSTATE	REASON	DATE	
1/1	up	down	VLACP_FLAP	09/19/19 02:22:58	



Note

VLACP Flap Detect and Damping does not support auto-recovery of VLACP ports; you must reenab the ports manually.

VLACP Flap Detect and Damping Considerations

Use the following information when configuring VLACP Flap Detect and Damping.

- Do not enable VLACP Flap Detect and Damping on Link Aggregation Control Protocol (LACP) enabled ports.



Note

Link flap detection takes priority if you enable both VLACP Flap Detect and Damping and link flap detection on the same port.

VLACP Configuration using CLI

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. VLACP detects end-to-end failures in the switch. Virtual LACP cannot interoperate with Link Aggregation Control Protocol (LACP).

Configure VLACP on a Port

Configure VLACP on a port to ensure there is end-to-end reachability. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure there is

an end-to-end approach. After Hello packets are not received, VLACP transitions to a failure state and disables the port.



Important

Changes made at the global level override and reset all port level settings.

About This Task

Use the following information to prevent flooding VLACP packets across a defaulted switch:

- Use the default MAC address, 01:80:c2:00:11:00, for end-to-end connections that traverse an intermediate network.
- Use the reserved multicast MAC address 01-80-c2-00-00-0f for directly-connected, peer-to-peer links.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure optional parameters for the port. If you do not configure these parameters, the system uses the default values.

- a. Configure the protocol identification for the port:

```
vlacp ethertype <1536-65535 | 0x600-0xffff> [funcmac-addr  
0x00:0x00:0x00:0x00:0x00:0x00]
```

- b. Configure the fast or slow periodic times:

```
vlacp fast-periodic-time <100-20000> | slow-periodic-time  
<10000-30000>
```

You can configure both parameters in the same command entry.

- c. Configure the timeout parameters:

```
vlacp timeout <long|short> timeout-scale <2-10>
```

You can configure both parameters in the same command entry.

3. Enable VLACP on a port:

```
vlacp enable
```

Example

Configure VLACP on port 1/1:

```
Switch:1#configure terminal
Switch:1#interface GigabitEthernet 1/2
Switch:1#vlacp fast-periodic-time 400 timeout short
Switch:1#vlacp enable
```

Variable Definitions

Use the data in the following table to help you use the **vlacp** command.

Variable	Value
<i>enable</i>	Enables VLACP for this port. The default is disabled.
<i>ethertype</i> <1536-65535 0x600-0xffff>	Configures the VLACP protocol identification for this port. Enter the type in decimal or hexadecimal format. The default is 0x8103.
<i>fast-periodic-time</i> <100-20000>	Configures the fast periodic time (in milliseconds) for this port. The default is 200.
<i>funcmac-addr</i> <0x00:0x00:0x00:0x00:0x00:0x00>	Configures the multicast MAC address used for the VLACPDU. Specify a MAC address in the format 0x00:0x00:0x00:0x00:0x00:0x00. The default is 01:80:c2:00:11:00.
<i>slow-periodic-time</i> <10000-30000>	Configures the slow periodic time (in milliseconds) for a specific port type. The default is 30,000.
<i>timeout</i> {long short}	Configures the port to use the long or short timeout: <ul style="list-style-type: none"> <i>long</i> sets the port to use the timeout-scale value multiplied by the slow periodic time. <i>short</i> sets the port to use the timeout-scale value multiplied by the fast periodic time. <p>For example, if you specify a short timeout, set the timeout-scale value to 3, and the fast periodic time to 400 ms, the timer expires within 1000 to 1200 ms. The default is long.</p>
<i>timeout-scale</i> <2-10>	Configures a timeout scale for this port used to calculate the timeout. The default value is 3.

View the VLACP Port Configuration

View the VLACP port configuration to show the port VLACP configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View the VLACP port configuration for all interfaces:

```
show vlacp interface gigabitethernet [vid <1-4059>] [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]
```

Example

```
Switch:1>show vlacp interface gigabitethernet
=====
                        VLACP Information
=====
INDEX  ADMIN   OPER   PORT   FAST  SLOW  TIMEOUT  TIMEOUT  ETHER   MAC
      ENABLED ENABLED STATE   TIME   TIME   TIME     SCALE   TYPE
ADDR
-----
1/1    false   false   DOWN    200  30000  long     3       0x8103  01:80:c2:00:11:00
1/2    false   false   DOWN    200  30000  long     3       0x8103  01:80:c2:00:11:00
1/3    false   false   DOWN    200  30000  long     3       0x8103  01:80:c2:00:11:00
1/4    false   false   DOWN    200  30000  long     3       0x8103  01:80:c2:00:11:00
1/5    false   false   DOWN    200  30000  long     3       0x8103  01:80:c2:00:11:00
--More-- (q = quit)
=====
                        VLACP Flap Detect Information
=====
INDEX  FLAP    FLAP    FLAP    TOTAL    FIRST-FLAP    LAST-FLAP
      DETECT  FREQ    INTERVAL FLAP      TIME          TIME
-----
1/1    false   3       60      0         -- --         -- --
1/2    false   3       60      0         -- --         -- --
1/3    false   3       60      0         -- --         -- --
1/4    false   3       60      0         -- --         -- --
1/5    false   3       60      0         -- --         --
--
--More-- (q = quit)

Switch:1>show vlacp interface gigabitethernet 2/1 vid 2
=====
                        VLACP Information
=====
INDEX  ADMIN   OPER   PORT   FAST  SLOW  TIMEOUT  TIMEOUT  ETHER   MAC
      ENABLED ENABLED STATE   TIME   TIME   TIME     SCALE   TYPE   ADDR
-----
2/1    true    true    up     500    30000  short    3       0x8103  01:80:c2:00:00:0f
=====
                        VLACP Flap Detect Information
=====
INDEX  FLAP    FLAP    FLAP    TOTAL    FIRST-FLAP    LAST-FLAP
      DETECT  FREQ    INTERVAL FLAP      TIME          TIME
-----
2/1    true    10      10      4         09/20/11 08:24:10  09/20/11 08:26:10

Switch:1>show vlacp interface gigabitethernet 1/7
=====
                        VLACP Information
=====
INDEX  ADMIN   OPER   PORT   FAST  SLOW  TIMEOUT  TIMEOUT  ETHER   MAC
      ENABLED ENABLED STATE   TIME   TIME   TIME     SCALE   TYPE   ADDR
-----
1/7    false   false   DOWN    200    30000  long     3       0x8103  01:80:c2:00:11:00
=====
```


VLACP Flap Detect Information						
INDEX	FLAP DETECT	FLAP FREQ	FLAP INTERVAL	TOTAL FLAP	FIRST-FLAP TIME	LAST-FLAP TIME
1/7	false	3	60	0	-- --	-- --

Variable definitions

Use the data in the following table to use the **show vlacp interface gigabitethernet** command.

Variable	Value
<code>vid <1-4059></code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>The VLAN ID is in one of the following formats: A single VLAN ID (vlan-id), a range of VLAN IDs [(vlan-id)-(vlan-id)] or a series of VLAN IDs (vlan-id, vlan-id, vlan-id).</p>
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Specifies a port or list of ports to show only the VLACP information for those ports. Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Display VLACP Statistics for Specific Ports

Display VLACP statistics for specific ports to manage network performance.

About This Task

You can enable sequence numbers for each VLACPDU to assist in monitoring performance. The switch counts mismatched PDU sequence numbers to determine packet loss information. By default, sequence numbers are enabled.

You can use the show commands from Privileged EXEC mode but must enter Global Configuration mode to enable or disable the sequence numbers.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Confirm sequence numbers are enabled:

```
show vlacp
```
3. (Optional) Enable sequence numbers for VLACPDUs:

```
vlacp sequence-num
```
4. View VLACP statistics:

```
show interfaces gigabitEthernet statistics vlacp [{slot/port[/sub-port]
[-slot/port[/sub-port]][,...]} ]
```
5. (Optional) View VLACP statistics history:

```
show interfaces gigabitEthernet statistics vlacp history [{slot/port[-slot/port]
[,...][slot/all][all]]}
```
6. (Optional) Clear VLACP statistics:

```
clear vlacp stats [port {slot/port[-slot/port][,...][slot/all][all]]}
```
7. (Optional) Disable sequence numbers for VLACPDUs:

```
no vlacp sequence-num
```

Example

Determine if sequence numbers are enabled, and then view port statistics. Port numbering may differ depending on your product and configuration.

```
Switch:1(config)#show vlacp
```

```
=====
```

```
                        Vlacp Global Information
```

```
=====
```

```
      SystemId: 32:11:9f:20:00:00
```

```
        Vlacp: enable
```

```
  Vlacp Sequence Number: enable
```

```
Switch:1(config)#show interfaces gigabitEthernet statistics vlacp
```

```
=====
```

```
                        Port Stats Vlacp
```

```
=====
```

PORT NUM	TX VLACPDU	RX VLACPDU	SEQNUM MISMATCH
1/1	106058	105554	0
1/11	15	12	0
1/23	0	0	0

```
=====
```

Variable Definitions

Use the data in the following table to use the commands in this procedure.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configure VLACP Globally

Use VLACP as an extension to LACP for end-to-end failure detection. Enable or disable VLACP globally to reset the port level configuration. The default is disabled.

About This Task



Important

Changes you make at the global level override and reset all port level settings.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Enable VLACP globally:

```
vlACP enable
```

Configure VLACP Flap Detect and Damping on a Port

Before You Begin

- To configure VLACP Flap Detect and Damping on sub-ports, enable channelization. For more information about channelization, see [Channelization](#) on page 704.
- As a best practice, to modify an existing VLACP Flap Detect and Damping configuration, first disable VLACP on the interface.

About This Task

Perform the following procedure to control link state changes on VLACP ports. By default, VLACP Flap Detect and Damping is disabled on all VLACP ports.



Important

Do not enable VLACP Flap Detect and Damping on Link Aggregation Control Protocol (LACP) enabled ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable VLACP Flap Detect and Damping on the specified port:

```
vlACP flap-detect enable
```

3. Configure the time interval (in seconds) to record VLACP flaps:

```
vlACP flap-interval <10-600>
```

4. Configure the VLACP flap frequency:

```
vlACP flap-frequency <3-30>
```

5. View VLACP Flap Detect and Damping configuration on the specific port:

```
show vlACP interface gigabitEthernet {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...][slot/all][all]}
```

Example

Configure VLACP Flap Detect and Damping on port 1/3:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 1/3
Switch:1(config-if)#vlACP flap-detect enable
Switch:1(config-if)#vlACP flap-interval 120
Switch:1(config-if)#vlACP flap-frequency 6
Switch:1(config-if)#show vlACP interface gigabitEthernet 1/3
```

VLACP Information									
INDEX	ADMIN ENABLED	OPER ENABLED	PORT STATE	FAST TIME	SLOW TIME	TIMEOUT TIME	TIMEOUT SCALE	ETHER TYPE	MAC ADDR
1/3	false	false	DOWN	200	30000	long	3	0x8103	01:80:c2:00:11:00
VLACP Flap Detect Information									

INDEX	FLAP DETECT	FLAP FREQ	FLAP INTERVAL	TOTAL FLAP	FIRST-FLAP TIME	LAST-FLAP TIME
1/3	enable	6	120	6	01/03/19 15:11:52	01/03/19 15:12:12

Variable Definitions

Use data in the following table to use the **vlacp** command.

Variable	Value
<i>flap-detect enable</i>	Enables VLACP Flap Detect and Damping on the specified port. By default, VLACP Flap Detect and Damping is disabled on all ports.
<i>flap-interval<10-600></i>	Configures VLACP Flap Detect and Damping time interval in seconds. The default value is 60 seconds.
<i>flap-frequency<3-30></i>	Configures the VLACP Flap count permitted during the configured time interval. The default value is 3.

Clear VLACP Flap Detect and Damping Statistics for a Port

Perform the following procedure to clear the VLACP Flap Detect and Damping statistics for a specific VLACP port or all VLACP ports on the switch.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Clear VLACP Flap Detect and Damping statistics:
clear vlapc flap-stats port [{slot/port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]

Example

```
Switch:1>enable
Switch:1#clear vlapc flap-stats port 2/1
```

Variable Definitions

Use data in the following table to use the **clear vlacp flap-stats** command.

Variable	Value
<code>port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

VLACP Configuration using EDM

Configure Virtual LACP (VLACP) to implement link status control protocol at the port level. VLACP cannot interoperate with Link Aggregation Control Protocol (LACP).

Enable VLACP Globally

About This Task

Enable VLACP globally to detect for end-to-end failure. VLACP uses the Hello mechanism of LACP to periodically send Hello packets to ensure there is an end-to-end approach. After Hello packets are not received, the VLACP transitions to a failure state and the port is disabled.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **MLT/LACP**.
3. Select the **VLACP Global** tab.
4. Select **VlacpEnable**.
5. Select **Apply**.

VLACP Global Field Descriptions

Use the data in the following table to use the **VLACP Global** tab.

Name	Description
VlacpEnable	Enables VLACP globally. The default is disabled.

Configure VLACP on a Port

Perform the following procedure to enable VLACP on a port. VLACP periodically checks the end-to-end condition of a point-to-point connection. You can also configure VLACP Flap Detect and Damping to control link state changes on VLACP ports. By default, VLACP Flap Detect and Damping is disabled on all VLACP ports.



Important

Changes made at the global level override and reset all port level settings.

Before You Begin

- To configure VLACP Flap Detect and Damping on sub-ports, enable channelization.
- As a best practice, to modify an existing VLACP or VLACP Flap Detect and Damping configuration, first disable VLACP on the interface.

About This Task

Use the following information to prevent flooding VLACP packets across a defaulted switch:

- Use the default MAC address, 01:80:c2:00:11:00, for end-to-end connections that traverse an intermediate network.
- Use the reserved multicast MAC address 01-80-c2-00-00-0f for directly-connected, peer-to-peer links.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **VLACP** tab.
5. Configure the parameters, as required.
6. Select **Apply**.

VLACP Field Descriptions

Use data in the following table to use the **VLACP** tab.

Name	Description
AdminEnable	Enables VLACP for the port. The default is disabled.
OperEnable	Specifies the VLACP operational status for the port. The default is false.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Use the same value for all LACP-enabled ports. The default value is 200.

Name	Description
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Use the same value for all LACP-enabled ports. The default value is 30000.
Timeout	Specifies the timeout control value. Specify long or short timeout. The default value is long.
TimeoutScale	Assigns the value used to calculate timeout from the periodic time for all VLACP-enabled ports. $\text{Timeout} = \text{PeriodicTime} \times \text{TimeoutScale}$. The default value is 3.
EtherType	Specifies the VLACP identification. The ID is in hexadecimal format. The default value is 0x8103.
EtherMacAddress	Specifies the multicast MAC address exclusively used for VLACPDUs. The default is 01:80:c2:00:11:00.
PortState	Specifies the VLACP port state.
FlapDetectEnable	Enables VLACP Flap Detect and Damping on the port. By default, VLACP Flap Detect and Damping is disabled on all ports.
FlapFrequency	Specifies the number of VLACP flaps detected on the port. The port is shut down if the recorded flaps are within the configured time interval. The default value is 3.
FlapInterval	Specifies the time interval (in seconds) during which the VLACP flaps on the port are monitored. The default value is 60 seconds.
TotalFlapCount	Shows the total number of VLACP flaps detected on the port since you enabled VLACP Flap Detect and Damping or cleared the counters and timers.
FirstFlapTimeStamp	Shows the timestamp of the first VLACP flap detected on the port since you enabled VLACP Flap Detect and Damping or cleared the counters and timers.
LastFlapTimeStamp	Shows the timestamp of the latest VLACP flap detected on the port since you enabled VLACP Flap Detect and Damping or cleared the counters and timers.
FlapClearStats	Clears all VLACP Flap Detect and Damping counters and timers.

Configure VLACP on an Extreme Integrated Application Hosting Port



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to enable Virtual Link Aggregation Control Protocol (VLACP) on an Extreme Integrated Application Hosting (IAH) port. VLACP periodically checks the end-to-end status of a point-to-point connection.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **VLACP** tab.
4. Select **AdminEnable**.
5. Configure the other parameters as required.
6. Select **Apply**.

VLACP Field Descriptions

Use data in the following table to configure the **VLACP** tab.

Name	Description
AdminEnable	Enables VLACP on the Extreme Integrated Application Hosting (IAH) port. The default is disabled.
OperEnable	Specifies the VLACP operational status for the IAH port. The default is false.
FastPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using short timeouts. Set the same value for all LACP enabled IAH ports. The default value is 200.
SlowPeriodicTimer	Specifies the number of milliseconds between periodic transmissions using long timeouts. Set the same value for all LACP enabled IAH ports. The default value is 30000.
Timeout	Specifies the timeout control value. The default value is long.
TimeoutScale	Specifies the value used to calculate timeout duration from the periodic time for all VLACP enabled IAH ports. Timeout = PeriodicTime x TimeoutScale. The default value is 3.

Name	Description
EtherType	Specifies the VLACP protocol identification. The ID is in hexadecimal format. The default value is 0x8103.
EtherMacAddress	Specifies the multicast MAC address exclusively used for VLACPDUs.
PortState	Specifies the VLACP port state.



VLAN Configuration

[VLAN Fundamentals](#) on page 3981

[VLAN Configuration Using CLI](#) on page 4007

[VLAN Configuration using EDM](#) on page 4043

A VLAN is a switched network that is logically segmented by functions, project teams, or applications without regard to the physical location of users. By using a VLAN, you can divide the Local Area Network into smaller groups without interfering with the physical network.

The following topics provide necessary concepts and procedures to configure VLANs. For information about SPBM B-VLANs, see [SPBM B-VLAN](#) on page 1201.

For information about how to enable IGMP Layer 2 querier on a switch in a VLAN, see [IGMP Layer 2 Querier](#) on page 1712.

VLAN Fundamentals

The practical applications of a VLAN include the following:

- You can create VLANs, or workgroups, for common interest groups.
- You can create VLANs, or workgroups, for specific types of network traffic.
- You can add, move, or delete members from these workgroups without making physical changes to the network.

By dividing the network into separate VLANs, you can create separate broadcast domains. This arrangement conserves bandwidth, especially in networks supporting broadcast and multicast applications that flood the network with traffic. A VLAN workgroup can include members from a number of dispersed physical segments on the network, improving traffic flow between them.

The switch performs the Layer 2 switching functions necessary to transmit information within VLANs, as well as the Layer 3 routing functions necessary for VLANs to communicate with one another. You can define a VLAN for a single switch or spanning multiple switches. A port can be a member of multiple VLANs. A VLAN is associated with a spanning tree group.

A VLAN packet is classified before it is forwarded. If the packet matches a classification rule, the port membership is checked. If the port is not an allowed member (potential, static, or active), the system drops the packet.

Port-Based VLANs

A port-based VLAN is a VLAN in which you explicitly configure the ports to be in the VLAN. When you create a port-based VLAN on a device, you assign a VLAN identification number (VLAN ID) and specify the ports that belong to the VLAN. These port members are always active port members. The VLAN ID is used to coordinate VLANs across multiple switches. Any type of frame can be classified to a port-based VLAN.

The example in the following figure shows two port-based VLANs: one for the marketing department, and one for the sales department. Ports are assigned to each port-based VLAN. A change in the sales area can move the sales representative at port 1/1 to the marketing department without moving cables. With a port-based VLAN, you only need to indicate in the Command Line Interface (CLI) that port 1/1 in the sales VLAN now is a member of the marketing VLAN.

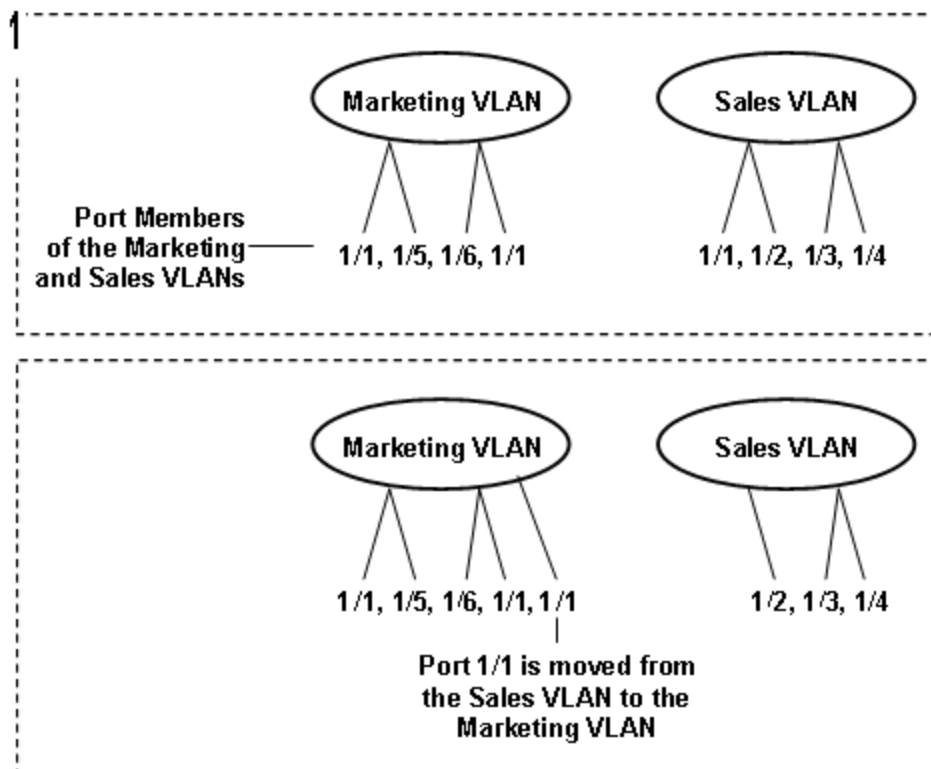


Figure 244: Port-based VLAN

Private VLANs

Table 335: E-Tree and Private VLANs product support

Feature	Product	Release introduced
E-Tree and Private VLANs	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Routing on Private VLANs	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 Not supported on 5320-24T-4X-XT
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.5
	VSP 7400 Series	VOSS 8.5

Private VLANs provide isolation between ports within a Layer-2 service.

The primary and secondary VLAN make the private VLAN. Standard VLAN configuration takes place on the primary VLAN. The secondary VLAN is used to label traffic ingressing isolated ports.

Ports in the private VLAN are configured as isolated, promiscuous, or trunk. The default value is None.

Port Types**Note**

Layer 2 and Layer 3 services based on port types also apply in any routing scenario between two private VLANs or E-trees.

Table 336: Layer 2 services based on port types

Port type	Description
Promiscuous (tagged or untagged ports)	Promiscuous ports communicate with all other ports within the private VLAN. Uses the primary VLAN.
Isolated (tagged or untagged ports)	Isolated ports communicate with the promiscuous ports, but not with any other isolated port. Uses the secondary VLAN.
Trunk (tagged ports)	Trunk ports carry traffic between other port members within the private VLANs. Accepts either primary or secondary VLAN.

Table 337: Layer 3 services based on port types

Port type	Description
Promiscuous (tagged or untagged ports)	Promiscuous ports communicate with all other ports of any Layer 3 VLAN.
Isolated (tagged or untagged ports)	Isolated ports communicate with any promiscuous or regular VLAN ports but not with other isolated ports.
Trunk (tagged ports)	Trunk ports carry traffic between other port members within the private VLANs.

Trunk ports must have VLAN encapsulation enabled. A port can be a single port or can belong to an MLT.

The following figure shows a basic private VLAN topology with private VLAN configured on five switches. All ports connecting to other switches are trunk type ports and all other ports are either promiscuous or isolated ports. On the secondary VLAN, spokes can communicate with hubs, and hubs can communicate with all spokes in the same private VLAN using the primary VLAN, however, spokes cannot communicate with other spokes.




Figure 245: Private VLAN Topology

E-Tree

The E-Tree allows private VLANs to traverse the Shortest Path Bridging MAC (SPBM) network.

For more information about E-Tree and SPBM configuration, see [E-Tree and Private VLAN Topology](#) on page 1187.

Routing on Private VLANs

You can route traffic out of a private VLAN (PVLAN) by applying an IP address to a PVLAN. IP routing is supported on private VLANs for edge switches. With PVLAN packet routing, two hosts located in different private VLANs can communicate over the network.

Packets can be routed over the network between the following:

- regular VLAN and a private VLAN
- different private VLANs

- different I-SIDs (E-tree)

**Note**

Routed private VLANs support Secondary IP Interfaces. For more information, see [Secondary IP Interfaces](#) on page 3992.

The number of private VLANs that you configure with an IP address influences the IPv4 Egress ACE count. For information about scaling information, see *Fabric Engine Release Notes*. Limits are not enforced; either number of private VLANs or number of private VLAN trunk ports can go beyond the recommended values. As a best practice, use the **show io resources filter** command to monitor allocated resources.

**Note**

You cannot enable IP Multicast on a private VLAN interface. **ip spb-multicast enable** and **ip pim enable** commands are not available on a private VLAN.

Private VLAN Configuration Rules

The following are private VLAN rules for the switch:

- Use private VLANs for Layer 2 and Layer 3 services.
- You cannot configure IP Source Guard (IPSG) on ports that are members of private VLANs.
- Do not use the *untag-port default vlan* parameter on private VLAN interfaces that are operating as trunk ports, because it impacts the private VLAN functionality.

Policy-Based VLANs

Received frames are classified into a policy-based VLAN based on certain fields of the frame that matches the associated VLAN policy.

Port Membership Types

In a policy-based VLAN, a port can be designated as a potential member, a static member, or one not allowed to be a member of the VLAN.

If a port is designated as a potential member of the VLAN, and the incoming traffic matches the policy, the system dynamically adds the port to the active port list of the VLAN, making the port an active member of the VLAN. After the system adds a port to the active list, it can remove the port from the active list due to time-out. Potential member ports that join the VLAN are removed (timed out) from the active port list of the VLAN after the timeout (aging time) period expires.

All members of the Spanning Tree Group associated with a protocol-based VLAN are automatically considered potential members of the VLAN. In addition, all tagged ports (trunk ports) become static ports. If you do not want all the tagged ports to be static members of a protocol-based VLAN, put the port in the disallowed list.

Static port members are always members of the VLAN. Static port members are not aged out due to inactivity and they are not removed from the active list. If a server

or router connects to a port, designate that port as a static member of a VLAN. If a server connects to a port that is only a potential member and the server sends very little traffic, a client fails to reach the server if the server port is timed out of the VLAN. As a best practice, make these ports static members of the VLAN.

A disallowed port can never become a member of the VLAN until you add it as a port-member. After you remove a port from the VLAN, the system adds the port to the disallowed list.

On any single spanning-tree instance, an access (untagged) port can belong to one port-based VLAN and many policy-based VLANs. A trunk (tagged) port can belong to many port-based and policy-based VLANs.

The following table describes port membership types for policy-based VLANs.

Table 338: Port membership types for policy-based VLANs

Membership type	Description
Potential	Potential members of a VLAN become active members upon receiving data matching the policy defined for the VLAN (a packet tagged with that VLAN, or an untagged packet matching the policy).
Static (always a member)	Static members are always active members of the VLAN after you configure them as belonging to that VLAN.
Not allowed to join (never a member)	Ports of this type cannot join the VLAN.

The following table lists supported policy-based VLANs.

Table 339: Supported policy-based VLAN types

VLAN type	Support
Protocol-based	supported

Protocol-Based VLANs

Protocol-based VLANs are an effective way to segment your network into broadcast domains according to the network protocols in use.

A port member of a port-based VLAN can belong to multiple protocol-based VLANs. Port tagging is not required for a port to be a member of multiple protocol-based VLANs.

The switch supports IPv6 protocol-based VLAN only.

VLAN Tagging and Port Types

The switch supports the IEEE 802.1Q specification for tagging frames and coordinating VLANs across multiple switches.

Figure 246 shows how an additional four octet (tag) header is inserted in a frame after the source address and before the frame type. The tag contains the VLAN ID associated with the frame.

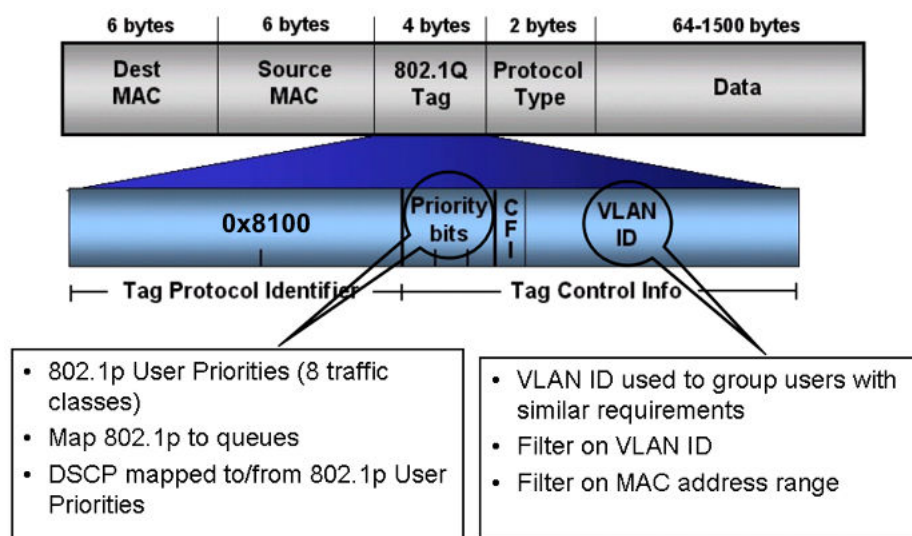


Figure 246: VLAN tag insertion

802.1Q Tagged Ports

Tagging a frame adds four octets to a frame, possibly making it bigger than the traditional maximum frame size. If a device does not support IEEE 802.1Q tagging, it can have problems interpreting tagged frames that it receives.

Whether tagged frames are sent depends on what you configure at the port level. Tagging is configured as true or false for the port, and is applied to all VLANs on that port.

A port with tagging enabled applies the VLAN ID tag to all packets sent on the port. Tagged ports are typically used to multiplex traffic belonging to multiple VLANs to other IEEE 802.1Q-compliant devices.

If you disable tagging on a port, it does not send tagged frames. A nontagged port connects a switch to devices that do not support IEEE 802.1Q tagging. If a tagged frame is forwarded to a port with tagging configured to false, the switch removes the tag from the frame before sending it to the port.

Treatment of Tagged and Untagged Frames

The switch associates a frame with a VLAN based on the data content of the frame and the configuration of the receiving port. The treatment of the frame depends on whether the frame is tagged or untagged.

If a tagged frame is received on a port, if the port is a static or potential member of the VLAN ID specified in the tag, the switch directs it to that VLAN. If the port is not a member of the VLAN that is identified by the tag in the packet, the switch discards the packet. If a port is untagged, you can configure it to discard tagged frames received on the port. In this case the tagged frame is discarded.

For untagged frames, VLAN membership is implied from the content of the frame itself. You can configure a tagged port to accept or discard untagged frames received on the port.

The default VLAN of a port is the VLAN to which untagged frames are classified if they do not match the criteria of any policy-based VLAN of which the port is a member. The default VLAN of the port can be any port-based VLAN a port belongs to, or the unassigned VLAN (1). Frames classified to the unassigned VLAN are discarded.

The frame is forwarded based on the VLAN on which the frame is received, and on the forwarding options available for that VLAN. The switch tries to associate untagged frames with a VLAN in the following order:

- Does the frame belong to a protocol-based VLAN?
- What is the default VLAN for the receiving port?
- Is the default VLAN for the port not the unassigned VLAN?

If the frame meets none of these criteria, it is discarded.

Untagging Default VLAN on a Tagged Port Feature

This feature provides the ability to connect two devices such as an IP phone and a PC to a single port of the switch. Most IP phones ship with an embedded three port switch, and traffic coming from the phone is generally tagged (VLAN ID configured statically or remotely). However, the traffic originating from a PC is usually untagged traffic and must be separated from the IP phone traffic. This separation ensures that broadcast traffic from the PC does not impact voice quality.

After an IP phone is attached to an untagged port, it can fail to register with a remote Internet Telephony Gateway (or equivalent device) dependent on the netmask of the destination IP address (Call Server subnet).

For more information about the Network with IP phone and PC, see the following figure.




Figure 247: Network with IP phone and PC

IP phones and PCs coexist on the same port due to the use of an embedded IP Phone Layer 2 switch. In this scenario if you configure the port as untagged, the egress traffic on this port is untagged and no separation exists between the traffic to the IP phone and the PC. To avoid this condition, the port that connects to the IP phone must be tagged. If the port is tagged, the traffic for the PC is tagged with the default VLAN ID for the port. This configuration creates a problem because the PC does not expect tagged packets. Untag the default VLAN on a tagged port (in this example, port 1/1 that connects to the IP phone) to ensure that the traffic to the PC is sent untagged.

VLAN Router Interfaces

When you configure routing on a VLAN, you assign an IP address to the VLAN, which acts as a virtual router interface address for the VLAN. This IP address is not associated with a physical port. You can reach the VLAN IP address through any of the VLAN port members. Frames are routed to another VLAN IP address within the device. A port can belong to multiple VLANs; some, all, or none can perform routing.

IP routing and VLANs

The switch supports IP routing on the following types of VLANs:

- Port-based VLANs
- IP protocol-based VLANs

VLAN Implementation

This section describes how to implement VLANs and describes default VLANs, the unassigned (NULL) VLAN, and brouter ports. This section also summarizes the defaults and rules regarding VLAN creation on the switch.

- [Default VLAN](#) on page 3991
- [NULL VLAN](#) on page 3991
- [Brouter Ports](#) on page 3991

Default VLAN

Devices are factory-configured so that all ports are in a port-based VLAN called the default VLAN. Because all ports are in the default VLAN, the device functions like a Layer 2 device. The VLAN ID of this default VLAN is always 1, and it is always a port-based VLAN. You cannot delete the default VLAN.

NULL VLAN

Internally, the switch creates a special port-based VLAN called NULL VLAN or unassigned VLAN. This is a place holder VLAN for ports that are not members of any port-based VLAN. When a port is removed from all port-based VLANs, it is added to the NULL VLAN as a port member. Ports can belong to policy-based VLANs as well as to the NULL VLAN. If a frame does not meet the policy criteria and no underlying port-based VLAN exists, the port belongs to the NULL VLAN and the frame is dropped.

Because it is an internal construct, the NULL VLAN cannot be deleted.

Brouter Ports

A brouter port is actually a one-port VLAN with an IP interface. The difference between a brouter port and a standard IP protocol-based VLAN configured to perform routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic and still route IP traffic. Because a brouter port is a single-port VLAN, it uses one VLAN ID. Each brouter port decreases the number of available VLANs by one.

Secondary IP Interfaces

Table 340: Secondary IP Interfaces product support

Feature	Product	Release introduced
Secondary IP Interfaces	4220 Series	Not Applicable
	5320 Series	Fabric Engine 8.10 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	Fabric Engine 8.10
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.10
	VSP 7400 Series	VOSS 8.10

The ability to add Secondary IP Interfaces on a Layer 2 VLAN means that you can have both primary and secondary IP subnets as local routes and increase the number of connected hosts. This feature is useful for subnets that have used all available host IP addresses.

You can configure Secondary IP Interfaces on VLAN interfaces only; you cannot configure a Secondary IP Interface for a brouter port or CLIP interface. Both the Global Router (VRF 0) and user-defined VRFs support Secondary IP Interfaces on associated VLANs. The secondary addresses share the same routeable MAC with the primary IP address.

A VLAN interface can support up to 32 secondary IP addresses. The total number of secondary IP addresses you can create on the same switch depends on available resources. For scaling information, see *Fabric Engine Release Notes*.

DHCP Smart Relay makes DHCP Relay aware of Secondary IP Interfaces. DHCP Smart Relay can also work with a subset of VRRP addresses. For more information, see [DHCP Smart Relay for Secondary IP Interfaces](#) on page 830.

Configuration Rules

- You must configure a primary IP address on the VLAN before you can add secondary addresses.
- You cannot remove a primary IP address if a secondary address exists.
- Primary and secondary subnets must be unique across all VLANs, including within the same VLAN. Subnets cannot overlap, meaning you cannot configure a subnet range that would be included within another subnet range.
- When the command originates on the switch, the **ping** and **traceroute** commands both support secondary IP addresses as the source IP address. If you do not specify a source address as part of the command, the switch automatically uses the a secondary address in the subnet.
- Protocol exchange, both unicast and multicast, occurs only on the primary subnet. Control packets use the primary IP address, including neighbor or peer relationships, and election processes.



Note

IGMP interfaces accept and process all IGMP reports received for subnets that belong to that interface, either primary or secondary.

- Routing protocols do not advertise secondary subnets; you must configure redistribution of direct routes to announce secondary subnets.
- The switch sends all broadcast packets within the VLAN to all hosts regardless of subnet.
- You cannot configure the following features on a secondary subnet:
 - BFD
 - IP security features
 - LLDP
 - Proxy ARP
 - UDP
 - Virtual Inter-Switch Trunk (vIST) peers
- The following feature restrictions exist for Secondary IP Interfaces:
 - You cannot configure a MAC offset for secondary IP addresses.
 - You cannot configure RSMLT on a VLAN with secondary IP addresses.
 - You cannot configure a DvR Gateway IP address that exists in a secondary subnet. DvR-One-IP always selects the primary IP address.
 - You cannot configure a secondary IP address as the source IP for a Fabric Extend tunnel.
 - Enabling Directed Broadcast on a VLAN only applies to the primary IP address on the VLAN and not to Secondary IP Interfaces on the VLAN.

For more information, see the following tasks:

- [Assign an IP Address to a VLAN](#) on page 4013 (CLI)
- [Configure a Secondary IP Interface for a VLAN](#) on page 4053 (EDM)

Network Load Balancing

Table 341: Network Load Balancing product support

Feature	Product	Release introduced
Network Load Balancing (NLB) - multicast operation	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
Network Load Balancing (NLB) - unicast operation	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Microsoft Network Load Balancing (NLB) is a clustering technology available with the Microsoft Windows 2000, Microsoft Windows 2003, Microsoft Windows 2008, and Microsoft Windows 2012 Server family of operating systems. You can use NLB to share the workload among multiple clustering servers. NLB uses a distributed algorithm to load balance TCP/IP network traffic across a number of hosts, enhancing the scalability and availability of the following mission critical, IP-based services:

- web
- VPN
- Streaming media
- Firewalls

NLB also provides high availability by detecting host failures and automatically redistributing traffic to remaining operational hosts.

NLB Considerations and Restrictions

Although the switch interoperates with NLB clusters that operate in Unicast mode and Multicast mode, the following restrictions apply:

- The software does not support true egress mirroring because packets are mirrored prior to the completion of packet processing, so egress mirrored packets can differ from the packets egressing the port.



Note

To mirror the egress traffic, you can use the NEXT-hop device ingress mirroring to capture the egress packets of the switch.

- Inter-VRF routing is not supported between an NLB client and an NLB cluster VLAN in Unicast mode or Multicast mode.
- You must configure NLB to use the same mode as the NLB Server.
- Static ARP entries are not supported for NLB Unicast or NLB Multicast.
- For interoperability with NLB, the switch provides configuration options at the VLAN level.
- ARP entries for NLB server IP addresses do not age out when there is still client traffic coming to the NLB servers, even after the NLB servers are no longer reachable.

NLB Clustering in Unicast Mode

When the cluster is running in NLB unicast mode, all servers in the cluster share a common virtual MAC address, which is 02-bf-x-x-x-x (where x-x-x-x is the cluster IP address in hexadecimal form). All traffic destined to this MAC address is sent to all the servers in the cluster. The virtual MAC address is specified in the Sender MAC Address field of the Address Resolution Protocol (ARP) reply from the cluster to the switch. ARP responses from the switch are sent to the virtual MAC address (rather than to the hardware MAC address).

You can configure the switch for NLB unicast mode support. After you enable the NLB unicast option, the switch floods traffic destined to the cluster IP address to all ports on the VLAN. Unicast mode supports connectivity to a secondary virtual IP address. For information about software scaling capabilities in unicast mode, see *Fabric Engine Release Notes*.

NLB Clustering in Multicast Mode

When the cluster is running in NLB multicast mode, a multicast virtual MAC address with the format 03-bf-x-x-x-x (where x-x-x-x is the cluster IP address in hexadecimal form) is bound to all cluster hosts but the real MAC address of the network adapter is retained. The multicast MAC address is used for client-to-cluster traffic, and the real MAC address of the adapter is used for network traffic specific to the host server.

You can configure the switch for NLB multicast mode support. When you enable NLB multicast mode on a VLAN, the routed traffic destined to the NLB cluster is flooded by default on all ports of the VLAN. All VLANs support multiple cluster IPs by default. You

can connect up to 200 NLB clusters to a single VLAN. For information about software scaling capabilities, see .



Note

Shortest Path Bridging MAC (SPBM) supports NLB Unicast and Multicast modes. For more information on SPBM, see [Fabric Basics and Layer 2 Services](#) on page 1178.

Supported NLB Topologies

The switch supports Network Load Balancing (NLB) in the following topologies.

Supported NLB Topology–Example 1

The switch supports NLB when the NLB Cluster connections use a different physical port on the switch than the NLB clients.

The following figure illustrates this configuration where the NLB Server and the NLB Client workstations connect to different aggregation switches, which connect to the switch using different VLANs and different ports.

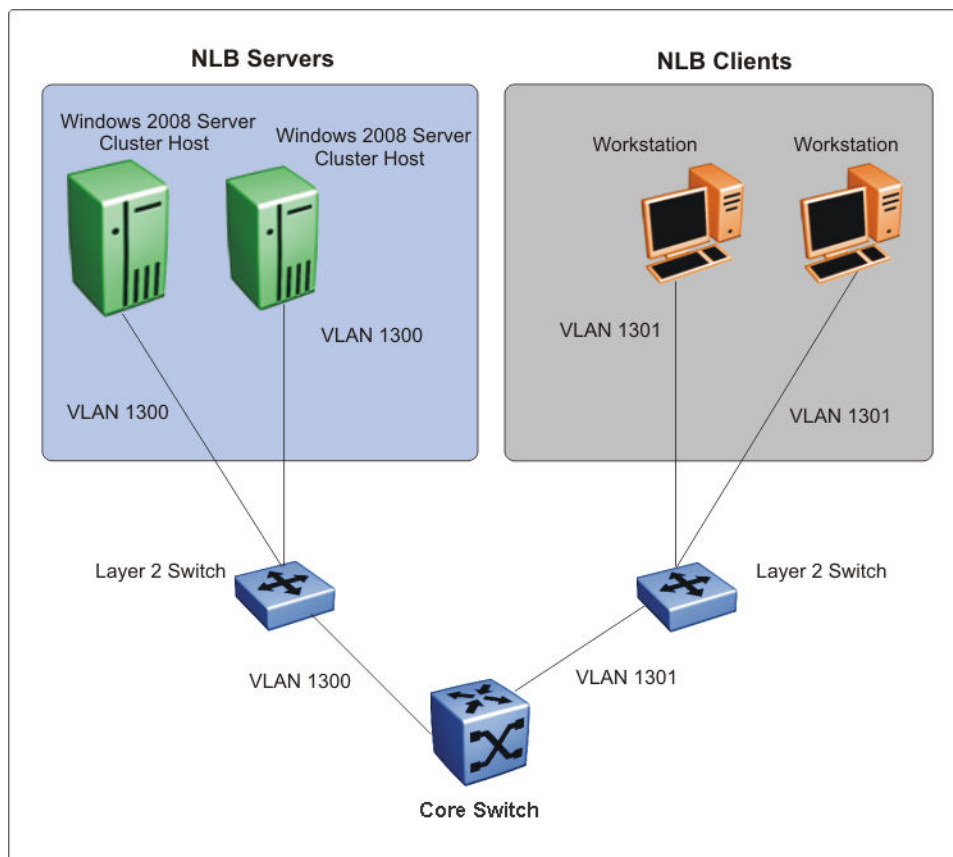


Figure 248: Supported NLB topology–example 1

Supported NLB Topology–Example 2

The switch also supports the following topology where the NLB Server and the NLB Client workstations connect to the same aggregation switch and then connect to the switch using the same port.



Note

The switch supports Layer 3 routing between an NLB-enabled VLAN and another VLAN on the same port.

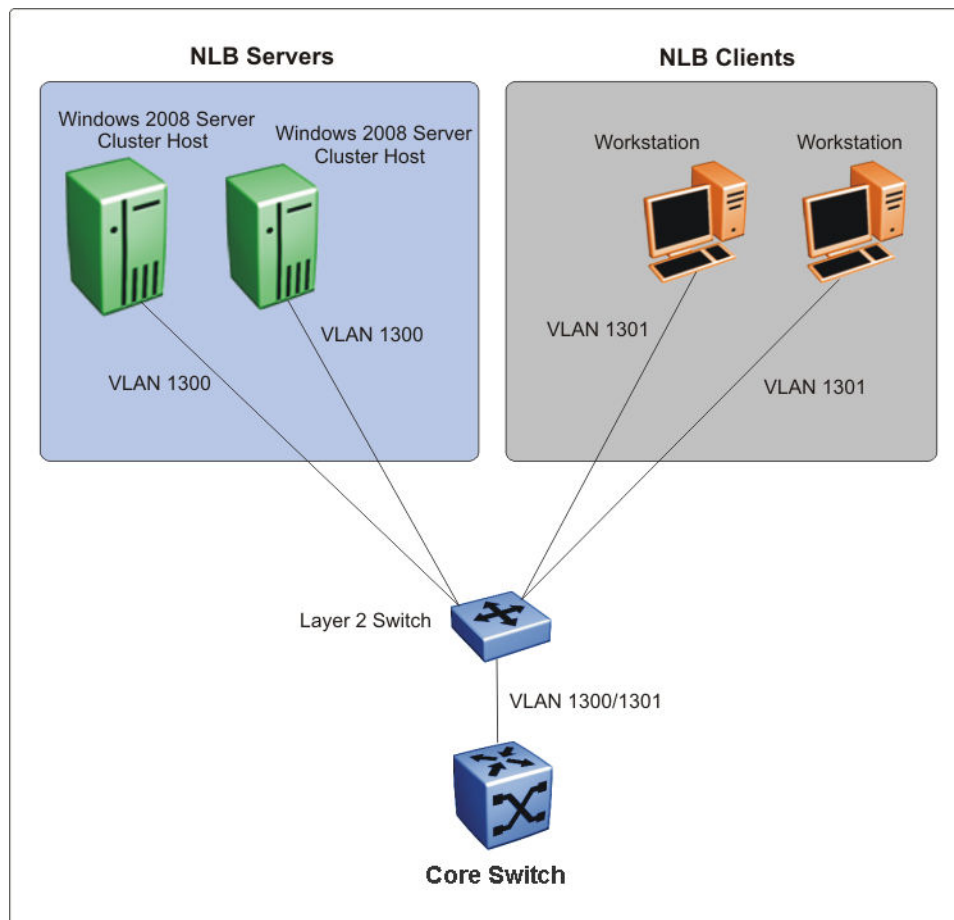


Figure 249: Supported NLB topology–example 2

Other Supported NLB Topologies

The switch supports NLB in the following other topologies:

- NLB cluster hosts and clients are connected to Layer 2 Ethernet switches that are SMLT connected to the SMLT cluster.
- NLB cluster hosts are directly connected and distributed between the switches and the clients are connected to Layer 2 Ethernet switch that is SMLT connected to the SMLT cluster.

- NLB cluster hosts and clients are directly connected and distributed between the switches in the SMLT cluster.
- NLB cluster hosts and clients are connected to Layer 2 Ethernet switches that are SMLT connected to the SMLT cluster core.

**Note**

For more information on these topologies, see *Technical Configuration Guide for Microsoft Network Load Balancing*.

NLB and Directed Broadcast Resource Limits

NLB and Directed Broadcast consume resources from the same pool of 200 resources. When you configure either NLB or Directed Broadcast, the switch uses one resource. If you configure both NLB and Directed Broadcast, the switch uses two resources.

To avoid a situation where there is a lack of resources, adhere to the following limits:

- The number of NLB cluster IP interfaces multiplied by the number of configured clusters must be equal to, or less than, 200. The number of NLB cluster IP interfaces is the key, not the number of VLANs. You can configure 1 VLAN with up to 200 NLB cluster IP interfaces or configure up to 200 VLANs with 1 NLB cluster IP interface per VLAN.

For example: 1 NLB cluster IP interface x 200 clusters = 200 or 2 NLB cluster IP interfaces x 100 clusters = 200

- If you configure VLANs with Directed Broadcast only, you can scale up to 200 VLANs.
- If you configure VLANs with both NLB and Directed Broadcast, you can only scale up to 100 VLANs assuming there is only 1 NLB cluster IP interface per VLAN.

For information on Directed Broadcast, see [Denial-of-Service Attack Prevention](#) on page 3372.

VLAN MAC-layer Filtering Database and MAC Security

Table 342: FDB Protected by Port product support

Feature	Product	Release introduced
FDB protected by port (MAC security limit-learning)	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.3
	5720 Series	Fabric Engine 8.7
	7520 Series	Not Supported
	7720 Series	Not Supported
	7830 Series	Not Supported
	VSP 4900 Series	VOSS 8.3
	VSP 7400 Series	Not Supported

To perform MAC-layer bridging, the device must know the destination MAC-layer address of each device on each attached network to forward packets to the appropriate destination. The system stores MAC-layer addresses in the bridge forwarding database (FDB) table, and can forward packet traffic based on the destination MAC-layer address information.

MAC Security

Use MAC security to control traffic from specific number of MAC addresses. You can also limit the number of allowed MAC addresses. You can enable this feature at the port level.

Port-level security applies to traffic for all VLANs received on that port.

Port-level MAC security provides limit—learning option:

- **limit-learning:** This option protects the FDB from traffic from too many MAC addresses, which fill the FDB table.

This option limits the number of MAC addresses a port learns. You can specify a maximum number of addresses. After the number of addresses reaches the maximum, learning stops. The port disables packet forwarding and drops packets from new source MAC address. MAC address learning resumes after enough existing addresses age out and there is room to learn new MAC addresses.



Note

If you configure a limit on a port that has already learned more than the new limit, packet forwarding for those additional MAC addresses continues to work until the port flaps, you flush the MAC address, or the MAC address disappears.

Restrictions

The following list identifies restrictions to MAC security limit-learning:

- This feature is not supported on:
 - MLT ports members
 - NNI ports
 - Transparent Port UNI or Switched UNI
- The switch supports MAC learning only in the VLAN domain; it is not supported in the I-SID domain.
- The port MAC limit does not count static MAC addresses.
- The switch supports the maximum number of MAC addresses a port can learn for non-SPBM configurations.

Prevention of IP Spoofing within a VLAN

VLAN IP as the default gateway

You can prevent VLAN logical IP spoofing by blocking the external use of the device IP address. A configurable option is provided, for each port, which detects a duplicate IP address (that is, an address that is the same as the device VLAN IP address) and blocks all packets with a source or destination address equal to that address.

If an ARP packet is received that has the same source IP address as the logical VLAN IP address of the receiving port, all traffic coming to that port (with this MAC address as source/destination address) is discarded by the hardware. After detecting a duplicate IP address, the device sends a gratuitous ARP packet to inform devices on the VLAN about the correct MAC address for that IP address. You can specify a time on a configurable global timer after which the MAC discard record is deleted, and the device resumes accepting packets from that MAC address.

VRRP IP as the default gateway

Similarly, you can prevent VRRP IP spoofing by blocking the external use of the virtual IP address. A configurable option is provided, for each port, which detects a duplicate IP address (that is, an address that is the same as the device virtual IP address) and blocks all packets with a source or destination address equal to that address.

If an ARP packet is received that has the same source IP address as the virtual IP address of the receiving port, all traffic coming to that port (with this MAC address as source/destination address) is discarded by the hardware. After detecting a duplicate IP address, the device sends a gratuitous ARP packet to inform devices on the VRRP subnet about the correct virtual router MAC address for that IP address. You can specify a time on a configurable global timer after which the MAC discard record is deleted, and the device resumes accepting packets from that MAC address.

Packet spoofing

You can stop spoofed IP packets by configuring the switch to forward only IP packets that contain the correct source IP address of your network. By denying all invalid source

IP addresses, you minimize the chance that your network is the source of a spoofed DoS attack.

A spoofed packet is one that comes from the Internet into your network with a source address equal to one of the subnet addresses on your network. The source address belongs to one of the address blocks or subnets on your network. To provide spoofing protection, you can use a filter that examines the source address of all outside packets. If that address belongs to an internal network or a firewall, the packet is dropped.

To prevent DoS attack packets that come from your network with valid source addresses, you need to know the IP network blocks in use. You can create a generic filter that:

- Permits valid source addresses
- Denies all other source addresses

To do so, configure an ingress filter that drops all traffic based on the source address that belongs to your network.

If you do not know the address space completely, it is important that you at least deny private (see RFC1918) and reserved source IP addresses. The following table lists the source addresses to filter.

Table 343: Source addresses to filter

Address	Description
0.0.0.0/8	Historical broadcast. High Secure mode blocks addresses 0.0.0.0/8 and 255.255.255.255/16. If you enable this mode, you do not need to filter these addresses.
10.0.0.0/8	RFC1918 private network
127.0.0.0/8	Loopback
169.254.0.0/16	Link-local networks
172.16.0.0/12	RFC1918 private network
192.0.2.0/24	TEST-NET
192.168.0.0/16	RFC1918 private network
224.0.0.0/4	Class D multicast
240.0.0.0/5	Class E reserved
248.0.0.0/5	Unallocated
255.255.255.255/32	Broadcast ¹

You can also enable the spoof-detect feature on a port.

VLAN loop prevention

Table 344: Simple Loop Prevention Protocol product support

Feature	Product	Release introduced
Simple Loop Prevention Protocol (SLPP)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Loop Prevention

Under certain conditions, such as incorrect configurations or cabling, loops can form. This is true mainly for Layer 2 bridged domains, such as VLANs.

Simple Loop Prevention Protocol (SLPP) provides active protection against Layer 2 network loops on a per-VLAN basis. SLPP uses a lightweight hello packet mechanism to detect network loops. Sending hello packets on a per VLAN basis allows SLPP to detect VLAN based network loops for untagged as well as tagged IEEE 802.1Q VLAN link configurations. After SLPP detects a loop, the port is shutdown.



Note

If you use SLPP in a vIST environment, you must enable it on both vIST peers. When an SLPP packet of a vIST peer is looped through UNI ports to the other device, that device will shut down its UNI port due to receiving SLPP packets from its peer. A device's own SLPP packets will go over a vIST connection but will not be forwarded by its vIST peer back onto its UNI ports.

Configure the SLPP functionality with the following criteria:

- **SLPP TX Process** – You decide on which VLANs a switch can send SLPP hello packets. The packets are then replicated out all ports which are members of the SLPP-enabled VLAN. As a best practice, enable SLPP on all VLANs.
- **SLPP RX Process** – You decide on which ports the switch can act when receiving an SLPP packet that is sent by the same switch or by its SMLT peer. You must enable this process only on Access SMLT ports. You can enable this process only when the design permits on SMLT CORE ports in the case of a square/full mesh core design.
- **SLPP Action** – The action operationally disables the ports receiving the SLPP packet. You can also tune the network failure behavior. You can choose how many SLPP packets a port needs to receives before a switch takes an action. You need to stagger these values to avoid edge switch isolation.

Instead of manually recovering a port that is shutdown by SLPP, you can configure an automatic recovery delay per port. The **auto-recover-delay** command specifies the minimum amount of time after which the port is brought back up. The maximum amount of time that can pass until the port is brought back up is two times the configured **auto-recover-delay** value.

Loops can be introduced into the network in many ways. One way is through the loss of an MLT/link aggregation configuration caused by user error or malfunctioning equipment. This scenario does not always introduce a broadcast storm, but because all MAC addresses are learned through the looping ports, does significantly impact Layer 2 MAC learning. Spanning Tree cannot, in all cases, detect such a configuration issue, whereas SLPP reacts and disables the malfunctioning links and limits network impact to a minimum.

The desire is to prevent a loop from causing network problems, while also attempting not to totally isolate the edge where the loop was detected. Total edge closet isolation is the last resort to protect the rest of the network from the loop. With this in mind, some administrators adopt the concept of an SLPP primary switch and SLPP secondary switch. These are strictly design terms and are not configuration parameters. The Rx thresholds are staggered between the primary and secondary switch. Therefore, the primary switch disables an uplink immediately upon a loop occurring. If this resolves the loop issue, then the edge closet still has connectivity back through the SLPP secondary switch. If the loop is not resolved, then the SLPP secondary switch disables the uplink and isolates the closet to protect the rest of the network from the loop.

As the number of VLANs running SLPP scale off of a specific uplink port, increase the Rx-threshold value to prevent complete isolation of the offending edge. The primary goal of SLPP is to protect the core at all costs. In certain loop conditions, what can occur is the secondary switch also detects the loop and SLPP Rx-threshold of the secondary switch is reached before the primary can stop the loop by taking its port down. Therefore, both switches eventually take their ports down and the edge is isolated. The larger the number of VLANs associated with the port, the more likely this can occur, especially for loop conditions that affect all VLANs.



Important

The loop detection functionality of the device must not be used under normal operating conditions. Only use it if directed by technical support personnel.

You cannot configure the EtherType for SLPP. The switch uses an EtherType of 0x8102 .

Spanning Tree and Protection against Isolated VLANs

Virtual Local Area Network (VLAN) isolation disrupts packet forwarding. The following figure illustrates the problem. Two VLANs (V1 and V2) connect four devices, and both VLANs are in the same spanning tree group. V2 includes three of the four devices, whereas V1 includes all four devices. After a spanning tree protocol detects a loop, it blocks the link with the highest link cost. In this case, the 100 Mbps link is blocked, which isolates a device in V2. To avoid this problem, either configure V2 on all four

devices or use MSTP with a different Multiple Spanning Tree Instance (MSTI) for each VLAN.




Figure 250: VLAN isolation

Switched UNI Layer 3

Create a platform VLAN using the command **vlan create <vlan-id> type port-mstprsp <msti-instance>**. Enable Layer 3 services on the platform VLAN and is associated with the Switched UNI (S-UNI) Service Instance Identifier (I-SID). All S-UNI ports are added to the platform VLAN.

You must associate the S-UNI I-SID to the platform VLAN. After you associate the platform VLAN with the I-SID, it becomes a CVLAN.

The switch performs MAC and ARP learning on the platform VLAN.



Note

You cannot add S-UNI ports or MLT to the S-UNI platform VLANs directly. Add the ports to the I-SID and assign the I-SID to the platform VLAN.
You can associate only port based VLAN with S-UNI I-SID.

VLAN Feature Support

The following table summarizes supported features.

For the latest scalability information, see *Fabric Engine Release Notes*.

Table 345: VLAN support

Feature	Description
Number of VLANs	4059
Port-based VLANs	Supported
Policy-based VLANs <ul style="list-style-type: none"> Protocol-based SPBM-based 	Supported
IEEE 802.1Q tagging	Supported
IP routing and VLANs	Supported
Special VLANs <ul style="list-style-type: none"> Default VLAN Null VLAN Router ports Private VLAN 	Supported

NEW! VLAN IDs

The following table provides information on VLAN ID availability.

Table 346: VLAN ID allocation

VLAN ID	Purpose
1	Default VLAN
2-4059	User-configurable VLANs
3500-3998	Reserved range on switches that enable the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags. Note: <i>vrf-scaling</i> does not apply to switch models that support a single active VRF with IP configuration, or to 5320-16P-2MXT-2X and 5320-24T-4X-XT.
4047	SD-WAN-associated VLAN

Table 346: VLAN ID allocation (continued)

VLAN ID	Purpose
4048-4049	Onboarding Private VLAN pair (primary/secondary)
4051-4052	SPBM B-VLANs
4060-4094	System-reserved VLANs

VLAN Configuration Rules

The following are VLAN rules for the switch:

- A tagged port can belong to multiple VLANs in multiple Spanning Tree Groups.
- Under the default configuration, the default Spanning Tree Group is number 1 if the chassis configuration permits multiple STGs.
- An untagged port can belong to only one port-based VLAN.
- You can configure only one protocol-based VLAN for a given protocol.
- The VLAN membership of a frame is determined by the following order of precedence, if applicable:
 1. IEEE 802.1Q tagged VLAN ID
 2. protocol-based VLAN
 3. port-based VLAN default VLAN of the receiving port
- You cannot configure a VLAN name that uses all numbers, for example, 222.

Interaction Between Anycast IP Gateway, Distributed Virtual Routing and SLPP Guard

On VLANs that use Anycast IP Gateway or Distributed Virtual Routing (DvR), the switch sends a small control packet every five seconds on UNI ports. Similar to VRRP hello packets, these packets ensure that the Anycast IP Gateway or DvR MAC address remains active on Layer 2 switches that connect to the SPB Fabric, such as those using Fabric Attach (FA).

Although these packets share the same EtherType as Simple Loop Prevention Protocol (SLPP), they differ in how the MAC addresses are constructed:

- Standard SLPP packets:
 - Source MAC: Chassis MAC address
 - Destination MAC: Chassis MAC address with the multicast bit set
- Anycast IP Gateway/DvR keep-alive packets:
 - Source MAC: Anycast IP Gateway or DvR MAC address
 - Destination MAC: Same Anycast IP Gateway or DvR MAC address (with multicast bit applied)

SLPP-RX does not use these packets but SLPP Guard does recognize them. When you enable SLPP Guard on a port, it can detect these packets, and take the necessary

action to prevent unwanted loops, and thus does not require you to configure SLPP-TX on the VXLAN.

SLPP-TX/RX can operate in parallel with these packets without restrictions, and the two mechanisms do not interfere with each other.

VLAN Configuration Using CLI

This section provides procedures to configure and to manage Virtual Local Area Networks (VLANs) using the command line interface (CLI).

Create a VLAN

Create a VLAN by port, protocol, or an SPBM B-VLAN. Optionally, you can choose to assign the VLAN a name and color.

Assign an IP address to the VLAN. You can also assign a MAC-offset value that allows you to manually change the default MAC address.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create a VLAN by specifying one of the following VLAN types:
 - a. Create a VLAN by port:

```
vlan create <2-4059> [name WORD<0-64>] type port-mstprstp <0-63> [color <0-32>]
```
 - b. Create a VLAN using a user-defined protocol and specify the frame encapsulation header type:

```
vlan create <2-4059> [name WORD<0-64>] type protocol-mstprstp <0-63> ipv6 [color <0-32>]
```
 - c. Create an SPBM B-VLAN:

```
vlan create <2-4059> [name WORD<0-64>] type spbm-bvlan [color <0-32>]
```
3. (Optional) Associate a CVLAN I-SID to the platform VLAN.

```
vlan i-sid <1-4059> <1-16777215> [force]
```
4. Log on to the VLAN Interface Configuration mode for the VLAN ID:

```
interface VLAN <1-4059>
```
5. Assign an IP address to a VLAN:

```
ip address {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>} [dvr-one-ip] [state-disabled]
```
6. (Optional) Assign a Secondary IP Interface to a VLAN:

```
ip address {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>} secondary
```
7. (Optional) Specify the MAC-offset value:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D> dvr-one-ip <0-511>
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#vlan create 2 type port-mstprstp 0 color 4
Switch:1(config)#vlan i-sid 2 100
Switch:1(config)#interface vlan 2
Switch:1(config-if)#ip address 192.0.2.0/24 dvr-one-ip
```

Variable Definitions

The following table defines parameters for the **vlan create** command.

Variable	Value
<code><2-4059></code>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.
<code>name WORD<0-64></code>	Specifies the VLAN name. The name attribute is optional.
<code>type port-mstprstp <0-63> [color <0-32>]</code>	Creates a VLAN by port: <ul style="list-style-type: none"> <code><0-63></code> is the STP instance ID from 0 to 63. <code>color <0-32></code> is the color of the VLAN in the range of 0 to 32. Note: MSTI instance 62 is reserved for SPBM if SPBM is enabled on the switch.
<code>type pvlan-mstprstp <0-63> [color <0-32>]</code>	Creates a private VLAN by port: <ul style="list-style-type: none"> <code><0-63></code> is the STP instance ID from 0 to 63. <code>color <0-32></code> is the color of the VLAN in the range of 0 to 32.
<code>type protocol-mstprstp <0-63> ipv6</code>	Creates a VLAN by protocol: <ul style="list-style-type: none"> <code><0-63></code> is the STP instance ID. <code>color <0-32></code> is the color of the VLAN in the range of 0 to 32.
<code>type spbm-bvlan</code>	Creates a SPBM B-VLAN.

The following table defines parameters for the **ip address** command.

Variable	Value
<code><A.B.C.D/X> <A.B.C.D> <A.B.C.D></code>	Specifies the IP address and subnet mask in the format A.B.C.D/X or A.B.C.D A.B.C.D.
<code>dvr-one-ip</code>	Specifies that the IP address will be used as the DvR gateway IP address and will be used by all other DvR Controllers for the DvR VLAN subnet.
<code><0-511></code>	Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated. Note: You cannot specify the mac-offset for a given secondary IP.
<code>secondary</code>	Identifies the IP address as a Secondary IP Interface.
<code>state-disabled</code>	Creates the IP interface in an administratively disabled state. This parameter is useful if you want to preconfigure an IP interface, including the routing protocol configuration for that interface, and enable it later during a maintenance window. This parameter applies only to the primary IP address. By default, the switch enables an IP interface after you configure the primary IP address.

The following table defines parameters for the **vlan i-sid** command.

Variable	Value
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<code><0-16777215></code>	Specifies the service instance identifier (I-SID). You cannot use I-SID 0x00ffffff. The system reserves this I-SID to advertise the virtual BMAC in an SMLT dual-homing environment. This value is the same for the primary and secondary VLANs.
<code>force</code>	Specifies the software must replace the existing VLAN-to-I-SID mapping, if one exists.

Create a Private VLAN

About This Task

You can create a private VLAN and set the port type. The primary and secondary VLAN IDs are associated with the same MTSI, the secondary VLAN inherits the primary VLAN configuration. You cannot create another VLAN with the same VLAN ID as the secondary VLAN. The secondary VLAN cannot be any other type of VLAN other than a secondary VLAN.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Create a private VLAN:

```
vlan create <2-4059> type pvlan-mstprstp secondary <2-4059>
```
3. Specify a name for the VLAN:

```
vlan create <2-4059> name
```


4. Enter GigabitEthernet Interface Configuration mode:

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

**Note**

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

5. Set the port type:

```
private-vlan <isolated|promiscuous|trunk>
```

**Note**

If the port is a member of an MLT, the port inherits the private VLAN port type of the MLT.

6. Exit to Global Configuration mode:

```
exit
```

7. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

8. Add ports to the primary VLAN:

```
vlan members add <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#vlan create 2 type pvlan-mstprstp 6 secondary 5
Switch:1(config)#interface gigabitethernet 1/36
Switch:1(config-if)#private-vlan isolated
Switch:1(config-if)#exit
Switch:1(config)#interface vlan 2
Switch:1(config-if)#vlan members add 2 1/36
```

Variable Definitions

Use the data in the following table to use the **vlan create** command.

Variable	Value
<code><2-4059></code>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.
<code>name WORD<0-64></code>	Specifies the VLAN name. The name attribute is optional.
<code>type pvlan-mstp <0-63></code>	Creates a private VLAN by port. The variable <code><0-63></code> is the STP instance ID from 0 to 63. Note: MSTI instance 62 is reserved for SPBM if SPBM is enabled on the switch.
<code>secondary<2-4059></code>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.

Use the data in the following table to use the `private vlan port type` command.

Variable	Value
<code><isolated promiscuous trunk></code>	Specifies the port type. If not specified, the port type defaults to None. <ul style="list-style-type: none"> Isolated: An Isolated port can belong only to one private VLAN Promiscuous: A Promiscuous port can belong to many private VLANs Trunk: A Trunk port can belong to many private VLANs, is tagged, and can also belong to non-private VLANs
<code>no private-vlan</code>	Port defaults to type None.
<code>default private-vlan</code>	Port defaults to type None.



Note

If there are other non-private VLANs using the defined port, the following message is displayed: All non private VLANs using this interface will be removed once this port becomes a member of a private VLAN. Ports with private-vlan type of isolated or promiscuous may only contain private VLANs. Do you wish to continue (y/n) ?

Use the data in the following table to use the `interface vlan` and `vlan members add` commands.

Variable	Value
<code><1-4059></code>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <code>vrf-scaling</code> and <code>spbm-config-mode</code> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

Assign an IP Address to a VLAN

Assign an IP address to a VLAN so that it supports routing operations.

Before You Begin

- You must create the VLAN.
- Activate IP forwarding globally.
- If you use this procedure to create a Secondary IP Interface, you must first create the primary IP address.

About This Task

If an IP interface is configured without specifying the VRF instance, it maps to VRF 0 by default.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Assign a primary IP address to a VLAN:

```
ip address {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>} [dvr-one-ip] [name WORD  
<0-64>] [state-disabled]
```

3. (Optional) Assign a Secondary IP Interface to a VLAN:

```
ip address {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>} secondary
```

4. (Optional) If required, associate the VLAN with a VRF:

```
vrf WORD<1-16>
```

Example

Add a primary IP address to the VLAN and configure it as the DvR-One-IP interface. Give the address a name.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface vlan 2
Switch:1(config-if)#ip address 192.0.2.5 255.255.255.0 dvr-one-ip name Boston
```

Variable Definitions

The following table defines parameters for the **ip address** command.

Variable	Value
<A.B.C.D/X> <A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask in the format A.B.C.D/X or A.B.C.D A.B.C.D.
dvr-one-ip	Specifies that the IP address will be used as the DvR gateway IP address and will be used by all other DvR Controllers for the DvR VLAN subnet.
name WORD <0-64>	Specifies the name associated with the IP address on a VLAN.

Variable	Value
<i>secondary</i>	Identifies the IP address as a Secondary IP Interface.
<i>state-disabled</i>	Creates the IP interface in an administratively disabled state. This parameter is useful if you want to preconfigure an IP interface, including the routing protocol configuration for that interface, and enable it later during a maintenance window. This parameter applies only to the primary IP address. By default, the switch enables an IP interface after you configure the primary IP address. Note: This parameter is not saved in the configuration file.

The following table defines parameters for the **vrf** command.

Variable	Value
<i>WORD<0-16></i>	Specifies the VRF of the VLAN.

Display All Configured IP Addresses

About This Task

Use this procedure to verify the configuration of all IP addresses, including Secondary IP Interfaces on VLANs. You cannot create Secondary IP Interfaces on brouter ports.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display all configured IP addresses:
`show ip address`
3. Display all primary IP addresses:
`show ip address primary`
4. Display all secondary IP addresses:
`show ip address secondary`
5. Display IP addresses associated with a VLAN:
`show ip address vlan [vlan-id[-vlan-id] [, ...]]<1-4059>`
6. Display IP addresses associated with a VRF name:
`show ip address vrf WORD<1-16>`
7. Display IP addresses associated with a VRF ID:
`show ip address vrfids WORD<0-512>`

View the IPv4 and IPv6 Interface Configuration on a VLAN

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Verify the interface configuration on the VLAN:

For IPv4, enter: `show interfaces vlan ip`

For IPv6, enter: `show ipv6 interface vlan`

Example

Verify the IPv4 interface configuration on the VLAN:

```
Switch:1(config-if)# show interfaces vlan ip
```

Vlan Ip										
VLAN ID	VRF NAME	IP ADDRESS	NET MASK	ADMIN STATUS	OPER STATUS	ADVERTISE WHEN_DOWN	DIRECTED BROADCAST	RPC	RPCMODE	RMON
1050	Globa~	192.0.2.9	255.255.255.0	enable	up	disable	disable	disable	exist-only	disable
1102	Globa~	198.51.100.1	255.255.255.0	enable	up	disable	disable	disable	exist-only	disable
1133	iir3	192.0.2.10	255.255.255.0	enable	up	disable	disable	disable	exist-only	disable
1500	spboip	192.0.2.11	255.255.255.0	enable	up	disable	disable	disable	exist-only	disable
1590	spboip	198.51.100.2	255.255.255.0	enable	up	disable	disable	disable	exist-only	disable
4057	Globa~	192.0.2.12	255.255.255.0	enable	up	disable	disable	disable	exist-only	disable

All 16 out of 16 Total Num of Vlan Ip Entries displayed

VLAN ID	VRF NAME	BCASTADDR FORMAT	REASM MAXSIZE	DVR STATE	DVR GW-IPv4
1050	GlobalRouter	ones	1500	disable	-
1102	GlobalRouter	ones	1500	disable	-
1133	iir3	ones	1500	disable	-
1500	spboip	ones	1500	disable	-
1590	spboip	ones	1500	disable	-
4057	GlobalRouter	ones	1500	disable	-

All 16 out of 16 Total Num of Vlan Ip Entries displayed

Verify the IPv6 configuration on the port:

```
Switch:1(config-if)# show ipv6 interface vlan
```

Vlan Ipv6 Interface													
IFINDX INDX	VLAN	PHYSICAL ADDRESS	ADMIN STATE	OPER STATE	TYPE	MTU	HOP LMT	REACHABLE TIME	RETRANSMIT TIME	MCAST STATUS	IPSEC	RPC	RPCMODE
3170	1122	2c:f4:c5:dc:b4:89	enable	up	ETHER	1500	64	30000	1000	disable	disable	disable	existonly
3174	1126	2c:f4:c5:dc:b4:8b	enable	up	ETHER	1500	64	30000	1000	disable	disable	disable	existonly
3185	1137	2c:f4:c5:dc:b4:90	enable	up	ETHER	1500	64	30000	1000	disable	disable	disable	existonly

Vlan Ipv6 Address					
IPV6 ADDRESS		VLAN-ID	TYPE	ORIGIN	STATUS

```

-----
2001:db8:0:0:0:0:0:1          V-1122      UNICAST  MANUAL    PREFERRED
2001:db8:0:0:0:2ef4:c5ff:fedc:b489  V-1122      UNICAST  LINKLAYER PREFERRED
2001:db8:0:0:0:0:0:1          V-1126      UNICAST  MANUAL    PREFERRED
2001:db8:0:0:0:2ef4:c5ff:fedc:b48b  V-1126      UNICAST  LINKLAYER PREFERRED
2001:db8:0:0:0:0:0:1          V-1137      UNICAST  MANUAL    PREFERRED
2001:db8:0:0:0:2ef4:c5ff:fedc:b490  V-1137      UNICAST  LINKLAYER PREFERRED

3 out of 4 Total Num of Interface Entries displayed.
6 out of 7 Total Num of Address Entries displayed.

```

Remove All Secondary IP Interfaces From a VLAN

Perform this procedure to remove all secondary IP addresses on a VLAN simultaneously.

About This Task

To remove specific secondary IP addresses from a VLAN, use the **no** operator with the **ip address {A.B.C.D} secondary** command on the VLAN.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Remove all Secondary IP Interfaces from a VLAN:
clear secondary-ips vlan <1-4059>

Perform a General VLAN Action

Perform a general VLAN action to initiate a specific function on a VLAN, such as clearing learned MAC addresses or ARP entries from the forwarding database by performing this procedure.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Perform a general VLAN action:
vlan action <1-4059> {none|flushMacFdb|flushArp|flushIp|flushDynMemb|triggerRipUpdate|all}

Example

Perform a general VLAN action:

```

Switch(config)#vlan action 1 none
Switch(config)#vlan action 1 flushMacFdb
Switch(config)#vlan action 1 flushIp
Switch(config)#vlan action 1 flushDynMemb

```

Variable Definitions

Use the data in the following table to use the **vlan action** command.

Variable	Value
<i>none</i>	Configures action to none. This action performs no updates.
<i>flushMacFdb</i>	Configures action to flushMacFdb. This action removes the learned MAC addresses from the forwarding database for the selected VLAN.
<i>flushArp</i>	Configures action to flushArp. This action removes the ARP entries from the address table for the selected VLAN.
<i>flushIp</i>	Configures action to flushIp. This action removes the learned IP addresses from the forwarding table for the selected VLAN.
<i>flushDynMemb</i>	Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN, and removes MAC addresses learned on those ports for this VLAN.
<i>flushDynMemb</i>	Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN, and removes MAC addresses learned on those ports for this VLAN.
<i>triggerRipUpdate</i>	Configures action to triggerRipUpdate.
<i>all</i>	Configures action to all and performs all preceding actions.

Configure Static MAC Addresses for a VLAN

Configure the static MAC address parameters.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Configure a static MAC address of a VLAN:

```
vlan mac-address-static <1-4059> <0x00:0x00:0x00:0x00:0x00:0x00>
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}
```


Example

Configure a static MAC address of a VLAN:

```
Switch(config)#vlan mac-address-static 1 0x00:0x00:0x00:0x00:0x00:0x01 1/1
```

Variable Definitions

Use the data in the following table to use the **vlan mac-address-static** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<0x00:0x00:0x00:0x00:0x00:0x00>	Indicates the MAC address.
{slot/port[/sub-port] [-slot/port[/sub-port]][, ...] [slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Limit MAC Address Learning

Configure the MAC security feature to control traffic from specific number of MAC addresses. The total number of MAC addresses that you can configure are fixed. The switch help text shows the maximum MAC addresses a port can learn in non-SPBM configurations. In an SPBM configuration, the maximum value is reduced by half.

About This Task

This feature limits the number of forwarding database (FDB) entries learned on a particular port to a user-specified value. After the number of learned forwarding database entries reaches the maximum limit, MAC learning stops on that port.



Note

This procedure does not apply to 7830 Series.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface mlt <1-512>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Protect the FDB from hits by too many MAC addresses:

```
mac-security [port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]] limit-learning enable [max-addr <1-64000>]
```

Example

Protect the FDB from hits by too many MAC addresses:

```
Switch:1(config)#interface gigabitethernet 1/1
Switch:1(config-if)#mac-security limit-learning enable
Switch:1(config-if)#mac-security limit-learning max-addr 5000
```

Variable Definitions

Use the data in the following table to use the **mac-security limit-learning** command.

Variable	Value
<i>enable</i>	Limits the MAC learning for the port. After the number of addresses reaches the maximum, the port disables packet forwarding and drops packets. If you enable <i>limit-learning</i> , the FDB entry for each port is limited to the number you specify in <i>max-addr</i> .
<i>max-addr</i> <1-64000>	Specifies the maximum number of MAC addresses to learn. After the number of addresses reaches the maximum, the port disables packet forwarding and drops packets. The default is 1024.
<i>port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all][all]}</i>	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <p>Use this parameter to apply the change to multiple ports without changing CLI modes.</p>

Configure the Forwarding Database Timeout Globally

Use the following procedure to configure the aging time globally for the forwarding database.

Procedure

1. Enter Global Configuration mode:
enable

configure terminal
2. Enter the following command:
mac-address-table aging-time <10-1000000>

Variable Definitions

Use the data in the following table to use the **mac-address-table** command.

Variable	Value
<i>aging-time</i>	Specifies the timeout period for dynamically learned mac addresses on the vlan. The default value is 300.
<i><10-1000000></i>	Specifies the range for the aging time.

Add or Remove Ports in a VLAN

Procedure

1. Enter Interface Configuration mode:
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Add ports in a VLAN:
vlan members add <1-4059> {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]} [{portmember|static|notallowed}]
3. Remove ports in a VLAN:
vlan members remove <1-4059> {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]} [{portmember|static|notallowed}]

Example

The following example illustrates the use of optional parameters with both commands:

```
Switch:1>enable
Switch:1#configure terminal
```

```

Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1<config>#interface vlan 21
Switch(config-if)#vlan members add 21 1/2 static
Switch(config-if)#vlan members remove 21 1/3 notallowed

```

Variable Definitions

Use the data in the following table to use the **vlan members add** and **vlan members remove** commands.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
portmember	Configures the port type as port member.

Configure NLB support

Use Microsoft Network Load Balancing (NLB) to share the workload among multiple clustering servers. For information about software scaling capabilities, see *Fabric Engine Release Notes*.

Before You Begin

- For all modes, configure an IP address on the VLAN enabled with NLB.
- To switch between Unicast NLB and Multicast NLB, you must first disable the NLB support.

About This Task

Use the following procedure to configure NLB support on an IP interface to enable or disable NLB support.

The default value is NLB support disabled.

**Note**

SPBM supports Network Load Balancing (NLB) Unicast and Multicast modes.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Enable NLB support on an interface:

```
nlb-mode unicast
```

Or

```
nlb-mode multicast
```

To switch from one nlb-mode to another, you must first disable the NLB support, and then enter the new nlb-mode.

3. (Optional) Disable NLB support on an interface:

```
no nlb-mode
```

Example

Configure unicast mode for VLAN 2, and then switch to multicast mode.

```
Switch:1(config)#interface vlan 2
Switch:1(config-if)#nlb-mode unicast
Switch:1(config-if)#no nlb-mode
Switch:1(config-if)#nlb-mode multicast
```

Configure a Tagged Port to Discard Untagged Frames

Configure a tagged port to discard all untagged packets so that the frame is not classified into the default VLAN for the port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure a tagged port to discard untagged frames:

```
untagged-frames-discard [port {slot/port[/sub-port] [-slot/port[/sub-  
port]][,...]]
```

3. Discard a tagged frame on an untagged port:

```
tagged-frames-discard [port {slot/port[/sub-port] [-slot/port[/sub-  
port]][,...]] enable
```

4. Untag the default VLAN on a tagged port:

```
untag-port-default-vlan [port {slot/port[/sub-port] [-slot/port[/sub-  
port]][,...]] enable
```

Example

Configure a tagged port to discard untagged frames:

```
Switch:1(config-if)#untagged-frames-discard port 1/1
```

Discard a tagged frame on an untagged port:

```
Switch:1(config-if)#tagged-frames-discard port 1/1 enable
```

Untag the default VLAN on a tagged port:

```
Switch:1(config-if)#untag-port-default-vlan port 1/2 enable
```

Variable Definitions

Use the data in the following table to use optional parameters with the **untagged-frames-discard** command.

Variable	Value
<code>{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Configuring SLPP

Enable the Simple Loop Prevention Protocol (SLPP) globally and for a VLAN to detect a loop and automatically stop it. The VLAN configuration controls the boundary of SLPP-PDU transmission.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Enable SLPP:

```
slpp enable
```
3. Configure the transmission interval:

```
slpp tx-interval <500-5000>
```
4. Add a VLAN to the transmission list:

```
slpp vid <1-4059>
```

Example

Enable SLPP:

```
Switch:1(config)#slpp enable
```

Configure the transmission interval to 5000 milliseconds:

```
Switch:1(config)#slpp tx-interval 5000
```

Add a VLAN, with the VLAN ID 2, to the transmission list:

```
Switch:1(config)#slpp vid 1
```

Variable Definitions

Use the data in the following table to use the **slpp** command.

Variable	Value
<i>enable</i>	Enables or disables the SLPP operation. You must enable the SLPP operation to enable the SLPP packet transmit and receive process. If you disable the SLPP operation, the system sends no SLPP packets and discards received SLPP packets. The default is disabled.
<i>500–5000</i>	Configures the SLPP packet transmit interval, expressed in milliseconds in a range from 500–5000. The default value is 500.
<i><1–4059></i>	Adds a VLAN, by VLAN ID, to a SLPP transmission list.

Job aid

The following table provides the SLPP values.

Table 347: SLPP values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configure SLPP packet-rx on a Port

Enable SLPP by port to detect a loop and automatically stop it.



Important

To provide protection against broadcast and multicast storms, as a best practice, enable Rate Limiting for broadcast traffic and multicast traffic.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]}
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure SLPP on a port:

```
slpp port {slot/port[-slot/port][,...][slot/all][all]} packet-rx  
[packet-rx-threshold <1-500>]
```

Example

```
Switch(config-if)#slpp port 1/1 packet-rx-threshold 5
```

Variable Definitions

Use the data in the following table to use the **slpp port** command.

Variable	Value
<1-500>	<p>Specifies the SLPP reception threshold on the ports, expressed as an integer. The packet reception threshold specifies how many SLPP packets the port receives before it is administratively disabled. The default value is 1.</p> <p>Note: As a best practice, configure the rx-threshold above 50 slpp packets only on lightly loaded switches. If you configure the rx-threshold to a value greater than 50 on a heavily loaded switch and a loop occurs, the system can experience high CPU utilization.</p>
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...] [slot/all] [all]}	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Job aid

The following table provides the SLPP values.

Table 348: SLPP values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configure SLPP packet-tx on a VLAN

Enable SLPP by VLAN to detect a loop and automatically stop it. This configuration controls the boundary of SLPP-PDU transmission.



Important

To provide protection against broadcast and multicast storms, as a best practice, enable Rate Limiting for broadcast traffic and multicast traffic.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable SLPP:

```
slpp enable
```

3. Configure the transmission interval:

```
slpp tx-interval <500-5000>
```

4. Add a VLAN to the transmission list:

```
slpp vid <1-4059>
```

Example

Enable SLPP:

```
Switch:1(config)#interface vlan 2  
Switch(config-if)#slpp enable
```

Configure the transmission interval to 500 milliseconds:

```
Switch:1(config-if)#slpp tx-interval 500
```

Add a VLAN, with the VLAN ID of 2, to the transmission list:

```
Switch:1(config-if)#slpp vid 2
```

Variable Definitions

Use the data in the following table to use the **slpp** command.

Variable	Value
<i>enable</i>	Activates or disables the SLPP operation. You must enable the SLPP operation to enable the SLPP packet transmit and receive process. If you disable the SLPP operation, the system sends no SLPP packets and discards received SLPP packets. The default is disabled.
<i>500–5000</i>	Configures the SLPP packet transmit interval, expressed in milliseconds in a range from 500–5000. The default value is 500.
<i><1–4059></i>	Adds a VLAN, by VLAN ID, to a SLPP transmission list.

View SLPP Information

Use SLPP information to view loop information.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View SLPP information:
`show slpp`

Example

```
Switch:1#show slpp
=====
                        SLPP Info
=====
      operation : enabled
    tx-interval : 500
         vlan : 2
```

View SLPP Information for a Port

Show SLPP information for a port so that you can view the loop information for a port.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View SLPP information for a port:
`show slpp interface GigabitEthernet [{slot/port[-slot/port]}[,...]
[slot/all][all]]`

3. Clear SLPP packet RX counters:

```
clear slpp stats port [{slot/port[-slot/port]][,...][slot/all][all]}
```

Example

```
Switch:1>enable
Switch:1#show slpp interface GigabitEthernet 1/7
=====
Port Interface
=====
PORT      PKT-RX      PKT-RX      INCOMING      SLPP PDU
NUM       THRESHOLD   VLAN ID      ORIGINATOR
-----
1/7       enabled     5
PORT      PKT-RX      TIME LEFT
NUM       COUNT      TO CLEAR RX COUNT
-----
1/7       29         300
```

Variable Definitions

Use the data in the following table to use the **show slpp interface GigabitEthernet** command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,...][slot/all][all]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none">• a single slot and port (slot/port)• a range of slots and ports (slot/port-slot/port)• a series of slots and ports (slot/port,slot/port,slot/port)• all ports on the same slot (slot/all)• all ports on the switch (all) If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configure Automatic Port Recovery

About This Task

Instead of manually recovering a port that is shutdown by SLPP, you can configure an automatic recovery delay per port. The **auto-recover-delay** command specifies the minimum amount of time after which the port is brought back up. The maximum amount of time that can pass until the port is brought back up is two times the configured **auto-recover-delay** value.

By default, automatic recovery is disabled on a port. When enabled, the default delay is 30 seconds.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the automatic recovery delay, in seconds:
`auto-recover-delay <5-3600>`
3. Enable automatic recovery on a port:
`interface gigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}`
`auto-recover-port enable`

Configure Spoof Detection

Configure spoof detection to prevent IP spoofing.


For more information about this feature, see [Prevention of IP Spoofing within a VLAN](#) on page 4000.



Important

If you use SMLT, configure spoof detection on both SMLT aggregation devices to avoid connectivity issues.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
`enable`
`configure terminal`
`interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}`
-  **Note**
If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
2. Enable or disable spoof detection:
`spoof-detect [port {slot/port[/sub-port]} [-slot/port[/sub-port]] [,...]
[slot/all][all]}] [enable]`
`no spoof-detect [port {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}} [enable]`

3. Enable or disable auto-recovery on a port:

```
auto-recover-port [port {slot/port[/sub-port]} [-slot/port[/sub-port]]
[,...][slot/all][all]}

no auto-recover-port [port {slot/port[/sub-port]} [-slot/port[/sub-
port]][,...][slot/all][all]] [enable]
```

Example

Enable spoof detection:

```
Switch(config-if)#spoof-detect port 1/1 enable
```

Enable autorecovery on a port:

```
Switch(config-if)#auto-recover-port port 1/1 enable
```

View VLAN Information

View the VLAN information to display the basic configuration for all VLANs or a specified VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View VLAN information:

```
show vlan basic <1-4059>
```
3. View advanced parameters:

```
show vlan advance <1-4059>
```

Example

View VLAN information for VLAN 2:

```
Switch:1>show vlan basic 2
=====
Vlan Basic
=====
VLAN      INST
ID  NAME      TYPE      ID  PROTOCOLID  SUBNETADDR  SUBNETMASK  VRFID  ORIGIN
-----
2    VLAN-2    byPort    0   none        N/A         N/A         0      CONFIG
```

View VLAN information:

```
Switch:1>show vlan basic
=====
Vlan Basic
=====
VLAN      INST
ID  NAME      TYPE      ID  PROTOCOLID  SUBNETADDR  SUBNETMASK  VRFID  ORIGIN
-----
1    Default    byPort    0   none        N/A         N/A         0      CONFIG
2    abc        byPort    0   none        N/A         N/A         0      CONFIG
```

3	VLAN-VRRP	byPort	0	none	N/A	N/A	0	CONFIG
4	VLAN-6	byPort	0	none	N/A	N/A	1	CONFIG
5	VLAN-7	byPort	0	none	N/A	N/A	1	CONFIG
6	VLAN-8	byPort	0	none	N/A	N/A	1	CONFIG
19	VLAN-9	byPort	0	none	N/A	N/A	0	CONFIG
10	VLAN-10	byPort	0	none	N/A	N/A	0	CONFIG
11	VLAN-11	byPort	0	none	N/A	N/A	0	CONFIG
12	VLAN-12	byPort	0	none	N/A	N/A	0	CONFIG
13	VLAN-13	spbm-bvlan	62	none	N/A	N/A	0	CONFIG
14	VLAN-14	spbm-bvlan	62	none	N/A	N/A	0	CONFIG
15	VLAN-15	byPort	1	none	N/A	N/A	0	CONFIG
--More-- (q = quit)								

View advanced parameters. The following example displays a VLAN that is not associated with a Segmented VRF, as seen in the L3 ACCESS of GlobalRouter:

```
Switch:1>show vlan advance
=====
                        Vlan Advance
=====
VLAN      IF      AGING  MAC      USER      L3
ID  NAME      INDEX TIME  ADDRESS  DEFINEPID ACCESS  ENCAP  DSAP/SSAP
-----
2    Default  2050  0    00:24:7f:9f:6a:03  0x0000  GlobalRouter
```

Variable Definitions

Use the data in the following table to use optional parameters with the **show vlan basic** and **show vlan advance** commands.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

View Private VLAN Information

You can view the private VLAN information to display the primary and secondary VLANs and I-SIDs, and also view the private VLAN port types.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View private VLAN information:
show vlan private-vlan <1-4059>
3. View private VLAN port information:
show interfaces gigabitethernet private-vlan

Example

View VLAN information for private VLAN:

```
Switch:1(config)#show vlan private-vlan
=====
PRIVATE VLAN
=====
Primary      Primary      Secondary      Secondary
VLAN         ISID         VLAN           ISID
-----
3            75           5              75
10           22           15             22

All 2 out of 2 Total Num of Private Vlans displayed
```

View port information for private VLAN:

```
Switch:1(config)#show interfaces gigabitethernet private-vlan
=====
Port Private Vlans
=====
PORT      TAGGING  ORIGIN  PVLAN  PVLAN TYPE  VID  VID
NUM              TYPE              TYPE      TYPE
-----
2/2       enable   CONFIG  enable  promiscuous -     -

All 1 out of 1 Total Num displayed
```

View Brouter Port Information

View the brouter port information to display the brouter port VLAN information for all VLANs on the device or for the specified VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View brouter port information:
show vlan brouter-port

Example

View brouter port information:

```
Switch:1>show vlan brouter-port
Vlan Id      Port      VrfId
=====
2202         1/11      0
```

All 1 out of 1 Total Num of Vlan Brouter Port Entries displayed

View VLAN Port Member Status

View the VLAN port member status to display the port member status for all VLANs on the device or for the specified VLAN.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View VLAN port member status:

```
show vlan members [<1-4059>][null-vlan][port {slot/port[-slot/port]}
[,...][slot/all][all]]
```

Example

View VLAN port member status:

```
Switch:1>show vlan members port 1/2
=====
Vlan Port
=====
VLAN PORT      ACTIVE      STATIC      NOT_ALLOW
ID  MEMBER      MEMBER      MEMBER      MEMBER
-----
2    1/2,1/5-1/8,1/11,  1/2,1/5-1/8,1/11,
    1/14,1/26,1/38    1/14,1/26,1/38

3    1/2,1/5-1/8,1/14,  1/2,1/5-1/8,1/14,
    1/26,1/38          1/26,1/38

4    1/1-1/2,1/5-1/8,  1/1-1/2,1/5-1/8,
    1/13-1/14,1/25-    1/13-1/14,1/25-
    1/26,1/37-1/38    1/26,1/37-1/38

100  1/2,1/14,1/23-    1/2,1/14,1/23-
    1/24,1/26-1/28,    1/24,1/26-1/28,
    1/38                1/38

300  1/2,1/5-1/8,1/14,  1/2,1/5-1/8,1/14,
    1/26,1/38          1/26,1/38
```

Variable definitions

Use the data in the following table to use optional parameters with the **show vlan members** command.

Variable	Value
<code>null-vlan</code>	Displays port members of the NULL VLAN. This is a place holder VLAN for ports that are not members of any port-based VLAN. When a port is removed from all port-based VLANs, it is added to the NULL VLAN as a port member. The NULL VLAN is an internal construct and cannot be deleted.
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...] [slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p> <p>If you do not specify a port, the command shows information for all the ports.</p>
<code><1-4059></code>	<p>Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.</p> <p>Note: Entering a VLAN ID is optional. If you enter a VLAN ID the command shows information for the specified VLAN or port. Without the VLAN ID the command shows information for all the configured VLANs.</p>

View VLAN Source MAC Addresses

View the VLAN source MAC addresses to display the source MAC address for a source MAC-based VLAN on the device or for the specified VLAN.

Procedure

View VLAN source MAC addresses:

```
show vlan src-mac [<1-4059>]
```

Example

View VLAN source MAC addresses:

```
Switch(config)#show vlan src-mac
=====
                        Vlan Srcmac
=====
VLAN_ID    MAC_ADDRESS
-----
10         00:00:00:00:00:11

All 1 out of 1 Total Num of Vlan Srcmac Entries displayed
```

View VLAN Forwarding Database Information

Use this procedure to display the MAC addresses that are learned or statically configured for a VLAN. In order to learn you have to be connected to another switch or host and receive some traffic.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View VLAN forwarding database information:

```
show vlan mac-address-entry [<1-4059>]
```

Example

View VLAN forwarding database information:

```
Switch:1>show vlan mac-address-entry
=====
                        Vlan Fdb
=====
VLAN      MAC              SMLT
ID  STATUS  ADDRESS          INTERFACE  REMOTE TUNNEL
-----
1    learned  f8:15:47:e1:80:0c Port-1/2   false  -
2    learned  32:20:d3:81:00:77 Port-1/9   false  -
4    learned  b4:a9:5a:2b:78:31 Port-2/1   false  -

3 out of 3 entries in all fdb(s) displayed.
```

View where entries are learned. The TUNNEL column indicates where in the SPBM network an entry is learned.

```
Switch:1>show vlan mac-address-entry spbm-tunnel-as-mac
=====
                        Vlan Fdb
=====
VLAN      MAC              SMLT
ID  STATUS  ADDRESS          INTERFACE  REMOTE TUNNEL
-----
1    learned  f8:15:47:e1:80:0c Port-1/2   false  -
2    learned  32:20:d3:81:00:77 Port-1/9   false  -
```

```

4      learned      b4:a9:5a:2b:78:31  Port-2/1      false  -
3 out of 3 entries in all fdb(s) displayed.

```

Variable Definitions

Use the data in the following table to use optional parameters with the **show vlan mac-address-entry** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
mac <0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the MAC address.
port {slot/port[-slot/port][,...]}	Specifies the port or port list.
spbm-tunnel-as-mac	Displays where entries are learned. The TUNNEL column indicates where in the SPBM network an entry is learned.

View Manual Edit MAC Addresses

Use the procedure to view the list of manual edit MAC addresses and the associated ports configured as allow-mac for MAC security.

Procedure

View manual edit MAC addresses:

```
show vlan manual-edit-mac
```

Example

View manual edit MAC addresses:

```

Switch(config)#show vlan manual-edit-mac
=====
Manual Edit Mac
=====
MAC ADDRESS      PORTS
-----
00:00:00:00:00:55  1/3
00:00:00:00:00:66  1/3

All 2 out of 2 Total Num of Manual Edit Mac Entries displayed

```

View Port-Level MAC Security

View port-level MAC security to review the configuration.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View port-level MAC security for limit-learning:
show interfaces gigabitethernet limit-fdb-learning [{slot/port[-slot/port]}[,...]]

Example

View port-level MAC security for limit-learning:

```
Switch:1#show interfaces gigabitethernet limit-fdb-learning 1/4-1/5
=====
Port limit-fdb-learning
=====
```

PORT NUM	FDB PROTECT	MAXMAC COUNT	MINMAC COUNT	PORT DOWN	CURMAC COUNT	MAC LEARN
1/4	dis	1024	512	dis	0	true
1/5	ena	5000	3000	dis	0	true

View NLB-Mode Information

View Network Load Balancing-mode (NLB-mode) information.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View NLB port information:
show interfaces vlan nlb-mode [<1-4059>]

Example

View NLB-mode information.

```
Switch:1#show interfaces vlan nlb-mode
=====
Vlan Nlb
=====
```

VLAN_ID	NLB_ADMIN_MODE	NLB_OPER_MODE	PORT_LIST	MLT_GROUPS
2	unicast	disable		
22	multicast	multicast	1/19-1/21	2 3

Total Entries: 2

Enable DvR on a Layer 2 VSN (VLAN)

Before You Begin

- Ensure that the VLAN on which to enable DvR exists and is associated with an I-SID.

About This Task

On a Controller, DvR must be manually enabled at the Layer 2 VSN (VLAN) level.

Use this procedure to configure a gateway IPv4 address for a Layer 2 VSN (VLAN) subnet and then enable DvR on it. This address is pushed, along with other Layer 3 configuration information, from the Controllers to the Leaf nodes within the domain, so that the Leaf nodes can create a gateway IP service for each DvR-enabled Layer 2 VSN.



Note

Optionally, you can now use a single IP address in the subnet for every Controller by configuring the DvR VLAN IP to be the same as the DvR GW IP.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable

configure terminal

interface vlan <1-4059>
```

2. Configure a gateway IPv4 address for the VLAN.

```
dvr gw-ipv4 {A.B.C.D}
```



Important

Ensure that you configure the same gateway IPv4 address on all Controllers in the domain that belong to a Layer 2 VSN (VLAN).

3. Enable DvR.

```
dvr enable
```

By default, DvR is disabled.

Example

Enable DvR on the Global Router VLAN.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config-if)#vlan create 200 type port-mstprstp 0
Switch:1(config-if)#vlan i-sid 200 20200
Switch:1(config-if)#interface vlan 200
Switch:1(config-if)#dvr gw-ipv4 192.0.2.1
Switch:1(config-if)#dvr enable
Switch:1(config-if)#ip address 192.0.2.2 255.255.0.0
```

Enable DvR on a VLAN associated to a VRF:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config-if)#vlan create 400 type port-mstprstp 0
Switch:1(config-if)#vlan i-sid 400 20200
```

```
Switch:1(config-if)#interface vlan 400
Switch:1(config-if)#vrf vrf500
Switch:1(config-if)#dvr gw-ipv4 198.51.100.1
Switch:1(config-if)#dvr enable
Switch:1(config-if)#ip address 198.51.100.2 255.255.0.0
```

Variable definitions

Use the data in the following table to use the **dvr gw-ipv4** command.

Variable	Value
{A.B.C.D}	Specifies the gateway IPv4 address for the VLAN.

Configure a Single IP Address for All DvR Controllers on a VLAN Subnet

Before You Begin

- Ensure that the VLAN on which to enable DvR exists and is associated with an I-SID.

About This Task

On a Controller, DvR must be manually enabled at the Layer 2 VSN (VLAN) level.

Use this procedure to configure a single IP address to be used for both the DvR GW IP and the VLAN IP in the DvR VLAN. This IP address represents the DvR GW IP address used by all DvR Controllers for the same DvR VLAN. When you configure a DvR VLAN in this way, the Controllers no longer have a unique VLAN IP address.



Note

Optionally, you can use two different IP address for the DvR VLAN and the DvR GW IP.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface vlan <1-4059>
```

2. Configure a single IP address for both the DvR GW IP and the VLAN IP in the DvR VLAN.

```
ip address {<A.B.C.D/X>|<A.B.C.D> <A.B.C.D>} dvr-one-ip [state-  
disabled]
```

3. Enable DvR.

```
dvr enable
```

By default, DvR is disabled.

Example

Enable DvR.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config-if)#vlan create 200 type port-mstprstp 0
Switch:1(config-if)#vlan i-sid 200 20200
Switch:1(config-if)#interface vlan 200
Switch:1(config-if)#ip address 192.0.2.1 255.255.0.0 dvr-one-ip
Switch:1(config-if)#dvr enable
```

Variable Definitions

The following table defines parameters for the **ip address** command.

Variable	Value
<A.B.C.D/X> <A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask in the format A.B.C.D/X or A.B.C.D A.B.C.D.
<i>dvr-one-ip</i>	Specifies the IP address used for the DvR gateway IP address and the VLAN IP in the DvR VLAN.
<i>state-disabled</i>	Creates the IP interface in an administratively disabled state. This parameter is useful if you want to preconfigure an IP interface, for example, to prepare for DvR migration, and enable it later during a maintenance window. By default, the switch enables an IP interface after you configure the primary IP address.

VLAN Configuration using EDM

This section describes how to configure and manage Virtual Local Area Networks (VLAN) using Enterprise Device Manager (EDM).

Configure a VLAN on a Port

Configure a VLAN on a port.

Procedure

1. In the Device Physical View tab, select a port or multiple ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **VLAN** tab.
5. To perform tagging, select **PerformTagging**.
6. To discard tagged frames, select **DiscardTaggedFrames**.
7. To discard untagged frames, select **DiscardUntaggedFrames**.

8. To use the Untag Default VLAN feature, select **UntagDefaultVlan**.



Important

Enable tagging on the port before you configure UntagDefaultVlans.

9. Enter a default VLAN ID.
 10. To enable spoof detect, select **SpoofDetect**.
 11. In the Classification area, select the types of VLAN to enable.
 12. In the **Classification** area, select the Private VLAN port type.
 See [Create a Private VLAN](#) on page 4049 for more information.
 13. Select **Apply**.
 14. Select **Close**.

VLAN Field Descriptions

Use the data in the following table to use the **VLAN** tab.

Name	Description
PerformTagging	If checked, this port is a tagged (Trunk) Port. It can belong to multiple port-based VLANs and a VLAN tag is inserted in every frame it transmits. If it is not checked, the port is an untagged (Access) port. The default is disabled.
VlanIdList	Identifies which VLANs this port is assigned.
DiscardTaggedFrames	If selected, and the port is untagged (an access port), tagged frames received on the port are discarded by the forwarding process. If clear, tagged frames are processed normally. The default is disabled.
DiscardUntaggedFrames	If selected and the port is tagged (a trunk port), untagged frames received on the port are discarded by the forwarding process. If clear, untagged frames are processed normally. The default is disabled.
UntagDefaultVLAN	If selected, even if the port is tagged (a trunk port), frames forwarded to the default VLAN for the port are not tagged. The default is disabled.
UntaggedVlanIds	Identifies which VLANs this port is associated with as untagged.
DefaultVlanId	Specifies the VLAN ID assigned to untagged frames received on this trunk port that match no policy-based VLAN to which the port belongs.
SpoofDetect	Enables or disables spoof detection on the specified port.

Name	Description
PrivateVlanPortType	Specifies the port type for a Private VLAN. If not specified, the port type defaults to None. <ul style="list-style-type: none">• Isolated: An Isolated port can belong only to one private VLAN• Promiscuous: A Promiscuous port can belong to many private VLANs.• Trunk: A Trunk port can belong to many private VLANs, is tagged, and can also belong to non-private VLANs.
Origin	Specifies the origin of VLAN configuration on the port, either manually configured through CLI or EDM, or dynamically configured through Auto-sense.

Configure the VLAN Feature on an Extreme Integrated Application Hosting Port

About This Task



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

Perform this procedure to configure the VLAN feature on an Extreme Integrated Application Hosting (IAH) port.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **VLAN** tab.
4. Select **PerformTagging** to enable tagging.
5. Select **DiscardTaggedFrames** to discard tagged frames.
6. Select **DiscardUntaggedFrames** to discard untagged frames.
7. Select **UntagDefaultVlan** to enable the Untag Default VLAN feature.



Important

You must enable tagging on the IAH port before you enable the Untagging Default VLANs feature.

8. In **DefaultVlanId**, enter a default VLAN ID.
9. Select **SpoofDetect** to enable spoof detection.
10. In **PrivateVlanPortType**, select a type.
11. Select **Apply**.

VLAN Field Descriptions

Use data in the following table to use the **VLAN** tab.

Name	Description
PerformTagging	Enables tagging (trunking) on the Extreme Integrated Application Hosting (IAH) port. It can belong to multiple port-based VLANs and a VLAN tag is inserted in every frame it transmits. If it is not checked, the IAH port is an untagged (Access) IAH port. The default value is different for each IAH port.
VlanIdList	Shows the VLAN ID to which the IAH port is assigned to.
DiscardTaggedFrames	Enables the functionality to discard tagged frames received on the IAH port by the forwarding process, for untagged (access) IAH port. The default is disabled.
DiscardUntaggedFrames	Enables the functionality to discard untagged frames received on the IAH port by the forwarding process, for tagged (trunk) IAH port. The default is disabled.
UntagDefaultVLAN	Enables the functionality to untag the frames forwarded to the default VLAN, even if the IAH port is tagged (trunk port). The default is disabled.
UntaggedVlanIds	Shows the untagged VLAN IDs associated with the IAH port.
DefaultVlanId	Specifies the VLAN ID assigned to untagged frames received on the IAH port that match no policy-based VLAN to which the EIAH port belongs.
SpoofDetect	Enables spoof detection on the specified IAH port.
PrivateVlanPortType	Specifies the following: <ul style="list-style-type: none"> • trunk: tagged IAH port. • isolated: Only private VLANs are permitted on isolated IAH ports. • promiscuous: Only private VLANs are permitted on promiscuous IAH ports. • none: default value.

View existing VLANs

View the configuration and current status of all defined VLANs.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.

2. Select **VLANs**.
3. View the configured VLANs in the **Basic** tab.
4. View the configured private VLANs in the **Private VLAN** tab.

Create a Port-Based VLAN

Create a port-based VLAN. To create a different type of VLAN, see one of the following procedures:

- [Create a Protocol-Based VLAN](#) on page 4054
- [Create an SPBM B-VLAN](#) on page 4055

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. In the **Basic** tab, select **Insert**.
4. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
5. In the **Name** box, type the VLAN name, or use the name provided.
6. In the **Color Identifier** box, select the down arrow and choose a color from the list, or use the color provided.
7. In the **MstpInstance** box, select the down arrow and choose an msti instance from the list.
8. In the **Type** box, select **byPort**.
9. In the **PortMembers** box, select the ellipsis (...).
10. Select the ports to add as member ports.
The ports that are selected are recessed. Nonselected ports are not recessed. You cannot select a Port number that is dimmed as a VLAN port members.
11. Select **OK**.
12. Select **Insert**.

Basic Field Descriptions

Use the data in the following table to use the **Basic** tab.

Name	Description
Id	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.

Name	Description
Color Identifier	Specifies a proprietary color scheme to associate a color with the VLAN. Color does not affect how frames are forwarded.
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-bvlan • private
MstpInstance	Identifies the MSTP instance.
VrfId	Indicates the VRF ID to which the VLAN belongs.
VrfName	Indicates the name of the VRF to which the VLAN belongs.
PortMembers	Specifies the slot/port of each VLAN member. The system displays the sub-port only for channelized ports.
ActiveMembers	Specifies the slot/port of each VLAN member. The system displays the sub-port only for channelized ports.
StaticMembers	Specifies the slot/port of each static member of a policy-based VLAN. The system displays the sub-port only for channelized ports.
NotAllowToJoin	Specifies the slot/ports that are never allowed to become a member of the policy-based VLAN. The system displays the sub-port only for channelized ports.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.

Name	Description
AgingTime	Specifies the timeout period, in seconds, to age out dynamic members of this VLAN. This parameter only applies to policy-based VLANs. The default is 600.
Clear Secondary IPs	Removes all secondary IP addresses on a VLAN simultaneously. To remove specific secondary IP addresses from a VLAN, use the Secondary IP Addresses tab (Configuration > VLAN > VLANs > Basic > IP). For more information, see Secondary IP Addresses Field Descriptions on page 4053.



Note

If you or another user changes the name of an existing VLAN using the VLAN **Basic** tab (or using CLI), the system does not display the new name initially in EDM. To display the updated name, perform one of the following actions:

- Refresh your browser to reload EDM.
- Log out of EDM and log in again to restart EDM.
- Select **Refresh** in the VLAN **Basic** tab toolbar. If the system displays the old VLAN name in other tabs, select **Refresh** on those tabs as well.

Create a Private VLAN

Before You Begin

- To create a private VLAN, you must configure the VLAN type to private and configure the private VLAN port type.
- The ports you add to a private VLAN must have a port type of isolated, promiscuous, or trunk.

About This Task

For additional information about the tabs referenced in this procedure, see the following sections:

- **VLAN > VLANs > Basic** tab: [Basic Field Descriptions](#) on page 4047
- **Edit > Port > General > VLAN** tab: [VLAN Field Descriptions](#) on page 4044

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select **Insert**.
5. For **Id**, type an unused VLAN ID, or use the ID provided.

6. (Optional) For **Name**, type the VLAN name, or use the name provided.
7. (Optional) For **Color Identifier**, select a color from the list, or use the color provided.
8. (Optional) For **MstpInstance**, select an msti instance from the list.
9. For **Type**, select **private**.
10. For **PortMembers**, select the ellipsis (...).
11. Select the ports to add as member ports.
The ports that are selected are recessed, while the non-selected ports are not recessed. Port numbers that are dimmed cannot be selected as VLAN port members.
12. Select **OK**.
13. For **Secondary Vlan**, type an unused VLAN ID.
14. Select **Insert**.
15. In the Device Physical View, select the Private VLAN port members.
16. In the navigation pane, expand **Configuration > Edit > Port**.
17. Select **General**.
18. Select the **VLAN** tab.
19. For **PrivateVlanPortType**, select the port type.
20. Select **Apply**.

View Private VLAN Information

You can view the private VLAN information for the primary and secondary VLANs, I-SIDs, and private VLAN port types.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select **Private VLAN**.

Private VLAN Field Descriptions

Use the data in the following table to use the **Private VLAN** tab.

Name	Description
Primary Vlan	Shows the VLAN ID for the primary VLAN.
Secondary Vlan	Shows the VLAN ID for the secondary VLAN.
Primary / Secondary I-sid	Shows the I-SID for the VLAN.

Configure a Primary IP Address for a VLAN

Assign an IP address to a VLAN to enable routing on the VLAN.

About This Task

The ability to create an IP interface in a disabled state in one step is available only using CLI.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. On the **Basic** tab, select the VLAN to configure.
4. Select **IP**.
5. On the **Primary IP Address** tab, select **Insert**.
6. Configure the required parameters.
7. Select **Insert**.

Primary IP Address Field Descriptions

Use the data in the following table to use the **Primary IP Address** tab.

Name	Description
Interface	Shows the interface to which this entry applies.
Ip Address	Specifies the IP address to associate with the VLAN.
Net Mask	Specifies the subnet mask associated with the IP address of this entry. The value of the mask is an IP address with all the network bits configured to 1 and all the hosts bits configured to 0.
Name	Specifies a name for the interface.
BcastAddrFormat	Shows the IP broadcast address format on this interface.
ReasmMaxSize	Shows the size of the largest IP datagram which this entity can reassemble from incoming IP fragmented datagrams received on this interface.
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
BrouterPort	Indicates whether this entry corresponds to a brouter port, as oppose to a routable VLAN.

Name	Description
MacOffset	Specifies the MAC offset value. Routable VLANs are assigned MAC addresses arbitrarily or by offset. Their MAC addresses are: <ul style="list-style-type: none"> • 24 bits: Vendor ID • 12 bits: Chassis ID • 12 bits: 0xA00-0xFFFF If you enter the MAC offset, the lowest 12 bits are 0xA00 plus the offset. If not, they are arbitrary.
IfType	Identifies the type of interface.
VrfId	Associates the VLAN with a VRF. VRF ID 0 is reserved for the administrative VRF.

Modify the Administrative State of an IP Interface on a VLAN

By default, the switch enables an IP interface after you configure the primary IP address. You can administratively disable, and then reenables, an IPv4 interface after you create it. For example, if you want to change the routing protocol configuration on the interface.

About This Task

This procedure applies to the IP Layer and functionality for the VLAN.

The ability to create an IP interface in a disabled state in one step is available only using CLI.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. On the **Basic** tab, select the VLAN you need to modify.
4. Select **IP**.
5. Select the **IP Interface** tab.
6. Change the **AdminStatus** value to change the administrative state.
7. Select **Apply**.

IP Interface Field Descriptions

Use the data in the following table to help use the **IP Interface** tab.

Name	Description
AdminStatus	Modifies the administrative state of the IP Interface. By default, the switch enables an IP interface after you configure the primary IP address.
OperState	Displays the operational state of the IP Interface.

Configure a Secondary IP Interface for a VLAN

Assign a secondary IP address to a VLAN to increase the number of connected hosts the VLAN can support.

Before You Begin

You must configure a primary IP address before you can add a Secondary IP Interface.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. On the **Basic** tab, select the VLAN to configure.
4. Select **IP**.
5. On the **Secondary IP Addresses** tab, select **Insert**.
6. Type the IP address.
7. Type the subnet mask.
8. Select **Insert**.

Secondary IP Addresses Field Descriptions

Use the data in the following table to use the **Secondary IP Addresses** tab.

Name	Description
Interface	Identifies the IP interface.
Ip Address	Specifies the secondary IP address.
Net Mask	Specifies the subnet mask for the address.
Vlan ID	Identifies the VLAN ID.
Vrf ID	Identifies the associated VRF ID, if applicable.

Change VLAN Port Membership

Modify VLAN port members to control access to the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
 2. Select **VLANs**.
 3. Select the **PortMembers** number for the VLAN for which you want to modify port membership.
 4. Select the port members you want to add or remove.
 5. Select **Ok**.
 6. Select **Apply**.
- The VLAN port membership is changed.

Create a Protocol-Based VLAN

Use a protocol-based VLAN so that the VLAN only carries certain traffic types.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. In the **Basic** tab, select **Insert**.
4. In the **Id** box, type the unique VLAN ID or use the ID provided.
5. In the **Name** box, type the VLAN name or use the name provided.
6. In the **Color Identifier** box, select the color or use the color provided.
This color is used to visually distinguish the VLANs in a network.
7. In the **MstpInstance** box, select the down arrow and choose an MSTI instance from the list.
8. In the **Type** box, select **byProtocolId**.
This activates additional fields needed to configure protocol-based VLANs.
9. To specify the VLAN port membership, select the ellipsis (...) for one of the following fields:

Port Members

OR

StaticMembers

OR

NotAllowToJoin

10. Select each port button to choose the desired membership color.
Yellow: Potential members—dynamic (potential members are treated as always members)

OR

Green: Always members—static

OR

Red: Never members—not allowed to join



Important

In a protocol-based VLAN, a potential member becomes an active member of the VLAN after a frame of the specified protocol is received.

11. Select **Insert**.

Create an SPBM B-VLAN

Create a Shortest Path Bridging MAC (SPBM) Backbone VLAN (B-VLAN). Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network. This VLAN is used for both control plane traffic and dataplane traffic.



Note

As a best practice, configure two B-VLANs in an SPBM dual-homing environment.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. In the **Basic** tab, Select **Insert**.
4. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
5. In the **Name** box, type the VLAN name, or use the name provided.
6. In the **Color Identifier** box, select the down arrow and choose a color from the list, or use the color provided.
7. In the **Type** box, select **spbm-bvlan**.
8. Select **Insert**.
9. Collapse the **VLANs** tab.

The VLAN is added to the **Basic** tab.

Configure Advanced VLAN Features

Use advanced VLAN features to configure the VLAN name, aging time, VLAN operation action, and QoS level. The VLAN Operation Action parameter can be useful for troubleshooting.

You can also configure a DvR Gateway IPv4 address on a VLAN, and enable DvR on it.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. Configure the parameters as required by double-clicking fields to make changes.
You cannot make changes to fields that are dim.
5. Select **Apply**.

Advanced Field Descriptions

Use the data in the following table to use the **Advanced** tab.

Name	Description
Id	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
Name	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-bvlan • private
Isid	Specifies the Service Instance Identifier (I-SID) number assigned to a customer VLAN (C-VLAN). The default value is 0, which indicates that no I-SID is assigned. <p>Note: When you add or remove an I-SID on a VLAN IP interface for 5520 Series or 5720 Series, the switch disables and enables routing on the VLAN interface.</p>
Isid Name	Specifies the name of the I-SID.
Origin	Specifies the origin of the VLAN configuration, either manually configured through CLI or SNMP, or dynamically configured through RADIUS VSA or Zero Touch Fabric Configuration.
ProtocolId	Specifies the network protocol for protocol-based VLANs. If the VLAN type is not protocol-based, the switch displays a value of None.
AgingTime	Specifies the timeout period for dynamic VLAN membership. A potential VLAN port is made ACTIVE after it receives a packet that matches the VLAN; if no such packet is received for AgingTime seconds, the port is no longer active. The default is 600.

Name	Description
MacAddress	Specifies the MAC address assigned to the virtual router interface for this VLAN, which is relevant only after the VLAN is configured for routing. This MAC address is used as the Source MAC in routed frames and ARP replies.
Vlan Operation Action	<p>Performs an operation on the VLAN. The values are:</p> <ul style="list-style-type: none"> • none • flushMacFdb: Configures action to flushMacFdb. This action removes the learned MAC addresses from the forwarding database for the selected VLAN. • flushArp: Configures action to flushArp. This action removes the ARP entries from the address table for the selected VLAN. • flushIp: Configures action to flushIp. This action removes the learned IP addresses from the forwarding table for the selected VLAN. • flushDynMemb: Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN and removes MAC addresses learned on those ports. • all: Configures action to all. This action performs all the supported actions; it does not perform the Snoop-related actions. <p>The default is none.</p>
Result	Specifies the result code after you perform an action.
NlbMode	Enables or disables Microsoft Network Load Balancing (NLB) operations on the VLAN. The default is disabled.
SpbMulticast	Enables or disables IP Multicast over Fabric Connect. The default is disabled.
SpbPimGatewayMulticast	Enables or disables SPB-PIM Gateway Multicast on a VLAN. The default is disabled.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
Ipv6FhsSnoopDhcpEnable	Enables or disables IPv6 dhcp snooping on a VLAN. The default is disabled.

Name	Description
Ipv6FhsNDInspectionEnable	Enables or disables neighbor discovery (ND) inspection on a VLAN. The default is disabled.
DvrEnable Note: Exception: not supported on 4220 Series.	Enables or disables Distributed Virtual Routing (DvR) on a VLAN that is configured on the DvR Controller. The default is disabled. Note: You must enable DvR on every VLAN that is configured on a DvR Controller.
DvrGwIpv4Addr Note: Exception: not supported on 4220 Series.	Specifies the DvR gateway IPv4 address for a VLAN. Important: Ensure that you configure the same gateway IPv4 address on all Controllers in the DvR domain that belong to a VLAN.
DvrGwIpv4Onelp Note: Exception: not supported on 4220 Series.	Enables or disables the DvR One IP for a VLAN. The default is disabled (false).
DvrVrrpElection Note: Exception: not supported on 4220 Series, 5320 Series, or 5420 Series.	Identifies the DvR VLAN, and associated I-SID, to become the primary Virtual Router Redundancy Protocol (VRRP) router through the election process. You can identify only one VLAN for election per DvR Controller. By default, no VLAN is configured as the DvR VRRP election VLAN.
MvpnIsidValue	Configures the MVPN I-SID value for the VLAN.
MvpnIsidStatus	Identifies the MVPN I-SID status. <ul style="list-style-type: none"> static—VRF was created by configuring an MVPN I-SID on a VLAN through CLI and EDM dynamic—VRF was created by configuring an MVPN I-SID on a VLAN through RADIUS notConfigured—MVPN I-SID is not configured on a VLAN.
MvpnIsidOffset	Allows you to add an offset parameter to the MVPN I-SID. This is an optional value and the range is 0-511.

Configure NLB Support

Use Microsoft Network Load Balancing (NLB) to share the workload among multiple clustering servers. For more information about software scaling capabilities, see *Fabric Engine Release Notes*.

Before You Begin

Ensure that the VLAN exists and has an associated IP address.

About This Task

Use the following procedure to configure NLB support on an IP interface to enable or disable NLB support. The default value is NLB support disabled.



Note

- SPBM supports Network Load Balancing (NLB) Unicast and Multicast modes.
- Static multicast ARP entries are not supported for NLB Unicast or NLB Multicast.
- Multicast MAC flooding is not supported for NLB.
- ARP entries for NLB server IP addresses do not age out when there is still client traffic coming to the NLB servers, even after the NLB servers are no longer reachable.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. In the row for the VLAN, double-click the value in the **NlbMode** column.
5. Select the appropriate value.
6. Select **Apply**.

Configure a Port to Accept Tagged or Untagged Frames

Configure a port to accept tagged or untagged frames.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **VLAN** tab.

5. To configure tagging on the port, select the **PerformTagging** check box.

This setting applies to all VLANs associated with the port.

**Note**

If the check box is selected, tagging is enabled. All frames sent from this port are tagged.

If the check box is cleared, tagging is disabled. The port does not send tagged frames. The switch removes the tag before sending the frame out of the port.

6. To discard tagged frames on a port for which tagging is disabled, select **DiscardTaggedFrames**.
7. To discard untagged frames on a port for which tagging is enabled, select **DiscardUntaggedFrames**.
8. To designate a default VLAN to associate with a packet that does not match a policy-based VLAN, enter a VLAN ID in the **DefaultVlanId** box or use the default VLAN 1.
9. Select **Apply**.
10. Select **Close**.

Configure Untagging Default VLAN on a Tagged Port

Configure an untagged default VLAN on a tagged port to separate untagged packets originating from a PC from the tagged packets originating from an IP phone.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **VLAN** tab.
5. Select **UntagDefaultVlan**.
6. In the **DefaultVlanId**, enter a default VLAN ID.
7. Select **Apply**.
8. Select **Close**.

Configure SLPP Globally

About This Task

Enable the Simple Loop Prevention Protocol (SLPP) to detect a loop and automatically stop it.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SLPP**.
3. Select the **Global** tab.
4. Select **GlobalEnable**.

5. In **TransmissionInterval**, type a value for the time interval for loop detection.
6. Select **Apply**.

Global Field Descriptions

Use the data in the following table to use the **Global** tab.

Name	Description
GlobalEnable	Activates or disables SLPP globally. The default is disabled.
TransmissionInterval	Configures the interval for which loop detection occurs. The interval is expressed in milliseconds in a range from 500–5000. The default value is 500.

Job aid

The following table provides the SLPP values.

Table 349: SLPP values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configure the SLPP by VLAN

Before You Begin

- Enable SLPP globally before you configure it on a VLAN.

About This Task

Activate SLPP on a VLAN to enable forwarding of the SLPP packet over the VLAN. This configuration controls the boundary of SLPP-PDU transmission.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SLPP**.
3. Select the **VLANS** tab.
4. Select **Insert**.

5. Select the **VlanId** ellipses (...).
6. Select the desired VLAN ID.
7. Select **Ok**.
8. Select **SlppEnable**.
9. Select **Insert**.

VLANS Field Descriptions

Use the data in the following table to use the **VLANS** tab.

Name	Description
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
SlppEnable	Activates SLPP on the selected VLAN. The SLPP packet transmission and reception process is active only if you enable the SLPP operation. If you disable the SLPP operation, the following occurs: <ul style="list-style-type: none">• the system sends no SLPP packets• the system discards received SLPP packets The default is enabled.

Configure the SLPP by Port

Before You Begin

- Enable SLPP globally before you configure it on a port.

About This Task

Use SLPP on a port to avoid traffic loops on the port.



Note

To provide protection against broadcast and multicast storms, as a best practice, enable Rate Limiting for broadcast traffic and multicast traffic.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **SLPP**.
3. Select the **Ports** tab.

4. Double-click **PktRxThreshold** for the desired port to edit the threshold value for packet reception.
5. Double-click **SlppEnable** for the desired port.
6. Select **true** to enable SLPP.
7. Select **Apply**.

Ports Field Descriptions

Use the data in the following table to use the **Ports** tab.

Name	Description
IfIndex	Specifies the interface index number for a port.
PktRxThreshold	Specifies the threshold for packet reception. Configure the SLPP packet receive threshold to a value (1- 500) that represents the number of received SLPP-PDUs to shut down the port. This variable is a port-level parameter, therefore if the port is tagged, SLPP-PDUs from the various VLANs increment this single threshold counter. The default is 1.
SlppEnable	Activates SLPP on the selected interface. The default is disabled.
IncomingVlanId	Shows the VLAN ID of the classified packet on a port disabled by SLPP.
SrcNodeType	Specifies the source node type of the received SLPP packet.
PktRxCount	Shows the total number of SLPP packets the port received.
TimeToClrPktRxCount	Specifies the timer to clear the SLPP receive counter. After you enable SLPP and the port receives SLPP PDUs, the timer starts. After the timer exceeds the configured value, the system resets the count to zero. The default is 21,600 seconds.
RemainingTimeToClrPktRxCount	Shows the time remaining before the SLPP receive counter is reset to zero.

Job aid

The following table provides the SLPP values.

Table 350: SLPP values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No

Table 350: SLPP values (continued)

Enable SLPP	Setting
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configure the Global Forwarding Database Timeout

Configure the forwarding database global timeout to age out dynamically learned forwarding information.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANS**.
3. Select the **FdbAging** tab.
4. Type an interval, in seconds, for aging out dynamically learned forwarding information.
5. Select **Apply**.

FDB Aging *Field Descriptions*

Use the data in the following table to use the **FDB Aging** tab.

Name	Description
FdbAging	Specifies the timeout period for dynamically learned mac addresses on the vlan. The default value is 300.

View VLAN Forwarding Database Information

About This Task

Perform this procedure to view forwarding database entries for all VLANs on the device.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANS**.
3. Select the **Forwarding** tab.

Forwarding Field Descriptions

Use the data in the following table to use the **Forwarding** tab.

Name	Description
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
Address	Specifies a unicast MAC address for which the VLAN has forwarding or filtering information.
Status	Specifies the status of the VLAN.
Port	Specifies either a value of zero (0) or the port number of the port on which a frame having the specified MAC address was seen. A value of cpp indicates a self-assigned MAC address.
BMac	Shows the backbone MAC address if the entry is learned from a Shortest Path Bridging MAC (SPBM) network.
Cvid	Specifies the customer VID.

View the Forwarding Database for a Specific VLAN

Use the forwarding database for VLANs to determine how the system forwards a received frame.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**
2. Select **VLANs**.
3. Select a VLAN.
4. Select **Bridge**.
5. Select the **Forwarding** tab and the VLAN forwarding database information is displayed.

Forwarding Field Descriptions

Use the data in the following table to use the **Forwarding** tab.

Name	Description
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
Address	Specifies a unicast MAC address for which the bridge has forwarding or filtering information.
Status	Specifies the status of this entry.
Port	Specifies either a value of zero (0) or the port number of the port on which a frame having the specified MAC address was seen. A value of cpp indicates a self-assigned MAC address.
Cvid	Identifies the customer VID for this interface.

Clear Learned MAC Addresses by VLAN

About This Task

Use the clear learned MAC addresses feature to flush the bridge forwarding database.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Advanced** tab.
4. Double-click **VLAN Operation Action**.
5. Select **FlushMacFdb** from the list.
6. Select **Apply**.

Clear Learned MAC Addresses for all VLANs by Port

About This Task

Use the following procedure to clear all the forwarding database (FDB) for VLANs associated with this port.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Interface** tab.
5. For **Action**, select **FlushMacFdb**.
6. Select **Apply**.

All learned MAC addresses are cleared from the forwarding database (FDB) for VLANs associated with this port.

7. Select **Close**.

Configure Static Forwarding

Configure static forwarding to specify the group of ports that are allowed to forward frames.



Important

Entries are valid for unicast and for group/broadcast addresses.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **Bridge**.
6. Select the **Static** tab.
7. Select **Insert**.
8. In **MacAddress**, type a forwarding destination MAC address.
9. In **Port**, select the ellipsis (...).
10. Select the port on which the frame is received.
11. Select **Ok**.
12. Select **Insert**.

Static Field Descriptions

Use the data in the following table to use the **Static** tab.

Name	Description
MacAddress	Specifies the destination MAC address in a frame to which the forwarding information for this entry applies. This object can take the value of a unicast address.
Port	Specifies the port number of the port on which the frame is received.

Name	Description
VlanId	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
Status	Specifies the status of the VLAN.

Configure Limit Learning

Limit MAC address learning to limit the number of forwarding database (FDB) entries learned on a particular port to a user-specified value. After the number of learned forwarding database entries reaches the maximum limit, MAC learning stops at that port.

Procedure

1. On the Device Physical View tab, select one or more ports.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.
4. Select the **Limit-Learning** tab.
5. Configure the parameters as required.
6. Select **Apply**.

Limit Learning Field Descriptions

Use the data in the following table to use the **Limit-Learning** tab.

Name	Description
PortNum	Shows the slot and port number to configure.
MaxMacCount	Configures the number of entries in the MAC table for the port that causes learning to stop. The default is 1024.
CurrentMacCount	Shows the number of entries currently in the MAC table for the port.
Enable	Enables or disables limit learning for the port. The default is disable.
MacLearning	Shows if MAC learning is enabled or disabled for the port.



VRF Lite

[VRF Lite Fundamentals](#) on page 4073

[VRF Lite Configuration using CLI](#) on page 4081

[VRF Lite Configuration Using EDM](#) on page 4092

Table 351: Virtual Routing and Forwarding product support

Feature	Product	Release introduced
Virtual Routing and Forwarding	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 351: Virtual Routing and Forwarding product support (continued)

Feature	Product	Release introduced
Inter-VRF Routing	4220 Series	Fabric Engine 9.2 The 4220 Series switches support only a single active VRF with IP configuration.
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	Fabric Engine 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 351: Virtual Routing and Forwarding product support (continued)

Feature	Product	Release introduced
Increased VRF and Layer 3 VSN scaling	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE Note: The boot config flag vrf-scaling command does not apply to 5320-16P-2MXT-2X or 5320-24T-4X-XT.
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 351: Virtual Routing and Forwarding product support (continued)

Feature	Product	Release introduced
IBGP over user-created VRFs	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.1
IPv6 Virtual Routing and Forwarding	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6 For the 5320 Series, only the following models support more than one VRF with IP configuration: <ul style="list-style-type: none"> • 5320-16P-2MXT-2X • 5320-24T-4X-XT • 5320-24T-24S-4XE-XT • 5320-48P-8XE • 5320-48T-8XE
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

VRF Lite provides secure customer data isolation.

VRF Lite Fundamentals

The switch supports what is termed as VRF Lite. Lite conveys the fact that the switch does not use Multiprotocol Label Switching (MPLS) for VRF; VRF Lite is a device virtualization feature, not a network-wide virtualization feature.

Use VRF Lite to offer networking capabilities and traffic isolation to customers that operate over the same node (router). Each virtual router emulates the behavior of a dedicated hardware router; the network treats each virtual router as a separate physical router. In effect, you can perform the functions of many routers using a single platform that runs VRF Lite. The result is a substantial reduction in the cost associated with providing routing and traffic isolation for multiple clients.

With multicast virtualization for IPv4, the switch can function as multiple virtual multicast routers.

The following figure shows one platform acting as multiple virtual routers, each serving a different customer network.

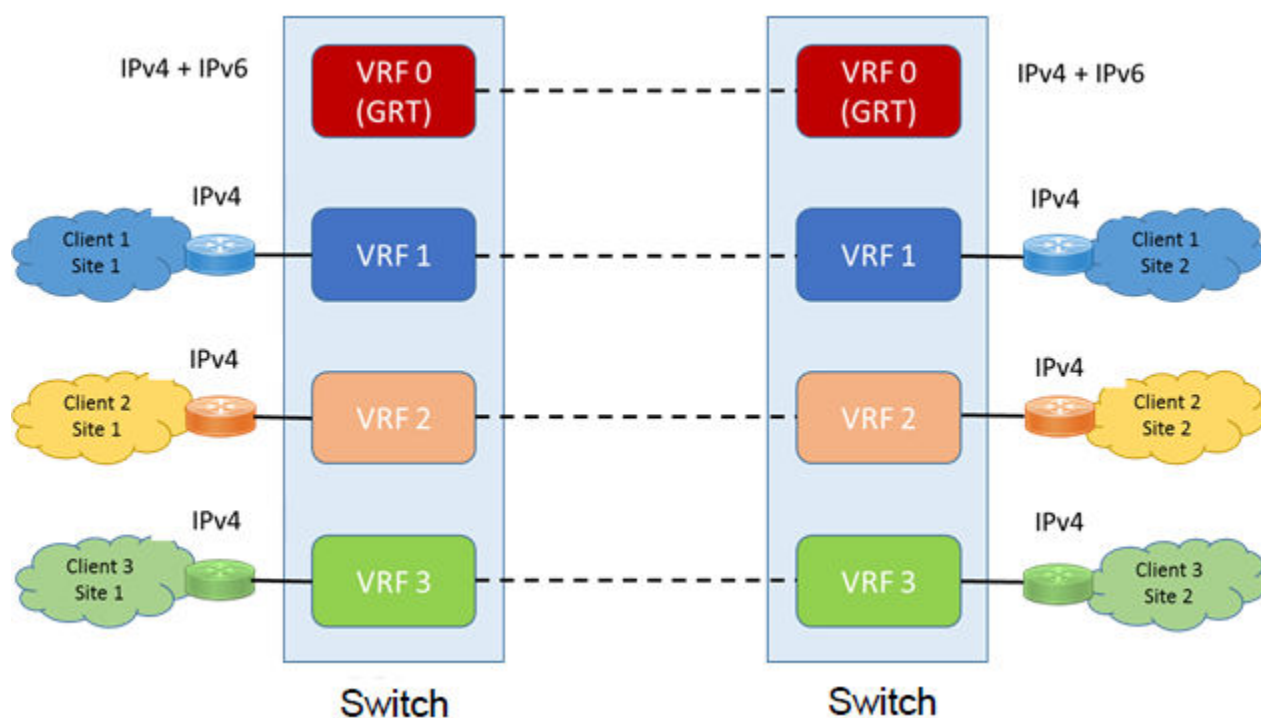


Figure 251: Multiple virtual routers in one system

A switch can support many virtual routers. Each virtual router instance is called a VRF instance. A VRF represents a single instance of a virtual router. Each instance maintains its own routing table. The term Multiple Virtual Router (MVR) is sometimes used to represent a router that contains many VRF instances.

The IPv6 Virtualization functionality adds IPv6 support on VRFs and Layer 3 VSNs. Each VRF instance has its own IPv6 interfaces, IPv6 address space, IPv6 routing table,

and IPv6 global parameters. For more information on Layer 3 VSN, see [Layer 3 VSN Configuration](#) on page 1654.

The Global Router, VRF 0, is the first instance of the router. When the system starts, it creates VRF 0 by default. VRF 0 provides all non-virtual and traditional routing services. You cannot delete this instance. You can create and configure other VRF instances, if required.

VRF 0 is the only VRF that you can log into through CLI. CLI requires you to specify the VRF when you enter commands.

You can associate one VRF instance with many IPv4 or IPv6 interfaces. These interfaces are unique for each VRF instance. An interface is an entity with an IPv4 or IPv6 address that has the following characteristics:

- A unique association with a VLAN.
- A unique association with a brouter, if not associated with a VLAN
- A unique association with a circuit

A VLAN can only be associated with a single VRF instance.



Note

- You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IPv4 or IPv6 address. You must first associate the port and VRF instance and then you can configure the IPv4 or IPv6 address.
- Use the **boot config flag vrf-scaling** command, on supported switches, to increase the total number of supported VRFs. For more information on route scaling, see *Fabric Engine Release Notes*.

VRF Lite Capability and Functionality

VRF Lite supports virtualization of IPv4 and IPv6 protocols and features. Use the VRF Router Configuration Mode in the CLI or the **Configuration > VRF** navigation pane in Enterprise Device Manager (EDM) to view the specific capabilities supported on your switch.

The switch uses VRF Lite to perform the following actions:

- Partition traffic and data and represent an independent router in the network
- Provide virtual routers that are transparent to end-users
- Support addresses that are not restricted to the assigned address space provided by host Internet Service Providers (ISP)
- Support overlapping IP address spaces in separate VRF instances



Note

If you enable multicast route redistribution between two VRFs, the switch does not support IP addresses that overlap within the two VRFs. The device does not generate an error if addresses overlap. You must avoid this situation.

VRF Lite interoperates with RFC 4364, Layer 3 VPNs. Split MultiLink Trunking (SMLT) and Routed SMLT (RSMLT) are also supported for VRF instances.

Although customer data separation into Layer 3 virtual routing domains is usually a requirement, sometimes customers must access a common network infrastructure. For example, they want to access the Internet, data storage, Voice over IP (VoIP)-public switched telephone network (PSTN), or call signaling services. To interconnect VRF instances, you can use an external firewall that supports virtualization, or use inter-VRF forwarding for specific services. With the inter-VRF solution, you can use routing policies and static routes to inject IP subnets from one VRF instance to another, and you can use filters to restrict access to certain protocols. The following figure depicts inter-VRF forwarding by the switch.




Figure 252: Inter-VRF forwarding

For more information about the latest VRF Lite scalability, see *Fabric Engine Release Notes*.

For configuration information about multicast virtualization, see [IP Multicast](#) on page 1696.

VRF Lite and inter-VRF route redistribution



Note

Some switches do not support Inter-VRF Routing because they support only a single active VRF. For more information about these models, see [VRF Lite Configuration Rules](#) on page 4078.

The switch supports three route redistribution functions:

- Intra-VRF inter-protocol route redistribution (redistribution within the same VRF instance), for example, redistribute RIP to OSPF.

- Inter-VRF inter-protocol redistribution (redistribution between two VRF instances), for example, redistribute RIP in VRF 2 to OSPF in VRF 4.
- Inter-VRF static routes (for example, a static route in a given VRF instance) configured as a typical static route but with the added parameter of a next-hop-vrf (the next-hop IP address is found in the next-hop-vrf instance).

With inter-VRF route redistribution, a user in one VRF instance can access route data in other VRF instances. You can redistribute routes within a VRF instance or between VRF instances; for example, one VRF instance can redistribute routes to all other VRF instances. You can redistribute Local, static, OSPF, RIP, and BGP routes and both dynamic (OSPF, BGP, and RIP) and static route redistribution is supported.

More than one routing protocol can be present in each VRF instance. Route redistribution can occur either between different protocol types, or between the same protocol types on different VRF instances.

An interface uses redistribution to announce routes that are learned by other protocols (OSPF or BGP, for example). Control route redistribution by using route policies. When you associate routing policies with route redistribution, the policy is checked before the target protocol is updated. Across VRF instances, the policy is checked at the source VRF instance, so only qualified routes are added to the routing table.

You can use static route commands to inject one specific route (including a default route) from one VRF instance to another. The route is added to the target VRF instance, while the next hop is resolved by the next-hop VRF instance.

Static routes are used to direct packets from a given source using a next-hop IP address. The next-hop-vrf option in a static route permits this path to proceed from one VRF to another. Overlapping IP addresses are supported within VRFs, thus it is possible for two VRFs to have identical IP addresses.

The following list describes interVRF route redistribution:

- Redistributed routes are added to the target VRF instance, and their next hop remains in the source VRF instance.
- If either the source or destination VRF instance is deleted, the redistribution configuration is automatically deleted.
- Redistributed routes are not further redistributed to another VRF instance.
- Route redistribution is unidirectional. You must configure route redistribution for the reverse direction if you require it. You can configure different route policies for each direction.
- After you configure interVRF route redistribution between two VRF instances, you must avoid using overlapping IP addresses between these two VRF instances.

Avoid overlapping addresses; the device does not generate an error if addresses overlap.

- Intra-VRF routes take precedence over inter-VRF routes.
- You can physically connect two VRF instances to distribute route across VRF instances (in this case, you do not need to configure route redistribution).

Route Redistribution Operation

To perform redistribution, the device maintains a route change list. The change list contains all the best routes that are either added to or deleted from the forwarding table. When a best route is added to or deleted from the forwarding table, the change list is updated to reflect the change and notify registered protocols. The registered protocols pick up the change from the change list when it becomes available.

An example scenario of interVRF redistribution follows. To redistribute OSPF routes in VRF 1 to RIP in VRF 0:

- Create, enable, and apply a RIP redistribution instance. The source protocol is OSPF and the VRF source is VRF 1.
- When an OSPF route is added in VRF 1, the Routing Table Manager (RTM) in VRF 1 puts the new route into the change list.
- The device notifies RIP in VRF 0, because RIP is registered with the RTM of VRF 1 for OSPF route changes.
- To send OSPF routes from VRF 1 through the RIP interface in VRF 0, the interface uses a route policy with match VRF criterion of VRF 1.

The switch also supports inter-domain multicast routing. For more information, see [IP Multicast](#) on page 1696.

Port parameters and VRF Lite management

You can configure each VRF instance as a separate router, this means that you can configure different routing protocols and associated parameters for each instance. You can associate non0 VRF instances with ports.

The port parameters that you can edit for a VRF instance depend on whether the port belongs to only one, or more than one, VRF instance. For example, if a port belongs to only one VRF, you can edit the port parameters of the VRF. If a port belongs to more than one VRF instance, you cannot edit the port parameters of that port unless you are accessing the port through the Global Router with read-write-all access. If you do not have read-write-all access, you can only edit the GlobalRouter port parameters. If a port belongs to a single non0 VRF, the port parameters can be changed by this VRF. If a port belongs to multiple VRF instances, only a user with read-write-all access who is accessing the port through the Global Router can change this port configuration.

Virtualized Management Applications

Virtualized management applications, such as **ping** and **tracert**, operate using the specified user-created VRF context. If you do not specify the VRF context in the command syntax, **ping** defaults to the Global Router VRF.

Ping test for IPv4:

```
Switch:1(config)#ping 192.0.2.1 vrf vrfRED
192.0.2.1 is alive
```

Ping test for IPv6:

```
Switch:1(config)#ping 2001:db8::1 vrf vrfRED
2001:db8::1 is alive
```

Traceroute test for IPv4:

```
Switch:1#traceroute 192.0.2.1 vrf vrfRED
```

Traceroute test for IPv6:

```
Switch:1#traceroute 2001:db8::1 vrf vrfRED
```

If you want to ping a device through the out-of-band management port, you must include the *mgmt oob* command parameters to specify the Segmented Management Instance as the source for the outgoing ICMP ECHO packet. The packet goes out this specific interface only.

If you do not specify the management interface type, the **ping** command uses the management routing table to determine the best management interface, and selects the source IP based on the egress management interface.

For more information about how to use a Segmented Management Instance with these commands, see [Ping an IP Device](#) on page 3901 and [Run a Traceroute Test](#) on page 3905.

VRF Lite Configuration Rules

You must select the VRF for global IPv4 or IPv6 options before entering commands; not all Global Router parameters are configurable on other VRF instances.

Layer 1 and Layer 2 information (including VLAN information) is global and is not maintained for each VRF instance. However, you can associate a set of VLANs with a VRF instance.

A VLAN cannot belong to more than one VRF instance at a time. When you create a VLAN, more than one physical port can belong to it. You can associate a VRF instance with more than one IPv4 or IPv6 interface (a physical Ethernet port or a VLAN).

Perform physical port assignment at the VLAN and brouter port level. A VRF instance inherits all the ports assigned to its VLANs and brouter ports. You cannot directly assign a physical port to a VRF instance, but it is implicitly assigned when you associate the VRF with VLANs or brouter ports.

For IPv4, after you configure interVRF route redistribution between two VRF instances, avoid overlapping IP addresses between these two VRF instances.

When you configure VRF Lite, remember the following rules:

- You can connect two VRFs from the same system with an external cable.
- An IPv4 or IPv6 routable VLAN can become a member of a VRF.
- An IPv4 or IPv6 interface can belong to only one VRF.
- A VRF can exist even if no interfaces are assigned to it.
- Routing policies apply to VRFs on an individual basis.
- Multiple VRFs on the same node can function in different autonomous systems.

Following rules apply to IPv4 interfaces specifically:

- If you configure an IPv4 interface without specifying the VRF instance, it is mapped to VRF 0 by default.
- VRF Lite supports SMLT and RSMLT.
- VRF Lite supports RIP in and out policies.
- VRF Lite supports OSPF in and out (accept and redistribute) policies.
- Before you delete a VRF instance, disable OSPF. Deleting a VRF instance deletes the OSPF instance if OSPF is disabled.
- When you create a VRF instance, an OSPF instance is not automatically created. To activate OSPF on a VRF instance, first create an OSPF instance, and then enable OSPF.
- You can configure a VRF so it can have IP interfaces with OSPF, RIP, static routes, and policies simultaneously.
- Every IPv4 interface is a member of VRF 0 unless explicitly defined to belong to another VRF.

4220 Series VRF Support

4220 Series switches support only a single active VRF with IP configuration. The VRF can be the Global Routing Table (GRT) or a non-default VRF:

- The GRT becomes the active VRF if you attach an IP interface to a VLAN. Otherwise, the first non-default VRF you create becomes the active VRF.
- If the GRT is the active VRF, you cannot create a second VRF.
- You cannot create more than one non-default VRF.
- You cannot configure IP interfaces in the GRT if the active VRF is a non-default VRF.
- To transition the active VRF from the GRT to a non-default VRF, remove all IP interfaces and then create the non-default VRF.
- To transition the active VRF from a non-default VRF to the GRT, first remove all router, ARP, and IP interfaces. Delete the VRF, and then add IP interfaces to the GRT.

Support for a single active VRF affects other features that rely on VRFs with IP configuration:

- ExtremeCloud SD-WAN—You must specify the VRF name that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.
- Fabric Extend—on platforms that support multiple VRFs with IP configuration, you cannot configure the tunnel source address (**ip-tunnel-source-address** command) if the address belongs to a VRF with an attached I-SID, which presents an issue for single VRF platforms. On platforms with only one active VRF, you can use an *overlay* parameter to bypass this restriction.
- Fabric Extend—you must configure a route-map policy to suppress IS-IS redistribution of the FE tunnel subnet:
 - Configure route-maps to not permit redistribution of the local route used as the tunnel source address (**ip-tunnel-source-address** command).

- Configure an accept policy to deny IS-IS routes that overlap with the destination tunnel IP address.
- IP Multicast config-lite for Fabric Connect—A user-created VRF cannot be active if you configure **mvpn-isid <0-15999999>[0-511]** on a switch that supports only a single active VRF.
- IP Shortcuts and Layer 3 VSN—You cannot configure the IP source address (**ip-source-address** command) as an IP address in the GRT if the active VRF is a non-default VRF.

5320 Series VRF Support

For the 5320 Series, only the following models support more than one VRF with IP configuration:

- 5320-16P-2MXT-2X
- 5320-24T-4X-XT
- 5320-24T-24S-4XE-XT
- 5320-48P-8XE
- 5320-48T-8XE



Note

The supported VRF ID range on the 5320-16P-2MXT-2X differs from other Fabric Engine platforms. 5320-16P-2MXT-2X supports a range of 1-14. To verify the supported range on a specific model, use the CLI contextual help (?).

For VRF scaling information, see *Fabric Engine Release Notes*.

The other models support a single active VRF with IP configuration. The VRF can be the Global Routing Table (GRT) or a non-default VRF:

- The GRT becomes the active VRF if you attach an IP interface to a VLAN. Otherwise, the first non-default VRF you create becomes the active VRF.
- If the GRT is the active VRF, you cannot create a second VRF.
- You cannot create more than one non-default VRF.
- You cannot configure IP interfaces in the GRT if the active VRF is a non-default VRF.
- To transition the active VRF from the GRT to a non-default VRF, remove all IP interfaces and then create the non-default VRF.
- To transition the active VRF from a non-default VRF to the GRT, first remove all router, ARP, and IP interfaces. Delete the VRF, and then add IP interfaces to the GRT.

Support for a single active VRF affects other features that rely on VRFs with IP configuration:

- ExtremeCloud SD-WAN—You must specify the VRF name that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.
- Fabric Extend—on platforms that support multiple VRFs with IP configuration, you cannot configure the tunnel source address (**ip-tunnel-source-address**

command) if the address belongs to a VRF with an attached I-SID, which presents an issue for single VRF platforms. On platforms with only one active VRF, you can use an *overlay* parameter to bypass this restriction.

- Fabric Extend—you must configure a route-map policy to suppress IS-IS redistribution of the FE tunnel subnet:
 - Configure route-maps to not permit redistribution of the local route used as the tunnel source address (**ip-tunnel-source-address** command).
 - Configure an accept policy to deny IS-IS routes that overlap with the destination tunnel IP address.
- IP Multicast config-lite for Fabric Connect—A user-created VRF cannot be active if you configure **mvpn-isid <0-15999999>[0-511]** on a switch that supports only a single active VRF.
- IP Shortcuts and Layer 3 VSN—You cannot configure the IP source address (**ip-source-address** command) as an IP address in the GRT if the active VRF is a non-default VRF.

VRF Lite Configuration using CLI

Use Virtual Router and Forwarding (VRF) Lite to provide many virtual routers using one switch.

The tasks in this section show you how to configure a VRF instance and how to associate ports and VLANs with VRF instances.

Create a VRF Instance

About This Task

Create a VRF instance to provide a virtual routing interface.



Note

You cannot configure BGP, OSPF, or RIP on a VRF instance unless you first configure the routing protocol trigger.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

2. Create a VRF instance by specifying a VRF name:

```
ip vrf WORD<1-16> [vrfid <1-511>]
```

**Note**

The single VRF models of the 5320 Series display the following warning and prompt after you enter this command: Warning: creating a new VRF will activate Unicast Routing on it and this will de-activate Unicast routing on existing VRF GlobalRouter. Are you sure about creating new VRF? (y/n) ?

3. Configure the maximum number of routes:

- For IPv4: `ip vrf WORD<1-16> max-routes <max-routes>`
- For IPv6: `ip vrf WORD<1-16> ipv6-max-routes <max-routes>`

4. Enable max-routes traps:

- For IPv4: `ip vrf WORD<1-16> max-routes-trap enable`
- For IPv6: `ip vrf WORD<1-16> ipv6-max-routes-trap enable`

**Note**

The maximum number of IPv6 routes for the Global Router are fixed and cannot be changed.

5. Enter VRF Router Configuration mode:

```
router vrf WORD<1-16>
```

6. Configure the IP routing protocol triggers for the VRF:

Use one of the following commands on your switch:

- `ip bgp` creates both IPv4 and IPv6 instances.
- `ip ospf` creates an OSPFv3 instance.
- `ip rip` creates an IPv4 instance. RIPng is not virtualized

7. Ensure that the instance is configured correctly:

```
show ip vrf [detail] [WORD<1-16>]
```

Examples

Create a VRF instance:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip vrf vrfRED
```

Configure the maximum number of IPv4 routes and enable max-routes traps.

```
Switch:1(config)#ip vrf vrfRED max-routes 12000
Switch:1(config)#ip vrf vrfRED max-routes-trap enable
```

Enter Router Configuration mode and configure the routing protocol triggers for the VRF:

```
Switch:1(config)#router vrf vrfRED
Switch:1(router-vrf)#ip bgp
```



```
Switch:1(router-vrf)#ip ospf
Switch:1(router-vrf)#ip rip
```

To Configure OSPFv3 instance for the VRF:

```
Switch:1(config)#router vrf vrfRED
Switch:1(router-vrf)#ipv6 ospf
```

Exit to Global configuration mode:

```
Switch:1(router-vrf)#exit
```

Configure the maximum number of IPv6 routes and enable IPv6 max-routes traps.

```
Switch:1(config)#ip vrf vrfRED ipv6-max-routes 7700
Switch:1(config)#ip vrf vrfRED ipv6-max-routes-trap enable
```

Ensure that the instance is configured correctly:

```
Switch:1#show ip vrf vrfRED
```

```
=====
VRF INFORMATION
=====
```

VRF COUNT	VLAN COUNT	ARP COUNT	RIP COUNT	OSPF COUNT	BGP COUNT	PIM COUNT	NBRv6 COUNT	RIPng COUNT	OSPFv3 COUNT	PIM6 COUNT	VRF IDs ALLOCATED
2	5	0	1	1	1	1	0	1	1	1	2

```
=====
```

VRF NAME	VRF ID	VLAN COUNT	ARP COUNT	RIP	OSPF	BGP	PIM	NBRv6 COUNT	RIPng	OSPFv3	PIM6	UNICAST ACTIVE	SD-WAN BREAKOUT	LOCAL	VRF ORIGIN	SEGMENTED
vrfRED	1	0	0	TRUE	TRUE	TRUE	TRUE	0	FALSE	FALSE	TRUE	TRUE	FALSE		CONFIG	FALSE

```
=====
1 out of 2 Total Num of VRF Entries displayed.
```

View the VRF configuration on the switch:

```
Switch:1#show ip vrf
```

```
=====
VRF INFORMATION
=====
```

VRF COUNT	VLAN COUNT	ARP COUNT	RIP COUNT	OSPF COUNT	BGP COUNT	PIM COUNT	NBRv6 COUNT	RIPng COUNT	OSPFv3 COUNT	PIM6 COUNT	VRF IDs ALLOCATED
3	9	1	1	2	1	1	0	1	2	1	3

```
=====
```

VRF NAME	VRF ID	VLAN COUNT	ARP COUNT	RIP	OSPF	BGP	PIM	NBRv6 COUNT	RIPng	OSPFv3	PIM6	UNICAST ACTIVE	SD-WAN BREAKOUT	LOCAL	VRF ORIGIN	SEGMENTED
GlobalRouter	0	7	1	TRUE	TRUE	TRUE	TRUE	0	TRUE	TRUE	TRUE	TRUE	FALSE		DYNAMIC	FALSE
l3vsn-100	1	1	0	FALSE	TRUE	FALSE	FALSE	0	FALSE	TRUE	FALSE	TRUE	FALSE		DYNAMIC	FALSE
vrfred	1	0	0	TRUE	TRUE	TRUE	TRUE	0	FALSE	FALSE	TRUE	TRUE	FALSE		CONFIG	FALSE
ospf	3	2	6	FALSE	TRUE	FALSE	FALSE	0	FALSE	FALSE	FALSE	TRUE	FALSE		CONFIG	FALSE
sd-wan	511	1	3	FALSE	FALSE	FALSE	FALSE	0	FALSE	FALSE	FALSE	FALSE	FALSE		DYNAMIC	FALSE
MgmtRouter	512	1	0	FALSE	FALSE	FALSE	FALSE	0	FALSE	FALSE	FALSE	TRUE	FALSE		DYNAMIC	FALSE

```
=====
6 out of 6 Total Num of VRF Entries displayed.
```

Variable Definitions

The following table defines parameters for the **ip vrf** command.

Variable	Value
<code>max-routes <max-routes></code> Note: Exception: not supported on 5320-16P-4XE, 5320-24P-8XE, or 5320-24T-8XE.	Specifies the maximum number of IPv4 routes for the VRF. Depending on the hardware platform, the parameter range and default value can vary. Use the CLI Help to see the available range for the switch. For route scaling information, see <i>Fabric Engine Release Notes</i> .
<code>ipv6-max-routes <max-routes></code> Note: Exception: not supported on 5320-16P-4XE, 5320-24P-8XE, or 5320-24T-8XE.	Specifies the maximum number of IPv6 routes for the VRF. Depending on the hardware platform, the parameter range and default value can vary. Use the CLI Help to see the available range for the switch. For route scaling information, see <i>Fabric Engine Release Notes</i> .
<code>max-routes-trap enable</code>	Enables SNMP traps after the maximum number of IPv4 routes are reached.
<code>ipv6-max-routes-trap enable</code>	Enables SNMP trap generation based on the configured number of maximum IPv6 routes. The default is enabled.
<code>name WORD<0-16></code>	Renames the VRF instance.
<code>vrf-trap</code>	Enables the device to send VRF-related traps.
<code>vrfid <1-511></code>	Specifies a VRF ID. Note: The supported VRF ID range on the 5320-16P-2MXT-2X differs from other Fabric Engine platforms. 5320-16P-2MXT-2X supports a range of 1-14. To verify the supported range on a specific model, use the CLI contextual help (?).
<code>WORD<1-16></code>	Names the VRF instance.

The following table defines parameters for the **show ip vrf** command.

Variable	Value
<code>detail</code>	Displays detailed VRF information, including Segmented VRF information.
<code>max-routes [vrfids WORD<0-512>] [WORD <1-16>]</code>	Displays the maximum number of routes for the specified VRFs. <ul style="list-style-type: none"> • <code>vrfids WORD<0-512></code> specifies a list of VRFs by VRF IDs. • <code>WORD<1-16></code> specifies a VRF by name.

Variable	Value
<code>vrfids WORD<0-512></code>	Specifies a list of VRFs by VRF IDs.
<code>WORD<1-16></code>	Specifies a VRF by name.

Associate a VLAN or Port with a VRF Instance

You can assign a VRF instance to a port or VLAN. You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IP address. You can configure the IP address after you associate the port and VRF instance.

Before You Begin

- Ensure the VRF is already configured.

About This Task

You can assign a VRF instance to a port after you configure the VRF. The system assigns ports to the Global Router, VRF 0, by default.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Associate the port or VLAN with a VRF instance:

```
vrf WORD<1-16>
```

Example

```
Switch:1> enable  
Switch:1# configure terminal
```

Create a VRF named Two:

```
Switch:1(config-if)# ip vrf Two
```

Create a VLAN of type byport:

```
Switch:1(config-if)# vlan create 33 name vlan-30 type port-mstprstp 0
```

Enter VLAN Interface Configuration mode:

```
Switch:1(config-if)# interface vlan 33
```

Assign the VLAN to VRF Two:

```
Switch:1(config-if)# vrf Two
```

Assign an IP address to the VLAN:

```
Switch:1(config-if)# ip address 192.0.2.2/24
```

Assign a secondary IP address to the VLAN:

```
Switch:1(config-if)# ip address 192.0.4.2/24
```

Enter VRF configuration mode:

```
Switch:1(config-if)# router vrf Two
```

Variable Definitions

Use the data in the following table to use the **vrf** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the VRF instance by name.

Create an IP VPN Instance on a VRF

Create an IP VPN instance to advertise IP routes from one VRF to another across a Shortest Path Bridging MAC (SPBM) network.

Before You Begin

- The VRF must exist.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create an IP VPN instance on the VRF:

```
ipvpn
```

3. Assign a service instance identifier (I-SID) to the IP VPN:

```
i-sid <0-16777215>
```

4. Enable IP VPN on the VRF:

```
ipvpn enable
```

By default, a new IP VPN instance is disabled.

5. Display all IP VPNs:

```
show ip ipvpn [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

From Global Configuration mode, log on to Router VRF Configuration mode:

```
Switch:1(config)# router vrf red
```

Create the IP VPN instance:

```
Switch:1(router-vrf)# ipvpn
```

Enable IP VPN:

```
Switch:1(router-vrf)# i-sid 100
```

Enable IP VPN:

```
Switch:1(router-vrf)#ipvpn enable
```

```
=====
                        IPv4  IPVPN
=====
VRF Name          VRF ID   IPv4  IPVPN   IPv6  IPVPN   I-SID   I-SID Name
-----
red                1         enabled  disabled  100     ISID-1
-----
1 out of 1 Total active IPv4 L3 VSN displayed, including 0 active IPv6 L3 VSN.
```

Variable definitions

Use the data in the following table to use the **show ip ipvpn** command.

Variable	Value
<i>vrf WORD<1-16></i>	Specifies the VRF name.
<i>vrfids WORD<0-512></i>	Specifies the VRF ID.

Use the data in the following table to use the **i-sid** command.

Variable	Value
<i>i-sid <0-16777215></i>	Assigns an I-SID to the VRF to configure. Use the no or default option to remove the I-SID to VRF allocation for this VRF.

Create an IPv6 VPN Instance

Before You Begin

The VRF must exist.

About This Task

Create an IPv6 VPN instance to advertise IPv6 routes from a VRF to Shortest Path Bridging MAC (SPBM) network. For more information about Layer 3 Virtual Services Networks (VSNs) and SPBM, see [Layer 3 VSN Configuration](#) on page 1654.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
```

```
configure terminal
```

```
router vrf WORD<1-16>
```



Note

You cannot use the derived name for the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF to enter this mode.

2. Create an IPv6 VPN instance:

```
ipv6 ipvpn
```

3. Assign a service instance identifier (I-SID) to the IPv6 VPN:

```
i-sid <0-16777215>
```

4. Enable IPv6 VPN:

```
ipv6 ipvpn enable
```

5. Display all IPv6 VPNs:

```
show ipv6 ipvpn [vrf WORD<1-16> | vrfids WORD<0-512>]
```

Example

Create and enable IPv6 VPN instance:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf vrfred
Switch:1(router-vrf)#ipv6 ipvpn
Switch:1(router-vrf)#i-sid 100
Switch:1(router-vrf)#ipv6 ipvpn enable
Switch:1(router-vrf)#show ipv6 ipvpn
```

```
=====
                        IPv6  IPVPN
=====
VRF Name          VRF ID   IPv6  IPVPN   IPv4  IPVPN   I-SID   I-SID Name
-----
vrfred             2       enabled disabled  5555   ISID-5555
-----
1 out of 1 Total IPv6 L3 VSN, 1 active IPv6 and 0 active IPv4 displayed.
```

Variable definitions

Use the data in the following table to configure the **ipv6 ipvpn** command.

Variable	Value
<i>enable</i>	Enables IPv6 IPVPN. The default is disabled.

Use the data in the following table to configure the **i-sid** command.

Variable	Value
<code><0-16777215></code>	Assigns an I-SID to the VRF being configured.

Use the data in the following table to configure the **show ipv6 ipvpn** command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

Configure the Maximum Number of VRFs

Perform this procedure to change the maximum number of VRFs and Layer 3 VSNs that the switch supports. Increasing the number of VRFs or Layer 3 VSNs can be useful in a WAN scenario or other large network.

The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Fabric Engine Release Notes*.



Note

This procedure does not apply to all products. For feature support information, see *Fabric Engine and VOSS Feature Matrix*.

Before You Begin

- If the switch operates in SPBM mode, before you enable the boot config flag, perform the following actions:
 - Check in-VLAN filters. If a filter references a VLAN in the 3500 to 3998 range, you must delete the filter or the filter configuration fails when you restart the switch.
 - Delete VLANs in the 3500 to 3998 range.

About This Task



Important

If you enable this boot config flag, and the switch operates in SPBM mode (default configuration), the switch reduces the number of configurable VLANs. In such a configuration, the switch reserves VLANs 3500 to 3998 for internal use. You cannot use these VLANs as either platform VLANs or B-VLANs. You can still use the reserved VLAN range for customer VLANs (C-VLAN) on Flex UNI and B-VLANs on FE-VID.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Increase the maximum number of VRFs and Layer 3 VSNs:
`boot config flag vrf-scaling`
3. Verify the configuration:
`show boot config flags`
4. Save the configuration:
`save config`
5. Restart the switch for the change to take effect:
`reset`

Examples

Enable the boot config flag to increase the maximum number of VRFs and Layer 3 VSNs. In the following example, the switch operates in SPBM mode and reserves the VLAN ID range of 3500 to 3999. If the switch does not operate in SPBM mode, the system does not display the VLAN warning message when you enable VRF scaling.

```
Switch:1>en
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flag vrf-scaling
Warning: Vlan 3500 to 3999 will be reserved for internal use.

Warning: Please save the configuration and reboot the switch
         for this configuration to take effect.
```



Note

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags logging true
flags macsec false
flags nni-mstp false
flags reboot true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags spbm-node-scaling true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
```



```

flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false

```

The following example shows the message that the system displays if you try to enable the boot config flag and configured VLANs use IDs between 3500 and 3999.

```

Switch:1>en
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flag vrf-scaling

Error: Delete all configured platform vlans between 3500 and 3999 to enable vrf-scaling.

```

Enable IPv6 Trap Notifications

About This Task

Perform this procedure to enable SNMP traps when maximum number of IPv6 routes are reached.



Note

Different hardware platforms support different parameter ranges. Use the CLI Help to see the available range.

Procedure

1. Enter Global Configuration mode:


```
enable

configure terminal
```
2. Enable max-routes trap:


```
ip vrf WORD<1-16> ipv6-max-routes-trap enable
```

Example

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip vrf vrfRED ipv6-max-routes-trap enable

```

Variable definitions

Use the data in the following table to use the **ipv6-max-routes-trap** command.

Variable	Value
<i>enable</i>	Enables SNMP trap generation based on the configured number of maximum IPv6 routes. The default is enabled.

Display IPv6 max-routes Information

About This Task

Perform this procedure to display the maximum IPv6 routes configured.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display max-routes information:
show ip vrf ipv6-max-routes [vrfids *WORD*<0-512> | *WORD*<1-16>]

Example

```
Switch:1#show ip vrf ipv6-max-routes
```

=====					
VRF Specific Configuration					
=====					
VRF-ID	VRF-NAME	CONTEXT-NAME	IPV6-MAX-ROUTES	IPV6-MAX-ROUTES-TRAP	VRF-TRAP

0	GlobalRouter		8000	enable	NA
1	vrfred	vrf1	5000	enable	enable
2	vrfblue	vrf2	5000	enable	enable
3 out of 3 Total Num of VRF Entries displayed.					

VRF Lite Configuration Using EDM

Use VRF Lite to provide many virtual routers using a single switch.

Configure a VRF Instance

About This Task

Configure a VRF instance to provide a virtual routing interface for a user.



Note

The maximum routes of IPv4 and IPv6 for Global Router (GRT) are non-configurable and fixed at the system limits.

Maximum route traps are not generated on GRT. For non-default VRFs, the permitted maximum routes can be lower than system limits and traps generate when the limit is exceeded.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **VRF** tab.
4. Select **Insert**.
5. Specify the VRF ID.

6. Type a name for the VRF instance.
7. Configure VRF Lite-related traps.
8. Configure the optional parameters as required.
9. Select **Insert**.

VRF Field Descriptions

Use the data in the following table to help you use the **VRF** tab.

Name	Description
Id	Specifies the ID number of the VRF instance. VRF ID 0 is reserved for the GlobalRouter.
Name	Names the VRF instance.
ContextName	Identifies the VRF. The SNMPv2 Community String or SNMPv3 contextName denotes the VRF context and is used to logically separate the MIB port management.
TrapEnable	Enables the VRF to send VRF Lite-related traps (VrfUp and VrfDown). The default is enabled.
MaxRoutes Note: Exception: not supported on 5320-16P-4XE, 5320-16P-4XE-DC, 5320-24P-8XE, or 5320-24T-8XE.	Configures the maximum number of routes allowed for the VRF, which varies depending on the hardware. For scaling information, see <i>Fabric Engine Release Notes</i> . The default value varies for the GlobalRouter and non-default VRFs, depending on your hardware platform.
RpTrigger	Specifies the Routing Protocol (RP) triggers for the VRF. The triggers are used to initiate or shutdown routing protocols on a VRF. You can act upon multiple RPs simultaneously. You can also use this option to bring individual RPs up in steps.
Segmented Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Specifies if the VRF is a Segmented VRF. The default is false, which indicates the VRF is not segmented.
MaxRoutesTrapEnable	Enables the generation of the VRF Max Routes Exceeded traps. The default is enabled.
Ipv6MaxRoutes Note: Exception: not supported on 5320-16P-4XE, 5320-16P-4XE-DC, 5320-24P-8XE, or 5320-24T-8XE.	Configures the maximum number of IPv6 routes allowed for the VRF, which varies depending on the hardware. For scaling information, see <i>Fabric Engine Release Notes</i> . The default value varies for the GlobalRouter and non-default VRFs, depending on your hardware platform.

Name	Description
Ipv6MaxRoutesTrapEnable	Enables SNMP trap generation after the maximum number of IPv6 routes are reached. The default is enabled.
Active	Displays if the VRF is active (true). Some switch models support only a single active VRF with IP configuration; the VRF can be the GlobalRouter (GRT) or a non-default VRF. For information about VRF support, see <i>Fabric Engine Release Notes</i> or <i>Fabric Engine and VOSS Feature Matrix</i> .
Origin	Displays the VRF origin: <ul style="list-style-type: none"> • config • dynamic • segmented The Default VRF and the MGMT VRF are considered DYNAMIC. Note: Only manually configured VRFs are visible in running-config.
Sd-Wan Local Breakout Note: Not supported on switch models that support only a single active VRF with IP configuration.	Enables Local Breakout on the VRF. The default is false (disabled).
BgpOrigin	Specifies the origin of the BGP configuration in the VRF. config indicates the BGP instance was created manually. autoSense indicates the BGP instance was generated from Auto-sense. default is for Global Routing Table (GRT) only and indicates no BGP configuration.

Associate a Port to a VRF Instance

About This Task

Associate a port to a Virtual Router Forwarding (VRF) instance so that the port becomes a member of the VRF instance.

You can assign a VRF instance to a port after you configure the VRF. The system assigns ports to the GlobalRouter, VRF 0, by default.

Procedure

1. In the **Device Physical** View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **General**.

4. Select the **VRF** tab.
5. Next to **BrouterVrfId**, select the ellipsis (...).
6. Select the required VRF.
7. Select **OK**.
8. Select **Apply**.

VRF Field Descriptions

Use the data in the following table to use the **VRF** tab.

Name	Description
VrfIds	Shows the VRF ID.
VrfNames	Shows the VRF name.
VrfCount	Shows the number of VRFs to which the port is associated.
BrouterVrfId	Specifies the VRF ID for a brouter port.
BrouterVrfName	Shows the VRF name for a brouter port.

Associate a VLAN with a VRF

You can assign a VRF instance to a VLAN. You cannot associate a VLAN and a VRF instance if the VLAN has an IP address. You can configure the IP address after you associate the VLAN and VRF.

Before You Begin

- Ensure the VRF is already configured.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**
2. Select **VLANs**.
3. Select the **Basic** tab.
4. In **VrfId** in the row for the VLAN, select a VRF ID.
5. Select **Apply**.

Associate an Extreme Integrated Application Hosting Port to a VRF Instance



Note

This procedure only applies to 5720-24MXW, 5720-48MXW, 7520 Series, and 7720 Series.

About This Task

Perform this procedure to associate an Extreme Integrated Application Hosting (IAH) port to a Virtual Router Forwarding (VRF) instance.



Note

You can associate a VRF instance to an IAH port after you configure the VRF. By default, the IAH ports are associated to the GlobalRouter.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Insight Port**.
2. Select the IAH port you want to configure.
3. Select the **VRF** tab.
4. (Optional) In the **VrfNames** field, select the **ShowAll** button to view the VRF instances the IAH Port is associated with.
5. In the **BrouterVrflid** field, select the ellipsis (...) button, and select the required VRF instance(s).
6. Select **Ok**.
7. Select **Apply**.

VRF Field Descriptions

Use data in the following table to use the **VRF** tab.

Name	Description
Vrflids	Shows the VRF ID.
VrfNames	Shows the VRF name.
VrfCount	Shows the number of VRFs to which the Extreme Integrated Application Hosting (IAH) port is associated with.
BrouterVrflid	Specifies the VRF ID for a brouter port.
BrouterVrfName	Shows the VRF name for a brouter port.

Configure interVRF route redistribution policies

Configure interVRF route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF, RIP, or BGP. Use a route policy to control the redistribution of routes.

Before You Begin

- Ensure VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required.

Procedure

1. In the navigation pane, expand **Configuration > IP**.

2. Select **Policy**.
3. Select the **Route Redistribution** tab.
4. Select **Insert**.
5. Select the ellipsis (...) near the **DstVrflid** box to select the source and destination VRF IDs.
6. Select the ellipsis (...) near the **SrcVrflid** box to select the source and destination VRF IDs.
7. In the **Protocol** option box, select the protocol.
8. In the **RouteSource** option box, select the route source.
9. Select **Enable**.
10. Choose the route policy to apply to the redistributed routes.
11. Configure other parameters as required.
12. Select **Insert**.
13. Select the **Applying Policy** tab.
14. Select **RedistributeApply**.
15. Select **Apply**.

Route Redistribution Field Descriptions

Use the data in the following table to use the **Route Redistribution** tab.

Name	Description
DstVrflid	Specifies the destination VRF ID to use in the redistribution.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrflid	Specifies the source VRF ID to use in the redistribution.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution. The default is disabled.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine whether the system advertises a specific route to the specified protocol.
Metric	Specifies the metric announced in advertisements. The default is 0.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone. The default is type2.
Subnets	Indicates that all the subnets must be advertised individually. The values are allow(1), and suppress(2). The default value is allow. This variable applies to OSPF only.

View Brouter Port and VRF Associations

About This Task

You can view each port and associated VRFs. You can also change the VRFs associated with the port if the port has no IP address.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **VRF-Ports** tab.
4. To display the VRF names associated with a port, select a cell in one of the table rows, and then select **ShowVRFNames**.
5. To change the VRF, double-click the **BrouterVrfId** field for the port.



Tip

You can associate a port with more than one VRF.

6. Select the required VRFs, and then select **Ok**.
7. Select **Apply**.

VRF-Ports Field Descriptions

Use the data in the following table to use the **VRF-Ports** tab.

Name	Description
Index	Specifies the slot and port.
Type	Specifies the port type.
VrfIds	Identifies the set of VRF IDs to which this port belongs.
VrfCount	Shows the number of VRF instances associated with this port.
BrouterVrfId	Shows the VRF ID for this brouter port.
BrouterVrfName	Shows the VRF name for this brouter port.
Show VrfNames	You can use this toolbar button to identify the set of VRF names to which a port belongs.

Use the data in the following table to use the **Show VrfNames** button.

Name	Description
Index	Specifies the slot and port.
VrfNames	Shows the VRF name for this brouter port.

View Global VRF Status Information

About This Task

View global VRF status information to determine the number of VRFs that are configured and active.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **Global Status** tab.

Global Status Field Descriptions

Use the data in the following table to use the **Global Status** tab.

Name	Description
ConfigNextAvailableVrflid	Specifies the number of the next available Virtual Router ID (index).
ConfiguredVRFs	Specifies the number of VRFs configured on this network element.
ActiveVRFs	Specifies the number of VRFs that are active on the network element. These are VRFs for which the OperStatus is up.

View VRF Instance Statistics and Status Information

About This Task

View VRF instance status information to determine the operational status of each VRF, as well as other operational parameters.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **VRF Stats** tab.

VRF Stats Field Descriptions

Use the data in the following table to use the **VRF Stats** tab.

Name	Description
Id	Specifies the ID number of the VRF instance.
StatRouteEntries	Specifies the total number of routes for this VRF.
StatFIBEntries	Specifies the total number of Forwarding Information Base (FIB) entries for this VRF.
StatUpTime	Specifies the time in (in hundredths of a second) since this VRF entry has been operational.

Name	Description
OperStatus	Shows the operational status of the Virtual Router.
RouterAddressType	Specifies the router address type of this VRF.
RouterAddress	Specifies the router address of this VRF, derived from one of the interfaces. If a loopback interface is present, you can use the loopback interface address.
DhcpRelayNumRequests	Specifies the total number of DHCP requests on this VRF.
DhcpRelayNumReplies	Specifies the total number of DHCP replies on this VRF.

View Statistics for a VRF

About This Task

View VRF statistics to ensure the instance is performing as expected.

Procedure

1. In the navigation pane, expand **Configuration > VRF**.
2. Select **VRF**.
3. Select the **VRF** tab.
4. Select a VRF.
5. Select the **Stats** button.

Stats Field Descriptions

Use the data in the following table to use the **Stats** tab.

Name	Description
StatRouteEntries	Specifies the number of routes for this VRF.
StatFIBEntries	Specifies the number of Forwarding Information Base (FIB) entries for this VRF.

Select and Launch a VRF Context View

Use this procedure to open a VRF context view when you use EDM. GlobalRouter is the default view at log in. If you launch a VRF context view, EDM displays the VRF navigation in a new tab.



Note

You cannot launch a VRF Context View for the untrusted or unrestricted segments of a Segmented VRF.

The Set VRF Context view function is not available to users in a service provider deployment where only a tenant VRF view is assigned.

About This Task



Important

If you log out from the GRT view, the system generates a warning: all tabs close and your session terminates. If you close a VRF view tab, you close only that view.

You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view. You can open five tabs for each EDM session.

Procedure

1. In the navigation pane, expand **Configuration > VRF Context View**.
2. Select **Set VRF Context View**.
3. Select the **VRF** tab.
4. Select a context to view.
5. Select **Launch VRF Context view**.

A new browser tab opens containing the selected VRF view

VRF Field Descriptions

Use the descriptions in the following table to use the **VRF** tab.

Name	Description
Id	Shows the unique VRF ID.
Name	Shows the name of the virtual router.
ContextName	Shows the SNMPv3 context name that denotes the VRF context and logically separates the MIB port management.

Create an IP VPN Instance on a VRF

Create an IP VPN instance to advertise IP routes from one VRF to another across a Shortest Path Bridging MAC (SPBM) network.

Before You Begin

- You must configure the required SPBM IS-IS infrastructure.
- The VRF must exist.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP-VPN**.
3. Select the **VPN** tab.
4. Select **Insert**.
5. Select the ellipsis ([...]), and then select a VRF from the list.
6. Select **OK**.

7. Select **Insert**.

By default, the new IP VPN instance is disabled.

8. In **Enable**, select **enable**.9. In **Trusted Isid**, specify an I-SID to associate with the VPN.10. Select **Apply**.**VPN Field Descriptions**

Use the data in the following table to use the **VPN** tab.

Name	Description
Vrflid	Specifies the ID of the VRF to configure.
Enable	Enables or disables the IP VPN instance on the VRF. The default is disabled.
Trusted Isid Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Specifies the I-SID to associate with the VPN. By default, no I-SID is assigned. For a Segmented VRF, assigns an I-SID to advertise trusted routes to trusted and unrestricted segments.
Unrestricted Isid Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Assigns an I-SID to advertise unrestricted routes to all segments. This parameter applies only to Segmented VRF. If you do not assign an I-SID to the unrestricted segment, the switch automatically derives a value based on the trusted I-SID, (trusted I-SID + 1).
Untrusted Isid Note: Exception: not supported on 7830 Series or switch models that support a single active VRF.	Assigns an I-SID to advertise untrusted routes to restricted segments. This parameter applies only to Segmented VRF. If you do not assign an I-SID to the untrusted segment, the switch automatically derives a value based on the trusted I-SID, (trusted I-SID + 2).
Isid Name	Specifies the name of the trusted I-SID associated with the VPN.

Create an IPv6 VPN Instance on a VRF

Create an IPv6 VPN instance to advertise IPv6 routes from a VRF to Shortest Path Bridging MAC (SPBM) network.

Before You Begin

- You must configure the required SPBM IS-IS infrastructure.
- The VRF must exist.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **IPv6-VPN**.
3. Select the **VPN** tab.
4. Select **Insert**.
5. Select the ellipses [...] button, and then select a VRF from the list.
6. Select **OK**.
7. Select **Insert**.
8. In the **IsidNumber** column, double-click the **0** value, and then enter the service instance identifier (I-SID) to assign to the IPv6-VPN.
9. In the **Enable** column, double-click the value and select **true** or **false** from the drop-down list.
10. Select **Apply**.

VPN Field Descriptions

Use the data in the following table to use the **VPN** tab.

Name	Description
VrfId	Specifies the ID of the VRF to configure.
Enable	Enables or disables the IPv6 VPN instance on the VRF. The default is disabled.
Isid Number	Specifies the I-SID to associate with the IPv6 VPN. By default, no I-SID is assigned.
Isid Name	Specifies the name of the I-SID.

Configure the Maximum Number of VRFs

Perform this procedure to change the maximum number of VRFs and Layer 3 VSNs that the switch supports. By default, the switch supports 24 VRFs and Layer 3 VSNs. Increasing the number of VRFs or Layer 3 VSNs can be useful in a WAN scenario or other large network.

The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Fabric Engine Release Notes*.

Before You Begin

- If the switch operates in SPBM mode, before you enable this boot config flag, perform the following actions:
 - Check in-VLAN filters. If a filter references a VLAN in the 3500 to 3998 range, you must delete the filter or the filter configuration fails when you restart the switch.
 - Delete VLANs in the 3500 to 3998 range.
- Before you disable this boot config flag, delete additional VRFs if more than 24 exist.

About This Task



Important

If you enable this boot config flag, and the switch operates in SPBM mode (default configuration), the switch reduces the number of configurable VLANs. In such a configuration, the switch reserves VLANs 3500 to 3998 for internal use. You cannot use these VLANs as either platform VLANs or B-VLANs. You can still use the reserved VLAN range for customer VLANs (C-VLAN) on Flex UNI and B-VLANs on FE-VID.

Enabling the boot config flag to use more than 24 VRFs requires a Premier or Premier + MACsec license.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **Chassis**.
3. Select the **Boot Config** tab.
4. Perform one of the following actions:
 - a. To enable VRF scaling, select the **EnablevrfScaling** check box.
 - b. To disable VRF scaling, clear the **EnablevrfScaling** check box.
5. Select **Apply**.
6. Restart the switch for the change to take effect.



Virtual Router Redundancy Protocol

[VRRP Fundamentals](#) on page 4107
[VRRP for IPv6](#) on page 4115
[VRRPv3](#) on page 4120
[Configure VRRP Using CLI](#) on page 4121
[IPv6 VRRP Configuration using CLI](#) on page 4132
[Configure VRRP using EDM](#) on page 4142
[IPv6 VRRP Configuration using EDM](#) on page 4155
[Troubleshooting IPv6 VRRP](#) on page 4166

Table 352: Virtual Router Redundancy Protocol product support

Feature	Product	Release introduced
Virtual Router Redundancy Protocol (VRRP)	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 352: Virtual Router Redundancy Protocol product support (continued)

Feature	Product	Release introduced
VRRPv3 for IPv4	4220 Series	Fabric Engine 9.2
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.3
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

Table 353: Virtual Router Redundancy Protocol for IPv6 product support

Feature	Product	Release introduced
IPv6 VRRP	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0
VRRPv3 for IPv6	4220 Series	Not Supported
	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.4
	5520 Series	VOSS 8.2.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Fabric Engine 9.4
	VSP 4900 Series	VOSS 8.1
	VSP 7400 Series	VOSS 8.0

VRRP Fundamentals

Virtual Router Redundancy Protocol (VRRP) (RFC 2338) provides another layer of resiliency to your network design by providing default gateway redundancy for end users. VRRP eliminates the single point of failure that can occur when the single static default gateway router for an end station is lost. A loss of the default gateway router causes a loss of connectivity to the remote networks. If a VRRP-enabled router that connects to the default gateway fails, failover to the VRRP backup router ensures no interruption for end users who attempt to route from their local subnet.

VRRP shares a virtual IP address (transparent to users) between two or more routers that connect the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides dynamic default gateway redundancy in the event of failover.

The VRRP router that controls the IP addresses associated with a virtual router is the primary router and it forwards packets to these IP addresses. The election process provides a dynamic transition of forwarding responsibility if the primary router becomes unavailable. (The primary router is sometimes referred to as the *master*.)



Note

The first virtual IP configured on a VRRP interface is the primary virtual IP. This primary address is the equivalent of IPv6 VRRP link-local address and must be in the same subnet as the primary IP of the interface. The address owner of a VRRP interface is determined by the primary virtual IP.

In the following figure, the first three hosts install a default route to the R1 (virtual router 1) IP address and the other three hosts install a default route to the R2 (virtual router 2) IP address. This configuration not only shares the load of the outgoing traffic, but it also provides full redundancy. If either router fails, the other router assumes responsibility for both addresses.




Figure 253: Virtual Router Redundancy Protocol configuration

For information about the number of supported VRRP interfaces, see the scaling information in *Fabric Engine Release Notes*.

The following terms are specific to VRRP:

VRRP router

a router running the VRRP protocol

Virtual router

an abstract object acting as the default router for one or more hosts, consisting of a virtual router ID and a set of addresses

Primary IP address of the interface

the real IP address of the interface. This IP address is used as the source of the VRRP advertisement. In the case of multinetting, this is the primary IP address of the interface.

Primary IP VRRP address

a virtual IP address that is in the same subnet as the primary IP address of the interface. This address is mandatory for the creation of the VRRP instance. The primary VRRP router performs the following actions:

- adds this IP address to the VRRP advertisement
- responds to ARP requests for this IP address
- routes traffic destined to this IP address

Secondary IP VRRP address / Associated IP VRRP address

an IP address that is in the same subnet as the secondary IP address of the interface. If the primary VRRP address is the same subnet as the primary IP address of the interface, the secondary IP VRRP address also needs to be the same subnet as

the secondary IP address of the interface. This IP address is not mandatory for VRRP instance creation.

Virtual primary router

the router that assumes responsibility to forward packets sent to the IP address associated with the virtual router and answer ARP requests for these IP addresses. Only the VRRP Master router forwards traffic for a given subnet. (The primary router is sometimes referred to as the *master*.)

Virtual router backup

the virtual router that becomes the primary router if the current primary router fails. The backup VRRP router does not route traffic destined for the default gateway.

When a VRRP router is initialized it sends a VRRP advertisement. The VRRP router also broadcasts a gratuitous ARP request that contains the virtual router MAC address for each IP address associated with the virtual router. The VRRP router then transitions to the controlling state.

In the controlling state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. The VRRP router responds to ARP requests for these IP addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts packets addressed to IP addresses associated with the virtual router.

In the backup state, a VRRP router monitors the availability and state of the primary router. The backup router does not respond to ARP requests and must discard packets with a MAC address equal to the virtual router MAC address. The backup router does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, the backup router transitions back to the initialize state. If the primary router goes down, the backup router sends the VRRP advertisement and ARP request described in the preceding paragraph and transitions to the controlling state.

The router transitions to the backup state in the following situations:

- If the priority is greater than the local priority
- If the priority is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address

Otherwise, the router discards the advertisement. If a shutdown occurs, the primary router sends a VRRP advertisement with a priority of 0 and transitions to the initialize state.

Critical IP Address

Within a VRRP VLAN, one link can go down while the remaining links in the VLAN remain operational. Because the VRRP VLAN continues to function, a virtual router associated with that VLAN does not register a master router failure.

As a result, if the local router IP interface connecting the virtual router to the external network fails, this does not automatically trigger a master router failover.

**Note**

In this context, local implies an address from the same VRF as the IP interface where VRRP is being configured.

The critical IP address resolves this issue. If the critical IP address fails, it triggers a failover of the master router.

You can specify the local router IP interface uplink from the VRRP router to the network as the critical IP address. This ensures that, if the local uplink interface fails, VRRP initiates a master router failover to one of the backup routers.

In VRRP, the local network uplink interface on router 1 is shown as the critical IP address for router 1. As well, the same network uplink is shown as the critical IP address for router 2. Router 2 also requires a critical IP address for cases in which it assumes the role of the master router.

With the support of VRRP and the critical IP interface linked to VRRP, you can build reliable small core networks that provide support for converged applications, such as voice and multimedia.

**Note**

A Brouter port with a VLACP critical IP address in a VRRP is supported.

VRRP Backup Master

To enable both VRRP switches to route traffic, the switch software has an extension to VRRP, the BackupMaster, that creates an active-active environment for routing. If you enable BackupMaster on the backup router, the backup router no longer switches traffic to the VRRP Master. Instead the BackupMaster routes all traffic received on the BackupMaster IP interface according to the switch routing table.

Stagger VRRP instances on a network or subnet basis.

The following figure shows the VRRP Masters and BackupMasters for two subnets.




Figure 254: VRRP network configuration

The VRRP BackupMaster uses the VRRP standardized backup switch state machine. Thus, VRRP BackupMaster is compatible with standard VRRP.

Handling of IPv4 Layer 2 Unicast Packets

Only the VRRP Master handles IPv4 Layer 2 unicast packets, for example, ARP Request and Reply, with the destination MAC as the VRRP MAC.

The BackupMaster forwards all IPv4 Layer 2 unicast packets to the Master and the Master VRRP sends an ARP reply only.

Processing of IP unicast packets (for example, ICMP packets to VRRP IP) or IPv4 routed packets (with destination MAC as VRRP MAC) on VRRP BackupMaster stays the same. For example, the VRRP BackupMaster replies to ICMP requests and routes Layer 3 routed packets to the destination and does not forward these packets to the Master when they arrive at the BackupMaster.

To reflect these changes, the VRRP MAC entry on the BackupMaster points to the Master instead of itself, and the ARP entry for VRRP IP on the BackupMaster points to local.

VRRP Fast Hello Timers

You can configure the advertisement time interval (in seconds) between sending advertisement messages. This interval permits fast network convergence with standardized VRRP failover. However, losing connections to servers for more than a second can result in missing critical failures. Customer network uptime in many cases requires faster network convergence, which means network problems must

be detected within hundreds of milliseconds. Fast Advertisement Enable and Fast Advertisement Interval meet these requirements.

Fast Advertisement Enable acts like a toggle device for the Advertisement Interval and the Fast Advertisement Interval. When Fast Advertisement Enable is enabled, the Fast Advertisement Interval is used instead of the Advertisement Interval.

When you enable the Fast Advertisement Interval, VRRP can communicate with other switch ports and networking products that have the same configuration.

Configure all routers to the identical number of seconds for the holddown timer. You can manually force the preemption of the primary over the backup before the delay timer expires.

When you use the Fast Advertisement Interval to configure a master and backup device, you must enable the Fast Advertisement Interval on both systems for VRRP to work correctly. If you configure one device with the Advertisement Interval, and the other device with the Fast Advertisement Interval, it causes an unstable state and drops advertisements.

VRRP Holddown Timer

To modify the behavior of the VRRP failover mechanism, use the holddown timer to enable the router enough time to detect and update routes. The timer delays preempting VRRP primary over the backup, when the primary becomes available.

The transition of virtual router to Master state is delayed in the following conditions:

- The VRRP holddown timer runs only when the VRRP virtual router transitions from initialization to backup to master. This occurs only on a system startup.
- The VRRP holddown timer does not run if the amount of time passed since VRRP virtual router initialization is greater than preset holddown time. In such a case, VRRP virtual router transitions to Master happens irrespective of the holddown timer.

The VRRP holddown timer also applies to the VRRP BackupMaster feature.

VRRP and SMLT

The standard implementation of VRRP supports only one active master device for each IP subnet. All other VRRP interfaces in a network are in backup mode.

A deficiency occurs when VRRP-enabled switches use Split MultiLink Trunking (SMLT). If VRRP switches are aggregated into two Split MultiLink Trunk switches, the end host traffic is load-shared on all uplinks to the aggregation switches (based on the Multilink Trunk traffic distribution algorithm).

However, VRRP usually has only one active routing interface enabled. All other VRRP routers are in backup mode. Therefore, all traffic that reaches the backup VRRP router is forwarded over the vIST towards the master VRRP router. In this case, the vIST does not have enough bandwidth to carry all the aggregated traffic.

To resolve this issue, assign the backup router as the backup master router. The backup master router can actively load-share the routing traffic with a master router.

When the backup master router is enabled, the incoming host traffic is forwarded over the SMLT links as usual. When the backup master router is configured along with the critical IP interface and the critical IP interface goes down, the VRRP router transitions to be the backup router with the backup master state down. In this state, the VRRP router does not forward traffic.

VRRP and Spanning Tree

The switch can use one of two spanning tree protocols: Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP).

VRRP protects clients and servers from link or aggregation switch failures. Configure the network to limit the amount of time a link is out of service during VRRP convergence.

The following figure shows two possible configurations of VRRP and spanning tree; configuration A is optimal and configuration B is not.




Figure 255: VRRP and STG configurations

In this figure, configuration A is optimal because VRRP convergence occurs within 2 to 3 seconds. In configuration A, three spanning tree instances exist and VRRP runs on the link between the two routers. Spanning tree instance 2 exists on the link between the two routers, which separates the link between the two routers from the spanning tree instances found on the other devices. All uplinks are active.

In configuration B, VRRP convergence takes between 30 and 45 seconds because it depends on spanning tree convergence. After initial convergence, spanning tree blocks one link (an uplink), so only one uplink is used. If an error occurs on the uplink, spanning tree reconverges, which can take up to 45 seconds. After spanning tree reconvergence, VRRP can take a few more seconds to fail over.

VRRP Guidelines

The following list provides guidelines to use VRRP:

- If you configure VRRP and IP routing protocols on the same IP physical interface, you cannot select the interface address as the VRRP virtual IP address (logical IP address). You must use a separate dedicated IP address for VRRP.
- After you create a VRRP interface, you can add Secondary IP Interfaces, which are the equivalent of IPv6 VRRP global addresses. Secondary virtual IP addresses must be in the same subnet as a Secondary IP Interface on the VLAN.
- You can replace a primary virtual IP only if no secondary virtual IPs exist on the VRRP interface.
- Do not configure the virtual address as a physical interface that is used on the routing switches. Instead, use a third address, for example:
 - Interface IP address of VLAN A on Switch 1 = x.x.x.2
 - Interface IP address of VLAN A on Switch 2 = x.x.x.3
 - Virtual IP address of VLAN A = x.x.x.1



Note

The VRRP virtual IP address can be the same as the local IP address or Primary IP address of the interface on which VRRP is enabled. In this case, the router is the VRRP address owner on that interface and the priority is configured to the maximum value of 255.

- Configure the VRRP holddown timer with enough time that the Interior Gateway Protocol (IGP) routing protocol has time to update the routing table. In some cases, configuring the VRRP holddown timer to a minimum of 1.5 times the IGP convergence time is sufficient. For OSPF, as a best practice, use a value of 90 seconds if you use the default OSPF timers.
- Implement VRRP BackupMaster for an active-active configuration (BackupMaster works across multiple switches that participate in the same VRRP domain).
- Configure VRRP priority as 200 to configure VRRP Master.
- Stagger VRRP Masters between switches in the core to balance the load between switches.
- If you implement VRRP Fast, you create additional control traffic on the network and also create a greater load on the CPU. To reduce the convergence time of VRRP, the VRRP Fast feature allows the modification of VRRP timers to achieve subsecond failover of VRRP. Without VRRP Fast, normal convergence time is approximately 3 seconds.
- When you use the fast advertisement interval option to configure a master and backup device, you must enable the fast advertisement interval option on both systems for VRRP to work correctly. If you configure one device with the regular

advertisement interval, and the other device with the fast advertisement interval, it causes an unstable state and drops advertisements.

- Do not use VRRP BackupMaster and critical IP at the same time. Use one or the other.

VRRP for IPv6

For IPv6 hosts on a LAN to learn about one or more default routers, IPv6-enabled routers send router advertisements using the IPv6 ND protocol. The routers multicast these router advertisements every few minutes.

The ND protocol uses a mechanism called neighbor unreachability detection to detect the failure of a neighbor node (router or host) or the failure of the forwarding path to a neighbor. Nodes can monitor the status of a forwarding path by sending unicast ND neighbor solicitation messages to the neighbor node. To reduce traffic, nodes only send neighbor solicitations to neighbors to which they actively send traffic and only after the node receives no positive indication that the neighbors are up for a period of time. A host takes a minimum of 5 seconds to learn that a router is unreachable before it switches to another default router, but this minimum value increases ND traffic. This delay can cause service disruption.

VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol. With VRRP for IPv6, a backup router can take over for a failed default router in approximately three seconds (using default parameters). The switchover is accomplished without interaction with the hosts and with a minimum amount of VRRP traffic.

The IPv6 VRRP implementation is similar to the existing IPv4 VRRP operation, including support for holddown timer, critical IP, fast advertisements, and backup master. With backup master enabled, the backup switch routes all traffic according to its routing table. The backup master switch does not perform Layer 2 switching for the traffic to the VRRP master.

You must specify a link-local address to associate with the virtual router. Optionally, you can also assign global unicast IPv6 addresses to associate with the virtual router. Network prefixes for the virtual router are derived from the global IPv6 addresses assigned to the virtual router.

One active master switch exists for each IPv6 network prefix. All other VRRP interfaces in a network are in backup mode.

VRRP for IPv6 Operation

VRRP uses a virtual IP address shared between two or more routers connecting the common network prefix to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides dynamic default gateway redundancy in the event of failover.

The VRRP router with higher priority is called the master router. In case of equal priority the router with higher link-local address becomes the master router. The master router forwards packets sent to the virtual router IP addresses.

The following figure shows the minimum VRRP topology.




Figure 256: VRRP network topology

Traffic flows between users A and users B.

Router A uses VRRP global addresses as next hops for users B, and Router B for users A.

The VRRP master forwards the traffic and sends VRRP advertisements in the VLAN to announce to the backups that it is the master. If the master is no longer available, the backup takes over and becomes master. The only change occurs to the state of VRRP.

The VRRP router then transitions to the controlling state.



Note

The VRRP virtual IP address can be the same as the local IP address or Primary IP address of the interface on which VRRP is enabled. In this case, the router is the VRRP address owner on that interface and the priority is set to the maximum value of 255.

In the controlling state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. The router responds to ND neighbor solicitation and ND router solicitation messages for these IP addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts packets addressed to IP addresses associated with the virtual router.

If you initialize the VRRP router and the priority is not 255, the router transitions to the backup state to ensure that all Layer 2 switches in the downstream path relearn the new origin of the VRRP MAC addresses.

In the backup state, a VRRP router monitors the availability and state of the master router. The backup does not respond to ND neighbor solicitation and ND router

solicitation messages for virtual router IP addresses and discards packets with a MAC address equal to the virtual router MAC address. The backup does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, it transitions back to the initialize state. If the master router goes down, the backup router sends the VRRP advertisement and unsolicited ND neighbor advertisements and ND router advertisements described in the preceding paragraphs and transitions to the controlling state.

VRRP Advertisements and Master Router Failover

When you initialize a VRRP router, the master router continues to send advertisement messages at the advertisement interval period.

The other VRRP routers transition to the backup state in the following situations:

- if the priority in the received advertisement is greater than the local priority
- if the priority in the received advertisement is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address

The backup routers use the advertisements from the master router as a keepalive to monitor the status of the master router. If the backup router does not receive an advertisement during the master downtime interval, calculated as $3 * \text{advertisement interval}$, then the master router is declared down.

If a shutdown occurs, the master router sends a VRRP advertisement with a priority of 0 and transitions to the initialize state.

The priority value 0 indicates that the master router has stopped participating in VRRP. This value triggers the backup router to transition to the master state without waiting for the current master to time out.

Edge devices learn the VRRP MAC address towards the master. After the backup becomes master and sends the advertisements, this causes the VRRP MAC address to move on the edge devices and point towards the new master.

Critical IPv6 Address and Holddown Timer

The critical IPv6 address is an interface that has primary impact on VRRP. If you enable critical IPv6 and the status of the critical IP changes, the master and backup relationship also changes.

If you configure and enable critical IPv6 address, the master transitions to backup if the critical IPv6 is down, and the backup becomes the master. After the critical IPv6 address of the original master resumes, if the hold-down timer is configured to 0, it becomes the master immediately. Otherwise, the original master transitions to the master state after the hold-down timer time out.

You can specify the local router IP interface uplink from the VRRP router to the network as the critical IP address. This ensures that, if the local uplink interface fails, VRRP initiates a master router failover to one of the backup routers.

The critical address can be one of the global unicast IPv6 addresses assigned to any local IPv6 interfaces.

The holddown timer is a proprietary enhancement to VRRP.

After a master transitions to backup by critical IP changing, one of the backup routers will be elected as the master router. After the critical IPv6 of the original master is restored, the original master remains in the backup state for a period of time that you configure by using the **holddown-timer** parameter. The router becomes the master immediately if you use the command **ipv6 vrrp <1-255> action preempt**.

The holddown timer allows the master router enough time to detect and update the dynamic routes. The timer delays the preemption of the master over the backup, when the master becomes available. If the hold-timer is configured to 0, it becomes the master router immediately. Otherwise, it transitions to the master state only after the holddown timer times out.

The holddown timer does not apply during failovers caused by VRRP router priority change. The holddown timer applies only to failovers caused by a critical IP failure.

Configure all of your routers to use identical values for the holddown timer.



Important

Do not use VRRP backup master and critical IP at the same time. Use one or the other. The critical IP address must be a local address.

VRRP Backup Master with Triangular SMLT

The standard implementation of VRRP supports one active master switch for each IPv6 subnet. All other VRRP interfaces in a network are in backup mode.

A deficiency occurs when VRRP-enabled switches use SMLT. If VRRP switches are aggregated into two SMLT switches, the end host traffic is load-shared on all uplinks to the aggregation switches (based on the Multilink Trunk [MLT] traffic distribution algorithm).

However, VRRP usually has only one active routing interface enabled. All other VRRP routers are in backup mode. Therefore, all traffic that reaches the backup VRRP router is forwarded over Virtual Inter-Switch Trunk (vIST) toward the master VRRP router. In this case, vIST potentially does not have enough bandwidth to carry all the aggregated traffic.

To resolve this issue, assign the backup router as the backup master router. The backup master router can actively load-share the routing traffic with a master router.

Because the two VRRP peer nodes exchange MAC address tables, the VRRP backup master can forward traffic directly, on behalf of the master router. The switch in the backup master state routes all traffic received on the backup master IP interface according to its routing table. The backup master switch does not perform Layer 2 switching for the traffic to the VRRP master.

If you enable SMLT on the backup master router, the incoming host traffic is forwarded over the SMLT links as usual.

**Important**

Do not use VRRP backup master and critical IP at the same time. Use one or the other.

Fast Advertisement

You can configure the advertisement time interval (in seconds) between sending advertisement messages. This interval permits fast network convergence with standardized VRRP failover. However, losing connections to servers for more than a second can result in missing critical failures. Customer network uptime in many cases requires faster network convergence, which means network problems must be detected within hundreds of milliseconds.

To meet these requirements, the fast advertisement interval are provided.

The fast advertisement interval is similar to the advertisement interval parameter except for the unit of measure and the range. The fast advertisement interval is expressed in milliseconds and the range is from 200 to 1,000 milliseconds. This unit of measure must be in multiples of 200 milliseconds.

To configure fast advertisement, you must specify a fast advertisement interval and explicitly enable the fast advertisement option. After you enable fast advertisement, the fast advertisement interval is used instead of the advertisement interval.

If you enable fast advertisement, VRRP can only communicate with other products that have the same configuration.

Accept-mode

When you configure VRRP for IPv6 on an interface you can configure the `accept-mode` parameter, which controls whether the VRRP master or backup master accepts packets destined for the IPv6 address associated with the virtual router.

By default, `accept-mode` is disabled. The `accept-mode` parameter does not affect the Neighbor Discovery packets. The master router forwards packets with a destination link-layer MAC address that matches the virtual MAC address, and accepts packets forwarded over the virtual interswitch trunk (vIST) toward the master router, if `accept-mode` is enabled. If you disable `accept-mode`, you cannot ping the virtual IPv6 address. If you enable `accept-mode`, the master router accepts packets addressed to the IPv6 address that is associated with the virtual router.

When you configure VRRP for IPv6 on an interface, you can configure the `accept-mode` parameter. By default, `accept-mode` is disabled. If you disable `accept-mode`, the master router does not drop neighbor solicitations or neighbor advertisements. The master router forwards packets with a destination link-layer MAC address that matches the virtual MAC address. If you disable `accept-mode`, you cannot ping the virtual IPv6 address.

VRRPv3

VRRPv3 is a combined protocol for both IPv4 and IPv6. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router that controls the IPv4 or IPv6 addresses associated with a virtual router is called the Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Backups wait for a Master and take ownership when the Master is no longer detected.

The election protocol provides dynamic failover in the forwarding responsibility when the Master is unavailable. VRRP for IPv4 gains a higher-availability default path without configuring dynamic routing or router discovery protocols on every end-host. VRRP for IPv6 gains a quick switch-over to Backup routers compared to the standard IPv6 Neighbor Discovery mechanisms.



Note

The VRRP IPv6 link-local address must be the same for all VRRP routers sharing the same link and the same virtual router ID, that is, the same VRRP instance. It is the address the VRRP advertisements are sent from. Also, the Router Advertisement packets from the VRRP interface are transmitted using this address, so, when IPv6 Stateless Address Autoconfiguration is used, this address is added to host Default Router List and is used as a gateway.

The software supports VRRPv3 for IPv4 and VRRPv3 for IPv6. VRRPv3 for IPv6 is compliant to RFC 5798. The software also supports VRRPv2 for IPv4.

VRRPv3 guidelines

The switch also supports VRRPv2 for IPv4. If you configure VRRP IPv6 on an interface, it runs independently of the IPv4 version. Configure the version of the VRRP IPv4 on the interface before you configure any other IPv4 VRRP attributes. By default, the version is not configured to a particular value. However, when sourcing older configuration files that do not have the version saved, the router configures the version to VRRPv2 by default. If you change the version, all IPv4 configuration under that interface is automatically removed, and you are prompted for a confirmation before this operation.

Perform the CLI configuration through **ip vrrp** or **ipv6 vrrp** nodes; CLI commands for IPv4 are common for version 2 and version 3.

The following list identifies the features that make both IPv4 and IPv6 VRRPv3 features compliant to RFC 5798:

- Advertisement vs Fast-advertisement — Prior to RFC 5798, the minimum advertisement interval was 1 second, with Fast-advertisement a sub-second interval could be configured. When this feature is enabled, the VRRP ADVERTISEMENT packets are sent with type 7 instead of 1. With RFC 5798 the sub-second interval is standardised, and the switch sends all packets for VRRPv3 with type 1. The use

of Fast-advertisement remains the same. VRRPv2 packets send with type 7, if Fast-advertisement is enabled.

- Add Master-advertisement-interval — Prior to RFC 5798 compliance, all virtual routers on the same VLAN had the same Advertisement-Interval configured. RFC 5798 states that you can use different Advertisement Intervals on the Master and Backup. On the Master, the Master-advertisement-interval and the Advertisement-Interval have the same value. On the Backup, the Master-advertisement-interval is used to calculate the timers, and the locally configured Advertisement-Interval is ignored until the Backup transitions to Master. The Master-advertisement-interval value is put in the advertisement packet type sent by the Master
- Transition to master as specified in RFC 5798 — Prior to RFC 5798, if a Backup receives an advertisement with a lower priority (or same priority but lower IP), it immediately sends its own advertisement and transitions to Master. However, RFC 5798 states that such packets must be discarded, which means it will transition to Master after the Master_Down_Timer expires
- Add skew-time — RFC 5798 states that skew-time is calculated depending on the priority, and Master-advertisement-interval assures that the Backup with highest priority sends the first advertisement when the Master goes down

Skew time is calculated using the formula: $((256 - \text{priority}) * \text{Master_Adver_Interval}) / 256$.

- Add preempt-mode — Preempt-mode is different from the **ipv6 vrrp <vrid> action preempt** command, which is an operational command issued when you want to stop the hold-down timer. RFC 5798 states that preempt-mode should be set to false when you do not want a higher priority Backup to transition to Master. By default, it is set to true



Note

Accept-mode is not fully implemented for IPv4 VRRPv3. You can only ping the virtual IP address, the same way as it is for IPv4 VRRPv2.

Configure VRRP Using CLI

This section provides procedures to configure VRRP using Command Line Interface (CLI).

Before You Begin

- Assign an IP address to the interface.
- Enable VRRP globally.

Configure the VRRP Version on an Interface

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the VRRP version:

```
ip vrrp version <2-3>
```

Example

Configure VRRP version 3 for the specified interface:

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)# interface gigabitethernet 1/2  
Switch:1(config-if)# ip vrrp version 3
```

Variable definitions

Use the data in the following table to use the **ip vrrp version** command.

Variable	Value
<i>version</i> <2-3>	Configures the VRRP version (2 or 3) on the specified interface

Configure VRRP on a Port or a VLAN

Configure VRRP on a port or a VLAN to forward packets to the virtual IP addresses associated with the virtual router and customize the VRRP configuration.

Before You Begin

- Create primary and, optionally, Secondary IP Interfaces on the VLAN.
- Configure the VRRP version on the interface.

About This Task

You can create secondary virtual IP addresses for VRRP instances on VLAN interfaces using the same **ip vrrp address** command included in this procedure. The first address you configure is the primary virtual IP address. Secondary virtual IP addresses must be in the same subnet as a Secondary IP Interface on the VLAN.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure a primary virtual address and create the VRRP instance:

```
ip vrrp address <1-255> {A.B.C.D}
```

3. (Optional) Create a secondary virtual IP address:



Note

This step does not apply to brouter ports.

```
ip vrrp address <1-255> {A.B.C.D}
```

4. Enable the VRRP instance:

```
ip vrrp <1-255> enable
```

5. Show the global VRRP configuration:

```
show ip vrrp
```

Example

```
Switch:1>enable  
Switch:1#configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch:1(config)#interface gigabitethernet 1/2
```

Configure a primary VRRP address:

```
Switch:1(config-if)#ip vrrp address 28 192.0.2.1
```

Enable VRRP:

```
Switch:1(config-if)#ip vrrp 28 enable
```

Variable Definitions

Use the data in the following table to use the **ip vrrp** command.

Variable	Value
<i>1-255</i>	Specifies the number of the VRRP to create or modify.
<i>action {none preempt}</i>	Specifies the action of the VRRP interface: <ul style="list-style-type: none"> • none—do not override the holddown timer • preempt—override the holddown timer <p>Note: You can use this parameter only if the holddown timer is active.</p>
<i>address <1-255> {A.B.C.D}</i>	Configures the IP address of the VRRP physical interface that forwards packets to the virtual IP addresses associated with the virtual router. <i>A.B.C.D</i> is the IP address of the master VRRP.
<i>adver-int <1-255></i>	Configures the the time interval (in seconds) between sending VRRP advertisement messages. The default is 1. This value must be the same on all participating routers.
<i>backup-master enable</i>	Enables the VRRP backup master. When backup master functionality is enabled, the VRRP router forwards IP packets destined to the VRRP MAC even when the router is not the VRRP Master. Important: Do not enable backup master if you enable critical IP.
<i>critical-ip-addr <A.B.C.D></i>	Configures the critical IP address for VRRP. <i>A.B.C.D</i> is the IP address on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding. Note: In this context, <i>/local/</i> implies an address from the same VRF as the IP interface where VRRP is being configured.
<i>critical-ip enable</i>	Enables the critical IP address option. Important: Do not enable Critical IP if backup master is enabled.
<i>enable</i>	Enables VRRP on the port.
<i>fast-adv enable</i>	Enables the Fast Advertisement Interval. The default is disabled.

Variable	Value
<i>fast-adv-int</i> <200-1000>	Configures the Fast Advertisement Interval, the time interval between sending VRRP advertisement messages. <i>200-1000</i> is the range in milliseconds, and must be the same on all participating routers. The default is 200. You must enter values in multiples of 200 milliseconds.
<i>holddown-timer</i> <0-21600>	Specifies the time interval (in seconds) to wait before preempting the current VRRP primary router. The default is 0.
<i>preempt-mode</i>	Enables preempt mode for VRRP (v3).
<i>priority</i> <1-255>	Configures the port VRRP priority. <i>1-255</i> is the value used by the VRRP router. The default is 100. Assign the value 255 to the router that owns the IP address associated with the virtual router.

Display Information for all Configured VRRP Addresses

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display information about all addresses or filter the output based on a specified address, VRF name or ID, or VRRP ID:

```
show ip vrrp address [addr <A.B.C.D>] [vrf WORD<1-16>] [vrfids WORD<0-512>] [vrid <1-255>]
```

3. Display information for a specific VRRP version:

```
show ip vrrp address version <2-3>
```

4. Display version-based VRRP configuration for the specified VRF name:

```
show ip vrrp address vrf WORD<1-16> version <2-3>
```

5. Display version-based VRRP configuration for the specified VRF ID:

```
show ip vrrp address vrfids WORD<0-512> version <2-3>
```

Example

```
Switch:1>show ip vrrp address
```

```
=====
VRRP Info - GlobalRouter
=====
```

VRRP ID	P/V	IP	MAC	STATE	CONTROL	PRIO	ADV	VERSION	ORIGIN
10	2	203.0.113.2	00:00:5e:00:01:0a	Init	Disabled	255	1	2	config
1	222	11.1.1.1	00:00:5e:00:01:01	Init	Disabled	255	1	2	config
1	223	15.1.1.1	00:00:5e:00:01:01	Init	Disabled	255	1	3	config

3 out of 3 Total Num of VRRP Address Entries displayed.

VRRP ID	P/V	MASTER	UP TIME	HLD DWN	CRITICAL	IP (ENABLED)	VERSION	HOSTNAME
10	2	0.0.0.0	0 day(s), 00:00:00	0	0.0.0.0	(No)	2	

```

1      222      0.0.0.0      0 day(s), 00:00:00      0      0.0.0.0      (No)  2
1      223      0.0.0.0      0 day(s), 00:00:00      0      0.0.0.0      (No)  3

```

3 out of 3 Total Num of VRRP Address Entries displayed.

VRRP ID	P/V	BACKUP MASTER	BACKUP MASTER STATE	FAST ADV (ENABLED)	VERSION
10	2	disable	down	200 (NO)	2
1	222	disable	down	200 (NO)	2
1	223	disable	down	200 (NO)	3

3 out of 3 Total Num of VRRP Address Entries displayed.

VRRP ID	P/V	VIRTUAL ADDRESS
1	222	11.1.2.1
1	223	15.1.2.1

2 out of 2 Total Num of VRRP Virtual Address Entries displayed.

Variable Definitions

Use the data in the following table to use the **show ip vrrp address** command.

Variable	Value
<i>addr</i> <A.B.C.D>	Specifies the physical local address of the master VRRP.
<i>vrf</i> WORD<1-16>	Specifies the name of the VRF.
<i>vrid</i> <1-255>	Specifies a unique integer value that represents the virtual router ID. The virtual router acts as the default router for one or more assigned addresses.
<i>vrfids</i> WORD<0-512>	Specifies the ID of the VRF.
<i>version</i> <2-3>	Specifies the VRRP version to include in the command output.

Display Extended VLAN VRRP

Perform this procedure to display the extended VRRP configuration for all VLANs or a specified VLAN on the device.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the extended VRRP configuration for all VLANs on the device or for the specified VLAN:

```

show ip vrrp interface vlan [<1-4059>] [portList] verbose [vrf
WORD<1-16>] [vrfids WORD<0-512>]

```

Example

```

Switch:1>show ip vrrp interface vlan
=====

```

Vlan Vrrp						
VLAN ID	VRF NAME	VRRP ID	IP ADDRESS	VIRTUAL MAC ADDRESS	VERSION	ORIGIN
2	GlobalRouter	10	203.0.113.2	00:00:5e:00:01:0a	2	config
222	GlobalRouter	1	192.0.2.1	00:00:5e:00:01:01	2	config
223	GlobalRouter	1	198.51.100.1	00:00:5e:00:01:01	3	config

Variable Definitions

The following table defines parameters for the **show ip vrrp interface vlan** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<i>portList</i>	Specifies the slot or port number of a range of ports.
<i>vrf WORD</i> <1-16>	Specifies the name of the VRF.
<i>vrfids WORD</i> <0-512>	Specifies the ID of the VRF.

Display VRRP Interface Information

About This Task

If you enter a virtual router ID or an IP address when showing VRRP interface information, the system displays the information only for that virtual router ID or for that interface.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display VRRPv2 information about the specified interface:
show ip vrrp interface version <2-3>
3. Display additional VRRPv2 information about the specified interface:
show ip vrrp interface verbose version <2-3>
4. Display VRRPv2 information for the specified VRF:
show ip vrrp interface vrf WORD<1-16> version <2-3>
5. Display VRRPv2 information for the specified virtual router:
show ip vrrp interface vrfids WORD<0-512> [version <2-3>]

Examples

```
Switch:1#show ip vrrp interface
```

```
=====
```

Vlan Vrrp					
=====					
VLAN ID	VRF NAME	VRRP ID	IP ADDRESS	VIRTUAL MAC ADDRESS	VERSION ORIGIN

10	GlobalRouter	37	10.10.0.254	00:00:5e:00:01:25	2 dvr
400	GlobalRouter	1	40.40.0.254	00:00:5e:00:01:01	2 config

```
All 2 out of 2 Total Num of Vlan Vrrp displayed
```

```
=====
```

Port Vrrp				
=====				
PORT NUM	VRF NAME	VRRP ID	IP ADDRESS	VIRTUAL MAC ADDRESS VERSION

```
Switch:1#show ip vrrp interface verbose
```

```
=====
```

Vlan Vrrp Extended										
=====										
VLAN ID	VRRP ID	VRF NAME	STATE	CONTROL	PRIORITY	MASTER IPADDR	ADVERTISE INTERVAL	CRITICAL IPADDR	VERSION	HOSTNAME ORIGIN

10	37	Global~	backup	enable	100	0.0.0.0	1	0.0.0.0	2	dvr
400	1	Global~	master	enable	100	40.40.0.12	1	0.0.0.0	2	config

```
All 2 out of 2 Total Num of Vlan Vrrp Extended Entries displayed
```

VLAN ID	VRRP ID	VRF NAME	HOLDDWN TIME	ACTION	CRITICAL IP ENABLE	BACKUP MASTER STATE	BACKUP MASTER STATE	FAST ADV INTERVAL	FAST ADV ENABLE	VERSION

10	37	GlobalRouter	60	none	disable	disable	down	200	disable	2
400	1	GlobalRouter	0	none	disable	disable	down	200	disable	2

```
All 2 out of 2 Vlan Vrrp Extended Entries displayed
```

VLAN ID	VRRP ID	VRF NAME	MASTER ADV INTERVAL (ms)	PREEMPT MODE	PSEUDO-HEADER CHECKSUM	VERSION

10	37	GlobalRouter	N/A	enabled	disabled	2
400	1	GlobalRouter	N/A	enabled	disabled	2

```
All 2 out of 2 Vlan Vrrp Extended Entries displayed
```

```
=====
```

Port Vrrp Extended										
=====										
PORT NUM	VRRP ID	VRF NAME	STATE	CONTROL	PRIORITY	MASTER IPADDR	ADVERTISE INTERVAL	CRITICAL IPADDR	VERSION	

PORT NUM	VRRP ID	VRF NAME	HOLDDWN TIME	ACTION	CRITICAL IP ENABLE	BACKUP MASTER STATE	BACKUP MASTER STATE	FAST ADV INTERVAL	FAST ADV ENABLE	VERSION

PORT NUM	VRRP ID	VRF NAME	MASTER ADV INTERVAL (ms)	PREEMPT MODE	PSEUDO-HEADER CHECKSUM	VERSION				

Variable definitions

Use the data in the following table to use the **show ip vrrp interface** command.

Variable	Value
<i>gigabitethernet</i> { <i>slot/port</i> [- <i>slot/port</i>][,...]}	Specifies to show the VRRP information of which interface.
<i>verbose</i>	Specifies to show all available information about the VRRP interfaces.
<i>vlan</i>	Specifies the VLAN that contains the VRRP.
<i>vrf</i> <i>WORD</i> <1-16>	Specifies the name of the VRF.
<i>vrid</i> <1-255>	Specifies a unique integer value that represents the virtual router ID in the range 1–255. The virtual router acts as the default router for one or more assigned addresses.
<i>vrfids</i> <i>WORD</i> <0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
<i>version</i> <2-3>	Specifies the VRRP version (2 or 3) configured.

View IP VRRPv3 Statistics

Use the following procedure to view IP VRRPv3 statistics to monitor network performance.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Enter the following command to view VRRP statistics:
show ip vrrp statistics version <2-3>
3. Enter the following command to view VRRP statistics for the specified VRF:
show ip vrrp statistics vrf *WORD*<1-16> version <2-3>
4. Enter the following command to view VRRP statistics for the specified virtual router:
show ip vrrp statistics vrfids *WORD*<0-512> version <2-3>

Example

View IP VRRPv3 statistics:

```
Switch:1#show ip vrrp statistics
```

```
=====
                        VRRP Global Stats - GlobalRouter
=====

CHK_SUM_ERR   VERSION_ERR   VRID_ERR      VRRP_VERSION
-----
0              0              0              2
0              0              0              3
```

```
=====
VRRP Interface Stats - GlobalRouter
=====
```

VRRP ID	P/V	BECOME_MASTER	ADVERTISE_RCV	VERSION
3	3	1	0	2
2	1/1	1	0	3

VRRP ID	P/V	ADVERTISE_INT_ERR	TTL_ERR	PRIOR_0_RCV	VERSION
3	3	0	0	0	2
2	1/1	0	0	0	3

VRRP ID	P/V	PRIOR_0_SENT	INVALID_TYPE_ERR	ADDRESS_LIST_ERR	UNKNOWN_AUTHTYPE	VERSION
3	3	0	0	0	0	2
2	1/1	0	0	0	0	3

VRRP ID	P/V	AUTHTYPE_ERR	PACKLEN_ERR	VERSION
3	3	0	0	2
2	1/1	0	0	3

Variable Definitions

Use the data in the following table to use the **ip vrrp version** command.

Variable	Value
<i>version</i> <2-3>	Specifies the version of VRRP (2 or 3) to be configured on the specified interface.
<i>vrf</i> WORD<1-16>	Specifies the name of the VRF.
<i>vrfids</i> WORD<0-512>	Specifies the ID of the VRF.

Enable Ping to a Virtual IP Address

Use the following procedure to enable ping to a virtual IP address. The default is enabled.

Procedure

- Enter VRRP Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router vrrp
```
- Enable ping to a virtual IP address:

```
ping-virtual-address enable [vrf WORD<1-16>]
```

```
default ping-virtual-address enable [vrf WORD<1-16>]
```
- Display the configuration:

```
show ip vrrp [vrf WORD<1-16>]
```


Variable definitions

Use the data in the following table to use the `ping-virtual-address enable` and `show ip vrrp` commands.

Variable	Value
<code>enable</code>	Enables ping to a virtual IP address.
<code>vrf WORD<1-16></code>	Specifies the VRF.

Configure VRRP Notification Control

Use the following procedure to enable VRRP notification control. The generation of SNMP traps for VRRP events is enabled, by default.

About This Task

You can configure traps by creating SNMPv3 trap notifications, creating a target address to send the notifications, and specify target parameters. For more information about how to configure trap notifications, see [Logs and Traps](#) on page 2535.

Procedure

1. Enter VRRP Router Configuration mode:
`enable`
`configure terminal`
`router vrrp`
2. Enable a trap for VRRP events:
`send-trap enable [vrf WORD<1-16>]`
3. Disable a trap for VRRP events:
`no send-trap enable [vrf WORD<1-16>]`
4. Configure a trap for VRRP events to the default:
`default send-trap enable [vrf WORD<1-16>]`
5. Display the configuration:
`show ip vrrp [vrf WORD<1-16>]`

Variable definitions

Use the data in the following table to use the **`send-trap`** and **`show ip vrrp`** commands.

Variable	Value
<code>enable</code>	Enables generation of SNMP traps.
<code>vrf WORD<1-16></code>	Configures the send-trap for a particular VRF.

Enable IPv4 VRRP Preempt Mode

You can configure VRRP to preempt the existing router. If a new VRRP router is added to the network with a higher priority than the existing routers, then the new router becomes the master. If preempt-mode is disabled, then the new router does not become a master, it transitions to master only when the current master is down, that is when it does not receive any advertisement packets from the current master. By default, preempt-mode is enabled.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]} [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enter the following command:

```
ip vrrp <vrid> preempt-mode enable
```

Example

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)# interface gigabitethernet 1/2
```

Enabling preempt-mode on interface 1/2:

```
Switch:1(config-if)# ip vrrp 1 preempt-mode enable
```

Variable definitions

Use the data in the following table to use the **ip vrrp <vrid>** command.

Variable	Value
<i>preempt-mode enable</i>	Enables preempt-mode for VRRPv3 for IPv4.

IPv6 VRRP Configuration using CLI

This section provides procedures to configure IPv6 VRRP using the command line interface (CLI).

Configure the VRRP Interface

Configure VRRP to provide fast failover of a default router for IPv6 LAN hosts, in order to create a VRRP instance.

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Do not configure RSMLT on the VLAN.
- You must specify a link-local address to associate with the virtual router.

About This Task

VRRP supports a virtual IPv6 address shared between two or more routers that connect the common subnet to the enterprise network.

VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol.

Perform this procedure to also configure the additional addresses for which the virtual router acts as a backup.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]]  
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Associate an address with the virtual router for either link-local or global:

- `ipv6 vrrp address <1-255> link-local WORD <0-127>`
- `ipv6 vrrp address <1-255> global WORD <0-255>`



Note

You must configure the link-local address before you configure the global address.

3. Enable VRRP for the interface:

```
ipv6 vrrp <1-255> enable
```

Example

Associate a link-local address with the virtual router ID 12:

```
Switch:1(config-if)#ipv6 vrrp address 12 link-local fe80::1234
```

Associate a global address with the virtual router ID 12

```
Switch:1(config-if)#ipv6 vrrp address 12 global 3333::1234/64
```

Enable VRRP for the interface:

```
Switch:1(config-if)#ipv6 vrrp 12 enable
```

Variable Definitions

Use the data in the following table to use the **ipv6 vrrp address** command.

Variable	Value
<1-255>	Specifies the virtual router ID. The virtual router acts as the default router for one or more associated addresses.
<i>enable</i>	Enables IPv6 VRRP. The default is disabled.
<i>global WORD</i> <0-255>	Specifies a global IPv6 address and mask to associate with the virtual router.
<i>link-local WORD</i> <0-127>	Specifies a link-local IPv6 address to associate with the virtual router.

Use the data in the following table to use the **interface** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the <i>vrf-scaling</i> and <i>spbm-config-mode</i> boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{ <i>slot/port</i> [/ <i>sub-port</i>][<i>-slot/port</i> [/ <i>sub-port</i>]] [<i>,...</i>][<i>slot/all</i>][<i>all</i>]}	Identifies the slot and port in one of the following formats: <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

View VRRP Information

Display VRRP port or VLAN information to verify your configuration. Show VRRP information by IPv6 address or virtual router ID. If you enter a virtual router ID or an IPv6 address when you view VRRP information, the information applies only to that virtual router ID or for that interface.

Procedure

1. Enter Privileged EXEC mode:
enable
2. View the configuration information for all interfaces:
show ipv6 vrrp interface [verbose] [vrf *WORD*<1-16> | vrfids *WORD*<0-512>]

3. View the configuration information for one or more ports:

```
show ipv6 vrrp interface gigabitethernet {slot/port[/sub-port]}[-slot/
port[/sub-port]][[,...][slot/all][all]] [verbose] [vrf WORD<1-16> |
vrfids WORD<0-512>]
```

4. View the configuration information for one or more VLANs:

```
show ipv6 vrrp interface vlan [<1-4059>] [verbose] [vrf WORD<1-16> |
vrfids WORD<0-512>]
```

5. View the configuration information for one or more virtual router IDs:

```
show ipv6 vrrp interface vrid <1-255> [verbose] [vrf WORD<1-16> |
vrfids WORD<0-512>]
```

6. View VRRP address information:

```
show ipv6 vrrp address [vrf WORD<1-16> | vrfids WORD<0-512>]
```

7. View VRRP address information for a link-local address:

```
show ipv6 vrrp address link-local WORD<0-127> [verbose] [vrf WORD<1-
16> | vrfids WORD<0-512>]
```

8. View VRRP address information for a virtual router ID:

```
show ipv6 vrrp address vrid <1-255> [vrf WORD<1-16> | vrfids WORD<0-
512>]
```

Example

```
Switch:1>show ipv6 vrrp address
```

```
=====
                        VRRP Info - GlobalRouter
=====
```

VRID	P/V	IP	MAC	STATE	CONTROL
12	1/1	fe80:0:0:0:0:0:0:1234	00:00:5e:00:02:0c	Init	Disabled

```
=====
```

VRID	P/V	MASTER	PRIOR	ADV	UP TIME
12	1/1	0:0:0:0:0:0:0:0	100	1	0 day(s), 00:00:00

```
=====
```

VRID	P/V	CRITICAL IP	CRITICAL IP ENABLED	ACCEPT MODE
12	1/1	0:0:0:0:0:0:0:0	No	disable

```
=====
```

VRID	P/V	BACKUP MASTER	BACKUP-MASTER STATE	FAST (ENABLED) ADV	ACTION	HLD DWN	REM
12	1/1	disable	down	400 (YES)	none	30	0

```
=====
```

VRID	P/V	GLOBAL ADDRESS
12	1/1	1111::2222/64

```
=====
```

Flags Legend:
HLD DWN: Configured hold-down timer value, REM: REMaining hold-down timer value

--More-- (q = quit)

```
Switch:1#show ipv6 vrrp interface verbose
```

```
=====
                        Vlan Vrrp for IPv6 Extended
=====
```

VLAN ID	VRF NAME	VRRP ID	STATE	CONTROL	PRIORITY	MASTER IPADDR
40	GlobalRouter	1	init	disable	100	0:0:0:0:0:0:0:0
40	GlobalRouter	2	init	disable	100	0:0:0:0:0:0:0:0

```
All 2 Vlan Vrrp Extended Entries out of 18 Total Num of Vrrp displayed
```

VLAN ID	VRF NAME	VRRP ID	HOLDDWN	ACTION TIME	CRITICAL IP	CRITICAL ENABLE	CRITICAL IPADDR
40	GlobalRouter	1	0	none	disable		0:0:0:0:0:0:0:0
40	GlobalRouter	2	0	none	disable		0:0:0:0:0:0:0:0

```
All 2 Vlan Vrrp Extended Entries out of 18 Total Num of Vrrp displayed
```

VLAN ID	VRF NAME	VRRP ID	BACKUP MASTER	BACKUP MASTER STATE	ADVERTISE INTERVAL (s)	FAST ADV INTERVAL (ms)	FAST ADV ENABLE	MASTER ADV INTERVAL (ms)	PREEMPT MODE
40	GlobalRouter	1	disable	down	1	200	disable	1000	enable
40	GlobalRouter	2	disable	down	1	200	disable	1000	enable

```
All 2 Vlan Vrrp Extended Entries out of 18 Total Num of Vrrp displayed
```

```
=====
                        Port Vrrp for IPv6 Extended
=====
```

PORT NUM	VRF NAME	VRRP ID	STATE	CONTROL	PRIORITY	MASTER IPADDR
1/23	GlobalRouter	1	init	disable	100	0:0:0:0:0:0:0:0
1/23	GlobalRouter	2	init	disable	100	0:0:0:0:0:0:0:0

```
All 2 Port Vrrp Extended Entries out of 18 Total Num of Vrrp displayed
```

PORT NUM	VRF NAME	VRRP ID	HOLDDWN	ACTION TIME	CRITICAL IP	CRITICAL ENABLE	CRITICAL IPADDR
1/23	GlobalRouter	1	0	none	disable		0:0:0:0:0:0:0:0
1/23	GlobalRouter	2	0	none	disable		0:0:0:0:0:0:0:0

```
All 2 Port Vrrp Extended Entries out of 18 Total Num of Vrrp displayed
```

PORT NUM	VRF NAME	VRRP ID	BACKUP MASTER	BACKUP MASTER STATE	ADVERTISE INTERVAL (s)	FAST ADV INTERVAL (ms)	FAST ADV ENABLE	MASTER ADV INTERVAL (ms)	PREEMPT MODE
1/23	GlobalRouter	1	disable	down	1	200	disable	1000	enable
1/23	GlobalRouter	2	disable	down	1	200	disable	1000	enable

```
All 2 Port Vrrp Extended Entries out of 18 Total Num of Vrrp displayed
```

Variable Definitions

Use the data in the following table to use the **show ipv6 vrrp** commands.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...][slot/all] [all]}</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>link-local WORD<0-127></code>	Displays information by link-local IPv6 address.
<code>verbose</code>	Displays extended information.
<code>vlan [<1-4059>]</code>	<p>Identifies the slot and port in one of the following formats:</p> <ul style="list-style-type: none"> • a single slot and port (slot/port) • a range of slots and ports (slot/port-slot/port) • a series of slots and ports (slot/port,slot/port,slot/port) • all ports on the same slot (slot/all) • all ports on the switch (all) <p>If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>
<code>vrid <1-255></code>	Displays information by virtual router ID.
<code>vrf WORD<1-16></code>	Specifies the VRF name.
<code>vrfids WORD<0-512></code>	Specifies the VRF ID.

Configure VRRP Notification Control

Perform this procedure to configure VRRP notification control.

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.

About This Task

By default, generation of SNMP traps for VRRP events is enabled.

Procedure

1. Enter VRRP Router Configuration mode:
`enable`

`configure terminal`

`router vrrp`
2. Enable the VRRP-router to generate SNMP traps for events:
`ipv6 send-trap enable`

Example

Disable generation of SNMP traps for VRRP events:

```
Switch:1(config-vrrp)#no ipv6 send-trap enable
```

Configure Additional VRRP Parameters for an Interface

Configure VRRP to provide fast failover of a default router for IPv6 LAN hosts. VRRP supports a virtual IPv6 address shared between two or more routers that connect the common subnet to the enterprise network. VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol.

Configure the parameters in this procedure if the default values do not meet your requirements.

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Do not configure RSMLT on the VLAN.

About This Task

A switch that acts as a VRRP master does not reply to SNMP get requests to the VRRP virtual interface address. The switch will, however, respond to SNMP get requests to the physical IP address.

Procedure

1. Enter Interface Configuration mode:
`enable`

`configure terminal`

`interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]`
`[,...][slot/all][all]} or interface vlan <1-4059>`



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure the accept mode of the master router:

```
ipv6 vrrp <1-255> accept-mode enable
```

3. Determine if the router overrides the holddown timer:

```
ipv6 vrrp <1-255> action <none|preempt>
```

4. Configure the interval between advertisement messages:

```
ipv6 vrrp <1-255> adver-int <1-40>
```

5. Enable the backup VRRP switch for traffic forwarding:

```
ipv6 vrrp <1-255> backup-master enable
```

6. Configure the IP interface on the local router:

```
ipv6 vrrp <1-255> critical-ipv6-addr WORD<0-46> [critical-ipv6 enable]
```

7. Configure the fast advertisement interval:

```
ipv6 vrrp <1-255> fast-adv enable [fast-adv-int <200-1000>]
```

8. Configure the holddown timer:

```
ipv6 vrrp <1-255> holddown-timer <0-21600>
```

9. Configure the priority for the VRRP router:

```
ipv6 vrrp <1-255> priority <1-255>
```

Example

Configure the fast advertisement interval:

```
Switch:1(config-if)#ipv6 vrrp 12 fast-adv enable fast-adv-int 400
```

Configure the holddown timer:

```
Switch:1(config-if)#ipv6 vrrp 12 holddown-timer 30
```

Variable Definitions

Use the data in the following table to use the **ipv6 vrrp** command.

Variable	Value
<1-255>	Specifies a number that uniquely identifies a virtual router on a VRRP router. The virtual router acts as the default router for one or more assigned addresses.
<i>accept-mode enable</i>	Controls whether the VRRP master or backup master accepts packets (other than neighbor discovery packets) destined to the IPv6 address associated with the virtual router. The default value is disable.
<i>action <none preempt></i>	Lists options to override the holddown timer manually and force preemption: <ul style="list-style-type: none"> • <i>none</i> does not override the timer. • <i>preempt</i> preempts the timer. This parameter applies only if the holddown timer is active.
<i>adver-int <1-40></i>	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second. Only the master router sends advertisements.

Variable	Value
<i>backup-master enable</i>	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the vIST. The default is disabled.
<i>critical-ipv6 enable</i>	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
<i>critical-ipv6-addr WORD<0-46></i>	Specifies an IP interface on the local router configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.
<i>fast-adv enable</i>	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disable.
<i>fast-adv-int <200-1000></i>	Configures the interval between VRRP advertisement messages. You must configure the same value on all participating routers. This unit of measure must be in multiples of 200 milliseconds. The default is 200.
<i>holddown-timer <0-21600></i>	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.
<i>priority <1-255></i>	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.

View IPv6 VRRP Statistics

View IPv6 VRRP statistics to monitor network performance.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. View statistics for the device and for all interfaces:

```
show ipv6 vrrp statistics [link-local WORD<0-127>]] [vrid <1-255>]
```

Example

View IPv6 VRRP statistics for VRID 1.

```
Switch:1(config)#show ipv6 vrrp statistics vrid 1

=====
VRRP Interface Stats - GlobalRouter
=====

VRID  P/V   BECOME_MASTER  ADVERTISE_RCV
-----
1      84     2              17372
1      85     2              17372
1      86     1              0
1      87     1              0
1      1001   2              17372

VRID  P/V   ADVERTISE_INT_ERR  TTL_ERR      PRIO_0_RCV
-----
```

```

1      84      0          0          0
1      85      0          0          0
1      86      0          0          0
1      87      0          0          0
1     1001      0          0          0

VRID  P/V    PRIO_0_SENT  INVALID_TYPE_ERR  ADDRESS_LIST_ERR  UNKNOWN_AUTHTYPE
-----
--More-- (q = quit)

```

Variable Definitions

Use the data in the following table to use the **show ipv6 vrrp statistics** command.

Variable	Value
<i>link-local WORD<0-127></i>	Shows statistics for a specific link-local address.
<i>vrid <1-255></i>	Shows statistics for a specific VRID.

Enable IPv6 VRRP Preempt Mode

You can configure IPv6 VRRP to preempt the existing router. If a new VRRP router is added to the network with a higher priority than the existing routers, then the new router becomes the master. If preempt-mode is disabled, then the new router does not become a master, it transitions to master only when the current master is down. By default, preempt-mode is enabled.

Procedure

1. Enter Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
interface GigabitEthernet {slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]} or interface vlan <1-4059>
```



Note

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enter the following command:

```
ipv6 vrrp <vrid> preempt-mode enable
```

3. Use the following command to set the IPv6 VRRP preempt-mode to its default value:

```
default ipv6 vrrp <vrid> preempt-mode
```

4. Use the following command to disable the IPv6 VRRP preempt-mode:

```
no ipv6 <vrid> preempt-mode enable
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Enabling IPv6 VRRP preempt-mode for interface 1/2

```
Switch:1(config-if)# ipv6 vrrp 1 preempt-mode enable
```

Variable definitions

Use the data in the following table to use the **ipv6 vrrp <vrid>** command.

Variable	Value
<i>enable</i>	Enables preempt-mode for VRRPv3 for IPv6.
<i>vrid <1-255></i>	Specifies the virtual router ID.

Configure VRRP using EDM

This section provides procedures to configure VRRP using Enterprise Device Manager (EDM).

Before You Begin

- Assign an IP address to the interface.
- Enable VRRP globally.

Enable VRRP Global Variables

Enable VRRP global variables to enable the VRRP function.

Before You Begin

To perform this procedure on a non-default VRF, you must first change the VRF instance. You cannot launch a VRF Context View for the derived name of the untrusted (*-unt) or unrestricted (*-unr) segments of a Segmented VRF. For information about how to use EDM for a non-default VRF, see [Select and Launch a VRF Context View](#) on page 4100. All parameters might not be available in non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **Globals** tab.
4. Configure the required features.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
NotificationCntl	Indicates whether the VRRP-enabled router generates SNMP traps for events. <ul style="list-style-type: none">• enabled—SNMP traps are generated• disabled—no SNMP traps are sent The default is enabled.
PingVirtualAddrEnable	Configures whether this device responds to pings directed to a virtual router IP address. The default is enabled.

Modify VRRP Parameters for an Interface

Before You Begin

- You must enable VRRP on a brouter port or VLAN.

About This Task

You can manage and configure VRRP parameters for the routing interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **Interface** tab.
4. Double-click **HoldDownTimer**, and type the number of seconds for the timer.

The **HoldDownState** field displays active when the hold-down timer is counting down and preemption occurs. The field displays dormant when preemption is not pending. When the hold-down timer is active, the **HoldDownTimeRemaining** field displays the seconds remaining before preemption.
5. In **Action**, select an option.
6. Select **Apply**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Specifies the index value that uniquely identifies the interface to which this entry applies.
VrId	Specifies a number that uniquely identifies a virtual router on a Virtual Router Redundancy Protocol (VRRP) router. The virtual router acts as the default router for one or more assigned addresses.

Name	Description
IpAddr	Specifies the assigned IP addresses for which a virtual router is responsible to provide backup.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Specifies the state of the virtual router interface: <ul style="list-style-type: none"> • Initialize—waiting for a startup event • Backup—monitoring availability and state of the master router • Master—functioning as the forwarding router, the primary VRRP router, for the virtual router IP addresses.
Control	Specifies whether VRRP is enabled or disabled for the port (or VLAN). The default is enabled.
Priority	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvertisementInterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second. Only the primary VRRP router sends advertisements. The default is 1.
MasterIpAddr	Specifies the IP address of the physical interface of the primary virtual router that forwards packets sent to the virtual IP addresses associated with the virtual router.
VirtualRouterUpTime	Specifies the time interval, in hundredths of a second, since the virtual router was initialized.
Action	Lists options to override the delay timer manually and force preemption: <ul style="list-style-type: none"> • none does not override the timer • preemptHoldDownTimer preempts the timer
HoldDownTimer	Configures the amount of time, in seconds, to wait before preempting the current VRRP primary router.
HoldDownState	Indicates the hold-down state of this VRRP interface. If the hold-down timer is operational, this variable is set to active; otherwise, this variable is set to dormant.

Name	Description
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.
CriticalIpAddr	<p>Configures the critical IP address for VRRP. This command specifies an IP interface on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from primary to backup) in case the interface stops responding.</p> <p>Note: In this context, <i>local</i> implies an address from the same VRF as the IP interface where VRRP is configured.</p>
CriticalIpAddrEnable	Configures the IP interface on the local router to enable or disable the backup. The default is disabled.
BackUpMaster	Enables the backup VRRP system traffic forwarding. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP system traffic forwarding is enabled or disabled. The default is disabled.
FasterAdvInterval	<p>Configures the Fast Advertisement Interval between sending VRRP advertisement messages. You must configure the same value on all participating routers. Enter the value in multiples of 200 milliseconds. The default is 200.</p>
FasterAdvIntervalEnable	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disable.
MasterHostname Note: Exception: not supported on 4220 Series, 5320 Series, or 5420 Series.	Displays the SPBM host name of the VRRP primary router for a Distributed Virtual Routing (DvR) interface.
Origin Note: Exception: not supported on 4220 Series, 5320 Series, or 5420 Series.	Displays how the VRRP interface was configured, either manually or through DvR.

Add a Secondary Virtual IP for VRRPv2

The first virtual IP configured on a VRRP interface is the primary virtual IP. After you create a VRRP interface, you can create secondary virtual IPs.

Before You Begin

- Create primary and secondary IP addresses on the VLAN.
- Create a VRRP instance on the VLAN.
- Configure a primary virtual IP.

About This Task

Secondary virtual IP addresses must be in the same subnet as a Secondary IP Interface on the VLAN.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **Interface** tab.
4. Select an interface, and then select **AssociatedIPAddr**.
5. On the **Address List** tab, select **Insert**.
6. Type the IP address.
7. Select **Insert**.

Address List *Field Descriptions*

Use the data in the following table to use the **Address List** tab.

Name	Description
IpAddr	Specifies a secondary virtual IP address.

Configure VRRP on a V3 Interface

Before You Begin

- Assign an IPv4 address to the interface
- Enable routing globally
- Do not configure RSMLT on the VLAN

About This Task

Perform this procedure to configure VRRP on a V3 interface on either a brouter port or a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **V3 Interface** tab.

4. Select **Insert**.
5. Beside **IfIndex**, select **Port** or **VLAN**.
6. Select a port or VLAN.
7. Select **OK**.
8. Type the virtual router ID.
9. Type the primary IP address.
10. Type the advertisement interval.
11. Select **Insert**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
InetAddrType	Specifies the source network INET Address Type.
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router.
PrimaryIpAddr	Specifies the virtual address assigned to the VRRP.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Shows the state of the virtual router interface. The possible states are <ul style="list-style-type: none"> • initialize—waiting for a startup event • backup—monitoring availability and state of the master router • master—functioning as the forwarding router for the virtual router IP addresses
Control	Displays whether VRRP is enabled or disabled for the port or VLAN.
Priority	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvInterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second.
MasterIpAddr	Specifies the IP address of the physical interface of the Master's virtual router.
UpTime	Indicates the time interval since this virtual router exited the INIT state.

Name	Description
CriticalIpAddr	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.
CriticalIpAddrEnabled	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
BackUpMaster	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the v1ST. The default is disabled.
BackUpMasterState	Indicates if the Backup-Master is operational up. If the switch is in Master state but the Backup-Master is enabled, then the BACKUP MASTER STATE will be down.
FasterAdvIntervalEnabled	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disabled.
FasterAdvInterval	Configures the interval between VRRP advertisement messages. The default is 200. Enter the values in multiples of 200 milliseconds.
PreemptMode	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master. The default is enabled.
Action	Lists options to override the hold-down timer manually and force preemption: <ul style="list-style-type: none"> • none does not override the timer. • preemptHoldDownTimer preempts the timer. This parameter applies only if the holddown timer is active.
HoldDownTimer	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.

Name	Description
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.
MasterAdvInterval	On the VRRPv3 master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

Configure VRRPv3 Checksum

About This Task

Perform this procedure to configure VRRPv3 checksum on either a brouter port or a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **V3 Checksum** tab.
4. Select **Insert**.
5. In **Interface**, select **Port** or **Vlan**.
6. Select the ports or VLANs from the list, and then select **OK**.
7. Select a type of checksum computation.
8. Select a VRRP version.
9. Select **Insert**.

V3 Checksum Field Descriptions

Use the data in the following table to use the **V3 Checksum** tab.

Name	Description
Interface	Shows VRRPv3 information about a specified interface.
rcIpConflIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
ChkSumComputation	Specifies the type of checksum computation, with Pseudo Header or without Pseudo Header.
VrrpVersion	Specifies the VRRP version as unspecified, v2, or v3.

Configure Fast Advertisement Interval on a Port

About This Task

Configure the Fast Advertisement Interval to send VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation pane, expand **Configuration > Edit > Port**.
3. Select **IP**.
4. Select the **VRRP** tab.
5. Select **Insert**.
6. Select **FasterAdvIntervalEnable**.
7. In **FasterAdvInterval**, type a value using multiples of 200 milliseconds.
8. Select **Insert**.

Configure Fast Advertisement Interval on a VLAN

About This Task

Configure the Fast Advertisement Interval to send VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.

Procedure

1. In the navigation pane, expand **Configuration > VLAN**.
2. Select **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Select **IP**.
6. Select the **VRRP** tab.
7. Select **Insert**.
8. Select **FasterAdvIntervalEnable**.
9. In **FasterAdvInterval**, type a value using multiples of 200 milliseconds.
10. Select **Insert**.

View VRRP Statistics

About This Task

View VRRP statistics to monitor network performance.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **Stats** tab.

Stats Field Descriptions

The following table describes parameters on the **Stats** tab.

Name	Description
ChecksumErrors	Specifies the number of VRRP packets received with an invalid VRRP checksum value.
VersionErrors	Specifies the number of VRRP packets received with an unknown or unsupported version number.
VrIDErrors	Specifies the number of VRRP packets received with an invalid VrID for this virtual router.

View VRRP Interface Statistics

About This Task

View VRRP statistics to manage network performance.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **Interface** tab.
4. Select an interface.
5. Select **Graph**.

Interface Field Descriptions

The following table describes parameters on the **Interface** tab.

Name	Description
AdvertiseRcvd	Specifies the number of VRRP advertisements received by this virtual router.
AdvertiseIntervalErrors	Specifies the number of received VRRP advertisement packets with a different interval is than configured for the local virtual router.
IPtTlErrors	Specifies the number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
PriorityZeroPktsRcvd	Specifies the number of VRRP packets received by the virtual router with a priority of 0.

Name	Description
PriorityZeroPktsSent	Specifies the number of VRRP packets sent by the virtual router with a priority of 0'.
InvalidTypePktsRcvd	Specifies the number of VRRP packets received by the virtual router with an invalid value in the 'type' field.
AddressListErrors	Specifies the packets received address list the address list does not match the locally configured list for the virtual router.
AuthTypeMismatch	Specifies the count of authentication type mismatch messages.
PacketLengthErrors	Specifies the count of packet length errors.
AuthFailures	Specifies the count of authentication failure messages.

View IP VRRPv3 Statistics

About This Task

Use the following procedure to view IPv6 VRRPv3 statistics for monitoring the network performance.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **V3 Stats** tab.

V3 Stats Field Descriptions

Use the data in the following table to interpret the **V3 Stats** tab.

Name	Description
InetAddrType	Shows that the address type of the statistical entry is IPv4.
ChecksumErrors	Specifies the total number of VRRP packets received with an invalid VRRP checksum value.
VersionErrors	Specifies the total number of VRRP packets received with an unknown or unsupported version number.
VrIdErrors	Specifies the total number of VRRP packets received with an invalid VRID for the virtual router.

Graph IP VRRPv3 Statistics

About This Task

Use the following procedure to view and graph IP VRRPv3 statistics for monitoring the network performance.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **VRRP**.
3. Select the **V3 Interface** tab.
4. Select an interface, and then select **Graph**.
5. Select one or more values.
6. Select a graph type:
 - Line Chart
 - Area Chart
 - Bar Chart
 - Pie Chart

V3 Interface Field Descriptions

Use the data in the following table to use the **V3 Interface** tab.

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router.
PrimaryIpAddr	Specifies the virtual address assigned to the VRRP.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Specifies the state of the virtual router interface: <ul style="list-style-type: none">• Initialize—waiting for a startup event• Backup—monitoring availability and state of the master router• Master—functioning as the forwarding router for the virtual router IP addresses.
Control	Specifies whether VRRP is enabled or disabled for the port (or VLAN). The default is enabled.

Name	Description
Priority	Specifies the priority value used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvInterval	Specifies the time interval (in seconds) between sending advertisement messages. The range is 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements. The default is 1.
UpTime	Specifies the time interval (in hundredths of a second) since the virtual router was initialized.
CriticalIpAddr	<p>This command specifies an IP interface on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.</p> <p>Note: In this context, local implies an address from the same VRF as the IP interface where VRRP is being configured.</p>
CriticalIpAddrEnabled	Configures the IP interface on the local router to enable or disable the backup. The default is disabled.
BackUpMaster	Enables the backup VRRP system traffic forwarding. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP system traffic forwarding is enabled or disabled. The default is disabled.
FasterAdvIntervalEnabled	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disable.
FasterAdvInterval	Configures the Fast Advertisement Interval between sending VRRP advertisement messages. The interval is between 200 and 1000 milliseconds, and you must enter the same value on all participating routers. The default is 200. You must enter the values in multiples of 200 milliseconds.
PreemptMode	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master.

Name	Description
Action	Lists options to override the delay timer manually and force preemption: <ul style="list-style-type: none"> • none does not override the timer • preemptHoldDownTimer preempts the timer
HoldDownTimer	Indicates the hold-down state of this VRRP interface. If the hold-down timer is operational, this variable is set to active; otherwise, this variable is set to dormant.
HoldDownTimeRemaining	Indicates the amount of time (in seconds) left before the HoldDownTimer expires.
MasterAdvInterval	On the VRRPv3 master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

IPv6 VRRP Configuration using EDM

This section provides procedures to configure IPv6 VRRP using Enterprise Device Manager (EDM).

Configure VRRP for an Interface

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Do not configure RSMLT on the VLAN.
- Change the VRF instance as required to configure a VRRP for an interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Configure VRRP to provide fast failover of a default router for IPv6 LAN hosts. VRRP supports a virtual IPv6 address shared between two or more routers that connect the common subnet to the enterprise network. VRRP for IPv6 provides a faster switchover to an alternate default router than is possible using the ND protocol.

Perform this procedure to configure VRRP on either a brouter port or a VLAN.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **VRRP**.
3. Select **V3 Interface**

4. Select **Insert**.
5. Beside **IfIndex**, select **Port** or **VLAN**.
6. In the dialog box that opens, select a port or VLAN.
7. Select **OK**.
8. Type the virtual router ID.
9. To control the packets sent to the IPv6 address associated to the virtual router, select **AcceptMode**.
10. Type the primary IP address.
11. Select **Insert**.

V3 Interface *Field Descriptions*

Use the data in the following table to use the **V3 Interface** tab.

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
InetAddrType	Specifies the address type for the VRRP interface.
VrId	Specifies a number that uniquely identifies a virtual router on a VRRP router.
PrimaryIpAddr	Specifies the link-local address assigned to the VRRP.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Shows the state of the virtual router interface. The possible states are <ul style="list-style-type: none"> • initialize—waiting for a startup event • backup—monitoring availability and state of the master router • master—functioning as the forwarding router for the virtual router IP addresses
Control	Displays whether VRRP is enabled or disabled for the port or VLAN.
Priority	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvInterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second.
MasterIpAddr	Specifies the IP address of the physical interface of the Master's virtual router.

Name	Description
UpTime	Indicates the time interval since this virtual router exited the INIT state.
CriticalIpAddr	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.
CriticalIpAddrEnabled	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
BackUpMaster	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the v1ST. The default is disabled.
BackUpMasterState	Indicates if the Backup-Master is up. If the switch is in Master state, but Backup-Master is enabled, then the BACKUP MASTER STATE will be down.
FasterAdvIntervalEnabled	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disable.
FasterAdvInterval	Configures the interval between VRRP advertisement messages. The default is 200. Enter the values in multiples of 200 milliseconds.
AcceptMode	Controls whether the VRRP master or backup master accepts packets (other than neighbor discovery packets) destined to the IPv6 address associated with the virtual router. The default value is disable.
PreemptMode	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master. The default is enabled.
Action	Lists options to override the hold-down timer manually and force preemption: <ul style="list-style-type: none"> • none does not override the timer. • preemptHoldDownTimer preempts the timer. This parameter applies only if the holddown timer is active.
HoldDownTimer	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.

Name	Description
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.
MasterAdvInterval	On the VRRP master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

Configure VRRP Notification Control

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Change the VRF instance as required to configure VRRP notification control on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

Perform this procedure to configure VRRP notification control.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **VRRP**.
3. Select the **Globals** tab.
4. Select **enabled**.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
NotificationCntrl	Indicates whether the VRRP-enabled router generates SNMP traps for events. <ul style="list-style-type: none">• enabled: Generate SNMP traps.• disabled: Do not generate SNMP traps. The default is enabled.

View IPv6 VRRP Statistics

View IPv6 VRRP statistics to monitor network performance.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.

2. Select **VRRP**.
3. Select the **Stats** tab.

Stats Field Descriptions

Use the data in the following table to use the **Stats** tab.

Name	Description
InetAddrType	Shows the type of IP address (IPv4 or IPv6).
ChecksumErrors	Shows the number of VRRP packets received with an invalid VRRP checksum value.
VersionErrors	Shows the number of VRRP packets received with an unknown or unsupported version number.
VrIdErrors	Shows the number of VRRP packets received with an invalid VrID for this virtual router.

View IPv6 VRRP Statistics for an Interface

View IPv6 VRRP statistics for a VLAN or port.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **VRRP**.
3. Select the **V3 Interface** tab.
4. Select an interface.
5. Select **Statistics**.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.

Name	Description
MasterTransitions	Shows the total number of times that the state of this virtual router has transitioned to master. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcdAdvertisements	Shows the total number of VRRP advertisements received by this virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
AdvIntervalErrors	Shows the total number of VRRP advertisement packets received for which the advertisement interval is different than the one configured for the local virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
IpTtlErrors	Shows the total number of VRRP packets received by the virtual router with IPv4 TTL (for VRRP over IPv4) or IPv6 Hop Limit (for VRRP over IPv6) not equal to 255. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcvdPriZeroPackets	Shows the total number of VRRP packets received by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
SentPriZeroPackets	Shows the total number of VRRP packets sent by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.

Name	Description
RcvdInvalidTypePkts	Shows the number of VRRP packets received by the virtual router with an invalid value in the 'type' field. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
AddressListErrors	Shows the total number of packets received for which the address list does not match the locally configured list for the virtual router. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
PacketLengthErrors	Shows the total number of packets received with a packet length less than the length of the VRRP header. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcvdInvalidAuthentications	Shows the total number of packets received with an unknown authentication type.

View IPv6 VRRP Statistics

View IPv6 VRRP statistics for a VLAN or port.

Before You Begin

Change the VRF instance as required to view IPv6 VRRP statistics on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IPv6**.
3. Select the **VRRP** tab.
4. Select an interface.
5. Select **Statistics**.

Statistics Field Descriptions

Use the data in the following table to use the **Statistics** tab.

Name	Description
MasterTransitions	Shows the total number of times that the state of this virtual router has transitioned to master. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcdAdvertisements	Shows the total number of VRRP advertisements received by this virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
AdvIntervalErrors	Shows the total number of VRRP advertisement packets received for which the advertisement interval is different than the one configured for the local virtual router. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
IpTtlErrors	Shows the total number of VRRP packets received by the virtual router with IPv4 TTL (for VRRP over IPv4) or IPv6 Hop Limit (for VRRP over IPv6) not equal to 255. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcvdPriZeroPackets	Shows the total number of VRRP packets received by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
SentPriZeroPackets	Shows the total number of VRRP packets sent by the virtual router with a priority of 0. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.

Name	Description
RcvdInvalidTypePkts	Shows the number of VRRP packets received by the virtual router with an invalid value in the 'type' field. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
AddressListErrors	Shows the total number of packets received for which the address list does not match the locally configured list for the virtual router. Discontinuities in the value of this counter can occur at reinitialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
PacketLengthErrors	Shows the total number of packets received with a packet length less than the length of the VRRP header. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of DiscontinuityTime.
RcvdInvalidAuthentications	Shows the total number of packets received with an unknown authentication type.

Graph IPv6 VRRP Statistics

About This Task

Use the following procedure to graph IPv6 VRRPv3 statistics for monitoring the network performance.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **VRRP**.
3. Select the **Stats** tab.
4. Select an interface, and then select **Graph**.
5. Select one or more values.
6. Select a graph type:
 - Line Chart
 - Area Chart
 - Bar Chart
 - Pie Chart

Stats *Field Descriptions*

Use the data in the following table to use the **Stats** tab.

Name	Description
InetAddrType	Shows the type of IP address (IPv4 or IPv6).
ChecksumErrors	Shows the number of VRRP packets received with an invalid VRRP checksum value.
VersionErrors	Shows the number of VRRP packets received with an unknown or unsupported version number.
VrIdErrors	Shows the number of VRRP packets received with an invalid VrID for this virtual router.

Configure Additional Addresses on the VRRP Brouter Port

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.

About This Task

Perform this procedure to configure the additional addresses for which the virtual router acts as a back up.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Port**.
2. Select **IPv6**.
3. Select the **VRRP** tab.
4. Select **AssociatedIPAddr**.
5. Select **Insert**.
6. Type the address.
7. Type the prefix length.
8. Select **Insert**.

Address List *Field Descriptions*

Use the data in the following table to use the **Address List** tab.

Name	Description
IpAddr	Specifies an IP address that is associated with a virtual router. The number of rows on this tab equals the number of IP addresses associated (backed up) by the virtual router
IpAddrPrefixLength	Specifies the length of the prefix in bits.

Configure Additional Addresses on a VRRP Interface

Perform this procedure to configure the additional addresses for which the virtual router acts as a back up.

Before You Begin

- Assign an IPv6 address to the interface.
- Enable routing globally.
- Do not configure RSMLT on the VLAN.
- To create a secondary virtual IP address on a VLAN:
 - Create primary and secondary IP addresses on the VLAN.
 - Create a VRRP instance on the VLAN.
 - Configure a primary virtual IP.
- Change the VRF instance as required to configure additional addresses on the VRRP interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About This Task

The first virtual IP configured on a VRRP interface is the primary virtual IP. After you create a VRRP interface, you can create secondary virtual IPs. Secondary virtual IP addresses must be in the same subnet as a Secondary IP Interface on the VLAN.



Note

You can also complete this configuration through the **Configuration > IP > VRRP > V3 Interface** navigation path.

Procedure

1. In the navigation pane, expand **Configuration > IPv6**.
2. Select **VRRP**.
3. Select the **V3 Interface** tab.
4. Select an interface.
5. Select **AssociatedIPAddr**.
6. Select **Insert**.
7. Type the address.
8. Type the prefix length.
9. Select **Insert**.

Address List *Field Descriptions*

Use the data in the following table to use the **Address List** tab.

Name	Description
IpAddr	Specifies an IP address that is associated with a virtual router. The number of rows on this tab equals the number of IP addresses associated (backed up) by the virtual router
IpAddrPrefixLength	Specifies the length of the prefix in bits.

Troubleshooting IPv6 VRRP

The following sections describe troubleshooting information for IPv6 Virtual Router Redundancy Protocol (VRRP).

VRRP Transitions

When a VRRP transition takes place with the backup taking over as the master, look for the following message in the syslog on the new master, as well as the old master. This message provides information to allow you to determine the cause of the transition.

```
IPv6 Vrrp State Transition Trap(Port/Vlan=200, Type=masterToInitialize,
Cause=shutdownReceived, VrId=20,VrIpAddr=fe80:0:0:0:0:0:0:200,
Addr=fe80:0:0:0:224:7fff:fe9d:1a03)
```

In this message, see the Type and Cause fields.

**Note**

Although the possible causes and types are listed in the following table, the system does not display all of the listed causes and types in the trap/log message.

The following table describes the VRRP transition types.

Table 354: Transition type

Type value	Type definition
1	None
2	Master to backup
3	Backup to master
4	Initialize to master
5	Master to initialize
6	Initialize to backup
7	Backup to initialize

Table 354: Transition type (continued)

Type value	Type definition
8	Backup to backup master
9	Backup master to backup

The following table describes the VRRP transition causes.

Table 355: Transition cause

Cause value	Cause definition
1	None
2	Higher priority advertisement received
3	Shutdown received
4	VRRP address and physical address match
5	Master down interval
6	Preemption
7	Critical IP goes down
8	User disabling VRRP
9	VRRP status synced from primary
10	IPv6 interface on which VRRP is configured goes down
11	Lower priority advertisement received
12	Advertisement received from higher interface IP address with equal priority
13	Advertisement received from lower interface IP address with equal priority
14	User enabled VRRP
15	Transition because of any other cause

Enable Trace Messages for IPv6 VRRP Troubleshooting

Use this procedure to enable trace messages for IPv6 VRRP.

When VRRP is enabled on two routing switches, the master-backup relationship forms with one router taking the responsibility of routing. If the master-backup relationship is not formed between the VRRP virtual routers, look for the following trace messages to ensure that the master is sending the advertisements correctly and the backup is processing them.

Procedure

1. Enter Privileged EXEC mode:
`enable`

2. To troubleshoot IPv6 VRRP, you can enable RCIP6 trace messages with the following command:

```
trace level 66 3
```

3. Obtain additional trace information:

```
trace ipv6 nd enable
```

```
trace ipv6 base enable all
```

```
trace ipv6 forwarding enable all
```

```
trace ipv6 rtm enable all
```

```
trace ipv6 transport enable all
```

4. When VRRP is enabled on two routing switches, the master-backup relationship forms with one router taking the responsibility of routing. If the master-backup relationship is not formed between the VRRP virtual routers, look for the following trace messages to ensure that the master is sending the advertisements correctly and the backup is processing them. On the master router, look for the following RCIP6 trace messages.

- tMainTask RCIP6: rcip6_vrrp.c: 5118: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Am Master for Vrid 200 on IfIndex 2053 Timer 1

If VRRP is enabled on the interface, this timer kicks off every second and shows the state for the VRID.

- [11/18/09 15:08:20:383] tMainTask RCIP6: rcip6_vrrp.c: 5924: ipv6VrrpSendAdvertisement: for Vrid 200 on IfIndex 2053

```
[11/18/09 15:08:20:583] tMainTask RCIP6: rcip6_vrrp.c:
5175: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic:
ipv6VrrpSendAdvertisement
```

The preceding trace messages show that the VRRP master is sending the advertisements correctly at the end of advertisement interval for a VRID.

5. On the backup router, look for the following RCIP6 trace messages.

- tMainTask RCIP6: rcip6_vrrp.c: 5236: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Am Backup for VrId 200 on IfIndex 2052 Timer 1
- tMainTask RCIP6: rcip6_vrrp.c: 4854: ipv6VrrpIn: Vrid 200 on IfIndex 2052
- tMainTask RCIP6: rcip6_vrrp.c: 5545: VRF name: GlobalRouter (VRF id 0): rcIpVrrpProcessAdvt: Am backup for Vrid 200 on IfIndex 2052

The preceding trace messages show that the backup router is receiving the advertisements sent by the master and correctly processing them.

Risks Associated with Enabling Trace Messages

When traces are enabled on VRRP master, VrrpTic messages are logged for every second and any other configured traces keep displaying, so there is no guarantee that the backup will receive the advertisement from the master within 3 seconds, so it can

transit to master also. There is also the risk of toggling of VRRP states (from backup to master and back again).

Enable the limited traces based on whichever is required.

VRRP with Higher Priority Running as Backup

The VRRP router with the higher priority can display as the backup for the following reasons

- Hold-down timer is running.
- The configured Critical IP is not reachable or does not exist.

If the critical-IP is configured for VRRP master, and the critical interface goes down or is deleted, the master transitions to the backup state. In this case, the log shows the transition cause as 1 like many other cases.

If the holddown timer is configured for VRRP master, the holddown timer delays the preemption, giving the device, which is becoming the master enough time to construct routing tables.

Procedure

1. To determine that the issue is with the critical interface, look for the following trace message.

```
tMainTask RCIP6: rcip6_vrrp.c: 5152: VRF name: GlobalRouter (VRF id 0): ipv6VrrpTic: Becoming backup for Vrid 200 on IfIndex 2052 because of invalid critical IP
```

2. If the holddown Timer is configured for VRRP master, the holddown timer delays the preemption, giving the device, which is becoming the master enough time to construct routing tables.

```
tMainTask RCIP6: rcip6_vrrp.c: Enter in HoldDown processing,Vrid 200 LastRecvd 0 MasterDown 3, Holddown time remaining 970, Holddownstate 2
```



VXLAN Gateway

[VXLAN Gateway Fundamentals](#) on page 4171

[VXLAN Gateway configuration using the CLI](#) on page 4183

[OVSDB protocol support for VXLAN Gateway Configuration Using CLI](#) on page 4190

[VXLAN Gateway configuration using EDM](#) on page 4196

[OVSDB protocol support for VXLAN Gateway configuration using EDM](#) on page 4201

[Configuration Examples](#) on page 4204

VXLAN Gateway terminates virtual extensible LAN (VXLAN) tunnels that stretch emulated Layer 2 segments over an IP network.

The following topics provide necessary concepts and procedures to configure VXLAN Gateway and Open vSwitch Database Management Protocol (OVSDB) support for VXLAN Gateway.

Table 356: VXLAN Gateway product support

Feature	Product	Release introduced
VXLAN Gateway	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.0

Table 356: VXLAN Gateway product support (continued)

Feature	Product	Release introduced
OVSDB protocol support for VXLAN Gateway	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.0
ECMP support for VXLAN Gateway	4220 Series	Not Supported
	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	7830 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.0

VXLAN Gateway Fundamentals

VXLAN Gateway terminates virtual extensible LAN (VXLAN) tunnels that stretch emulated Layer 2 segments over an IP network.


VXLAN is an encapsulation protocol for running an overlay network over an existing Layer 3 infrastructure. This tunneling scheme uses a VLAN-like encapsulation mechanism to encapsulate MAC-based traffic. The VXLAN overlays separate workloads from physical networks, which facilitates the movement of VMs in a multi-tenant environment.

For more information about VXLAN, see the Internet Engineering Task Force (IETF) standard RFC 7348.

VXLAN Gateway as a VTEP

The VXLAN Gateway feature terminates VXLAN tunnels and operates as a VXLAN tunnel endpoint (VTEP). The VXLAN Gateway is a hardware-based VTEP that allows a VXLAN to communicate with VLANs, other VXLANs, as well as Fabric Connect I-SIDs.

The following figure shows a VXLAN Gateway (labeled as VTEP Source) connecting VXLAN segments to an SPB network. This type of deployment provides a solution for VXLAN in data centers to interoperate with Fabric Connect networks (labeled as SPB).



VXLAN Segments

VXLAN tunnels are also called VXLAN segments. VXLAN segments provide the ability to separate, abstract, and decouple the physical topology from a logical or virtual topology by using encapsulated tunneling. VXLAN segments allows only VMs within the same VXLAN tunnel to communicate with each other.

VNID

Each VXLAN segment has a 24-bit segment ID called a VXLAN Network Identifier (VNID). VNID allows up to 16 million VXLAN segments to coexist within the same administrative domain, which enables VMs to migrate between servers in the same data center or between distributed data centers. Each VTEP can support multiple VNIDs. For information about maximum scaling numbers, see *Fabric Engine Release Notes*.

The VNID uses the inner MAC frame originated by the individual VM. Overlapping MAC addresses can cross VXLAN segments, but never have traffic crossover because the traffic is isolated using the VNID qualifier. Due to this encapsulation, only the VTEP knows the VNID and its associated VXLAN tunnel. VMs are unaware of the VXLAN.

VM Traffic

Only VMs within the same VXLAN segment can communicate with each other. The following steps show how a VM communicates with a VM on a different host within the same VXLAN segment:

1. The VM sends a MAC frame destined to the target.
2. The VTEP looks up the VNID that the VM is associated with to verify that the destination MAC is on the same segment. If the destination MAC is on the same segment, the process continues.
3. The VTEP inserts an outer header comprising an outer MAC, outer IP address, and VXLAN header in front of the original MAC frame.
4. The VTEP transmits the final packet out to the destination IP address of the remote VTEP, connecting the destination VM addressed by the inner MAC destination address.
5. The remote VTEP verifies that the VNID is valid for the destination VM.
6. If the VNID is valid, the remote VTEP strips the packet outer header and passes the packet to the destination VM. The destination VM never knows about the VNID, or that the frame transmitted with VXLAN encapsulation.

VXLAN Gateway Management Methods

Two methods are available to manage VXLAN Gateway:

- Static management using CLI and local configuration files.
- Dynamic management using OVSDb protocol support for VXLAN Gateway.

For static management, you use CLI or local configuration files to configure and manage the VXLAN Gateway hardware-based VTEP functions. You must manually configure the Source-VTEP, VNID, VNID to I-SID bindings, Remote-VTEP, and VNI to Remote-VTEP associations, and MAC learning on the VNID only occurs at the data-plane level.

For dynamic management, you use OVSDb protocol support for VXLAN Gateway to configure and manage the VXLAN Gateway hardware-based VTEP functions with Open vSwitch Database Management Protocol (OVSDb). You must configure at least one physical Network Virtualization Controller (NVC) running VMware NSX with OVSDb features. You can add a Hardware VTEP, add a logical switch, configure VNID to VTEP bindings, and configure a replication cluster in NSX. The VXLAN Gateway must have OVSDb enabled and the Network Virtualization Controller (NVC) must communicate on the OVSDb managed interface. You must manually configure the Source-VTEP and NVC IP addresses, and the VNID to I-SID binding in NSX. The NVC can manage the VNID, Remote-VTEP, and VNI to Remote-VTEP associations. The NVC distributes the hosts MAC and IP addresses learned by the VTEP over the VXLAN.

The following table shows the functional differences between the two VXLAN Gateway management methods.

Table 357: VXLAN Gateway Management Methods

Function	Static management	Dynamic management
VNID provisioning	Manually from CLI or EDM	Automatically from NVC
VNID to I-SID binding	Manually from CLI or EDM	Automatically from NVC
Remote-VTEP	Manually from CLI or EDM	MAC learning from NVC
VNID to VTEP association for replicating BUM traffic	Manually from CLI or EDM	MAC learning from NVC
MAC learning at data-plane	Not distributed to other VXLAN Gateways	Distributed from NVC to other VXLAN Gateways
ARP suppression	Not supported	Not supported

VXLAN Modes

The VXLAN Gateway implementation is available in two modes: Base Interworking Mode and Full Interworking Mode. By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes through the **vxlan-gw-full-interworking-mode** boot flag.



Important

Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.

Base Interworking Mode

Base Interworking Mode is the default mode. In this mode, VXLAN Gateway supports Layer 2 gateway communication at line rate between VXLAN and traditional VLAN environments.

Base Interworking Mode has the following limitations:

- No support for VXLAN-to-VXLAN communication.
- No support for VXLAN-to-SPB communication.
- No support for SMLT, vIST, or Simplified vIST.
- No support for OVSDB protocol support for VXLAN Gateway.

Each VNID represents a service instance for the VXLAN Gateway. To enable a VNID to extend into the SPB network, you must use Full Interworking Mode.



Note

If your configuration does not have a VXLAN Gateway but has SMLT or vIST, you can enable VXLAN Gateway in the Full Interworking Mode only. SMLT and vIST are incompatible and mutually exclusive with the Base Interworking Mode.

Full Interworking Mode

Full Interworking Mode is not the default mode. To configure Full Interworking Mode, you must enable the **vxlan-gw-full-interworking-mode** boot flag, save the configuration, and then reboot the switch.

Full Interworking Mode supports the Base Interworking Mode communication between VXLAN and traditional VLAN environments as well as VXLAN-to-VXLAN communication and all SPB functionality including vIST and SMLT. Layer 2 and Layer 3 services extend across VXLAN, VLAN, and SPB domains using an internal loopback. In Full Interworking Mode, VXLAN traffic flows at less than full line rate and the total amount of available internal loopback bandwidth depends on the platform used.

Full Interworking Mode has the following limitations:

- No support for Simplified vIST.
- A VNID cannot have the same value as an I-SID. A consistency check ensures that this does not happen if you change the boot configuration flag to turn off Full Interworking Mode and go into Base Interworking Mode.

Types of VXLAN Gateway Deployments

VXLAN Gateway provides the following deployment solutions:

- [VXLAN-to-VLAN Layer 2 gateway deployment](#)
- [VXLAN-to-VXLAN Layer 3 gateway deployment](#)
- [VXLAN-to-SPB gateway deployment](#)

These gateways provide VXLAN tunnel termination functions for switches that could be Top of Rack (ToR) switches, access switches, or switches higher up in the data center network topology such as a core switch or even WAN edge devices. The last case (WAN edge) could involve a Provider Edge (PE) router that terminates VXLAN tunnels in a hybrid cloud environment.

VXLAN-to-VLAN Layer 2 Gateway Deployment

The VXLAN-to-VLAN deployment is where the VXLAN Gateway communicates between VXLAN segments and VLAN segments. Use this scenario when nodes on a VXLAN overlay network need to communicate with nodes on legacy networks that are VLAN based as shown in the following figure.




Figure 257: VXLAN-to-VLAN deployment

In this type of deployment, the VXLAN Gateway maps the VXLAN VNID to a VLAN in the following way:

- When packets traverse a VXLAN segment into the VLAN segment, the VXLAN VTEP removes the VXLAN header and forwards regular VLAN packets to a physical port based on the destination MAC address of the inner Ethernet frame.

The VTEP discards decapsulated frames with the inner VLAN ID unless you explicitly configure them to be passed on to the non-VXLAN interface. The VXLAN VNID to VLAN mapping is statically configured.

- In the reverse direction when packets traverse a VLAN segment into a VXLAN segment, the VXLAN VTEP maps the incoming frames from the non-VXLAN interfaces to a specific VXLAN overlay network based on the VLAN ID in the frame. This complies with the standard VXLAN RFC.

The VTEP also accepts both untagged and tagged inner frames and adds or replaces the tag of the inner frame with the local VLAN mapped to the VNID before sending out the frame.

The VTEP removes the VLAN ID of the original frame before it encapsulates the frame for VXLAN.

VXLAN-to-VXLAN Layer 3 Gateway Deployment

The VXLAN-to-VXLAN deployment is where the VXLAN Gateway communicates between VXLAN segments. Use this scenario when one VXLAN segment needs to communicate with another VXLAN segment.

When a packet traverses a VXLAN segment and enters the VXLAN Gateway, the VXLAN VTEP maps one VXLAN segment to another in the following way:

- Decapsulates the VXLAN header.
- Performs a route lookup of the customer packet.
- Identifies that the next hop VLAN is in another VXLAN segment.

- Encapsulates the packet with a new VXLAN header with a new VNID.
- Forwards the packet out.

The following figure shows two virtualized servers attached to a Layer 3 infrastructure. The servers could be on the same rack, on different racks, or across data centers within the same administrative domain.

There are four VXLAN overlay networks identified by the VNIDs: 22, 34, 74, and 98. Notice how the VMs are associated with a VNID. For example, VM1-1 in Server 1 and VM2-4 on Server 2 are on the same VXLAN overlay network identified by VNID 22.

The VMs do not know about the overlay networks and transport method because the encapsulation and decapsulation happen transparently at the VTEPs on Servers 1 and 2. The other overlay networks and the corresponding VMs are: VM1-2 on Server 1 and VM2-1 on Server 2 both on VNID 34, VM1-3 on Server 1 and VM2-2 on Server 2 on VNID 74, and finally VM1-4 on Server 1 and VM2-3 on Server 2 on VNID 98.




Figure 258: VXLAN-to-VXLAN deployment

VXLAN-to-SPB Gateway Deployment

The VXLAN-to-SPB deployment is where the VXLAN Gateway communicates between a VXLAN segment and an SPB domain. This type of deployment provides a solution for customers using VXLAN in their data-centers to interoperate with Fabric Connect as shown in the following figure.




Figure 259: VXLAN-to-SPB deployment

The VXLAN Gateway maps the VXLAN VNID to I-SID in the following way:

- When packets traverse a VXLAN segment into an SPB domain, the VXLAN VTEP decapsulates the VXLAN packet, encapsulates it with MAC-in-MAC encapsulation, and then forwards it into the SPB domain.
- In the reverse direction, the VXLAN VTEP decapsulates the MAC-in-MAC packet from an SPB domain, encapsulates it with a VXLAN header, and then forwards it into a VXLAN segment.

VXLAN Gateway Considerations and Limitations

Review the following considerations, limitations, and behavioral characteristics associated with VXLAN Gateway.

Unsupported Features

VXLAN Gateway does not support the following features:

- IP Multicast over Layer 3 VSN Fabric Connect traffic cannot be carried over into a VXLAN domain through a VXLAN Gateway. Layer 3 VSN Fabric Connect requires MAC-in-MAC encapsulation over VXLAN encapsulation, and the VXLAN Gateway data plane does not support double-encapsulation. However, IP Multicast over IP Shortcuts Fabric Connect can be carried over into a VXLAN domain through a VXLAN Gateway.
- SPB-PIM Gateway and VXLAN Gateway interoperability is not supported.
- VXLAN Gateway does not support Simplified vIST in any mode.

QoS and VXLAN Gateway

- If present in the packet, VXLAN Gateway honors user-configured ingress dot1p and DSCP mappings to derive the internal CoS.
- Customer packets are not remarked.
- Control packets are handled by high priority CoS queues.
- DSCP bits in the outer IP header of VXLAN-encapsulated packets are always derived from the internal QoS, irrespective of the ingress port DiffServ configuration.

Customer packet IP DSCP bits are not modified as part of VXLAN encapsulation. For more information about QoS, see [Quality of Service](#) on page 2983.

- DSCP bits in the inner IP header of VXLAN-encapsulated packets change when traffic comes in on a Layer 2 Trusted or a Layer 3 Untrusted port with DiffServ enabled.

ECMP Support for VXLAN Gateway and Fabric Extend

VXLAN Gateway requires ECMP support to communicate with remote VTEPs. The software extended this ECMP support to Fabric Extend Layer 3 core tunnels. Therefore, if your switch supports VXLAN Gateway, it also supports ECMP for both VXLAN Gateway and Fabric Extend.

Base Interworking Mode Considerations and Limitations

- After associating a VNID with an I-SID (**vnid <vnid value> i-sid <isid value>**), you cannot create an I-SID with the same value as the **<vnid value>** and vice versa.
- There is no support for SMLT, vIST, or Simplified vIST in Base Interworking Mode.
- Neither Base Interworking Mode nor Full Interworking Mode support Fabric Attach endpoints.
- Base Interworking Mode does not permit SMLT and vIST configurations. If you change from Full Interworking Mode to Base Interworking Mode, make sure there are none of these configurations.
- In Base Interworking Mode, the **untagged-traffic <port> << mlt>** support is limited to untagged traffic forwarding only. No control packet forwarding is supported on this port or MLT.

LACP MLT is not supported as part of **untagged-traffic <mlt>**.

Full Interworking Mode Considerations and Limitations

- Whenever you change from Base Interworking Mode to Full Interworking Mode, you have to reboot the switch. The software automatically changes the configuration to be in line with Full Interworking Mode.
- Full Interworking Mode does not have VXLAN endpoints. Configure ELAN endpoints under the ELAN I-SID only.
- Neither Base Interworking Mode nor Full Interworking Mode support Fabric Attach endpoints.

General Configuration Guidelines

- In VXLAN environments that have an underlying IP network with SMLT deployment, you must enable RSMLT in the underlying network on vIST VTEP devices.
- If you use more than one VXLAN Gateway node for the same I-SID, you must ensure that loops are prevented.
- The VTEP remote destination IP table in the datapath hardware is shared with the Fabric Extend IP core remote tunnel destination IP table.
 - The maximum number of FE tunnel destination IP addresses is 256.
 - The maximum number of VTEP IP addresses is 500.

For every FE tunnel destination IP you configure, you must reduce the number of VTEP IPs by one. For example, if you configure the maximum number of FE tunnel destination IPs then the maximum of VTEP remote destination that you can configure is $500-256=244$.

- The total number of MACs that the VTEP learns in the switch is constant and includes the MACs learned on VNID, I-SID, and on VLANs (normal and CVLANs).
 - MAC addresses that the VTEP learns from the VNID table are also learned when an I-SID is associated with the VNID. Therefore, the maximum number of MAC addresses that the VTEP can support is reduced accordingly.
 - If a VLAN is associated with an I-SID, and that I-SID is associated with a VNID, then the VTEP learns each customer MAC address from three different tables (VLAN, I-SID, and VNID). This consumes three MAC entry records. In a vIST scenario where all I-SIDs have associated VNIDs, then the max number of CMACs that the VTEP learns is reduced by $224K$ divided by $3 = 74K$ MAC addresses.
 - If a Switched-UNI endpoint is associated with an I-SID, then the VTEP learns the CMAC only once in the I-SID table. If this I-SID is associated with a VNID, then the VTEP learns the CMAC twice: once in the I-SID table and once in the VXLAN VNID table.
- IP filter rules that you create do not work for VXLAN tunnel terminated packets. This restriction also applies to the Fabric Extend Layer 3 core.
- All ports in an MLT configured as a Switched-UNI endpoint or a VXLAN endpoint have the same properties as configured under the MLT interface.
- A VXLAN tunnel VTEP source IP cannot be a broadcast or multicast IP address. Configure it as a loopback IP address and not a router IP address.
- The VXLAN source address (**vtep-source-ip**) should be different from the SPBM IP Shortcut (**ip-source-address**) and Fabric Extend (**ip-tunnel-source-address**) source addresses.
- If you configure the loopback IP address under a VRF, the **vtep-source-ip** address must be set up under the same VRF.
- The remote VTEP IP address (**vtep <id> ip <ip address>**) cannot be local to the system.
- You can configure a maximum of 500 remote VTEPs per VNID.
- You can change the VTEP source IP and remote VTEP IP dynamically.
- VTEPs have to be reachable for the tunnel to be up.
- The tunnel ID for a VTEP has to be the same on both vIST peers.
- If the VTEP source IP is configured under a VRF, then the VRF cannot have an I-SID associated with it or vice versa.
- Two VNIDs cannot be mapped to the same I-SID.
- An I-SID associated with a VNID cannot be a T-UNI or E-Tree I-SID.
- To prevent routing loops, the platform VLANs used for VXLAN interworking should NOT be in the same VRF or GRT as the IP interfaces in the underlay network.
- You cannot delete a VTEP source IP address (**vtep-source-ip**) when VTEPs are configured.

- If an ELAN I-SID attached to a VNID has a platform VLAN associated with it, then you cannot delete the platform VLAN or change the I-SID associated with that VLAN.
- You cannot delete the loopback IP associated with a VTEP source IP.
- If the VTEP source IP is configured under a VRF, then you cannot delete the VRF.
- You cannot delete or modify the VLAN with I-SID that is associated with a VNID.
- If an ELAN I-SID attached to a VNID also has a platform VLAN associated with it, then you cannot delete the platform VLAN or change the I-SID associated with that VLAN.
- VXLAN Gateway cannot coexist with virtual NNI links on boundary nodes. You cannot enable VXLAN Gateway on a loopback interface for a virtual NNI link and you cannot create a loopback interface for a virtual NNI link when VXLAN Gateway is enabled.

OVSDB protocol support for VXLAN Gateway Considerations and Limitations

Review the following considerations, limitations, and behavioral characteristics associated with OVSDB protocol support for VXLAN Gateway.

Network Virtualization Controllers Dependencies and Restrictions

OVSDB protocol support for VXLAN Gateway requires at least one Network Virtualization Controller (NVC) that runs OVSDB:

- OVS is an open source implementation of a virtual switch and OVSDB is a standard defined in RFC 7047.
- You require at least one NVC. The NVC is an ESXi host server that runs the VMware NSX Network Virtualization and Security Platform (minimum version 6.2.4).

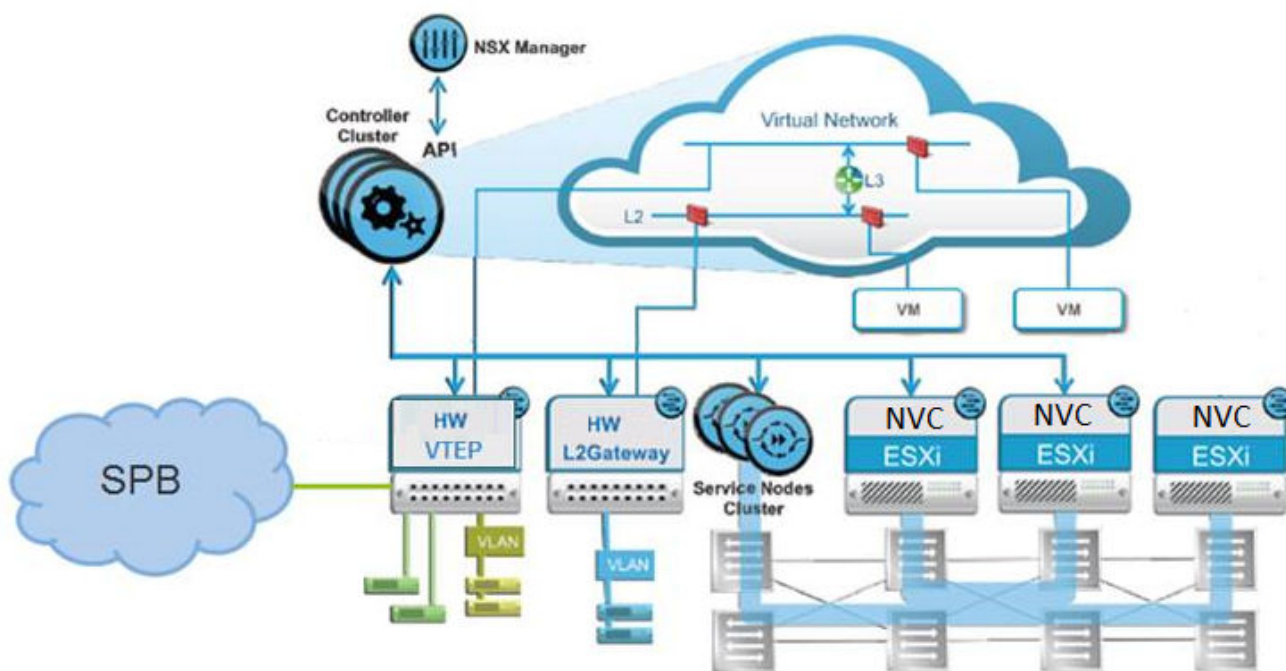


Note

VMware NSX 6.2.4 has known issues that can cause VXLAN discovery delays and NVC connectivity issues in certain scenarios, for more information see *Fabric Engine Release Notes*.

- You must configure VMware NSX. You can add a Hardware VTEP, add a logical switch, configure VNID to VTEP bindings, and configure a replication cluster in NSX.
- The NSX logical switch on the NVC is assigned a VNID. The VXLAN Gateway is assigned an OVSDB managed interface I-SID. You must configure the VNID to I-SID binding in NSX.
- SSL is the default protocol for NVC to VXLAN Gateway communications.
- You can configure more than one NVC in a controller cluster for high availability and load-balancing.
- The NVC performs MAC learning on the VXLAN tunnel.
- ARP suppression is not supported.
- One NVC manages all encapsulated Broadcast, Unknown Unicast, and Multicast (BUM) traffic. This NVC replicates the BUM traffic to all other VTEPs in the network. If connectivity to the NVC managing BUM traffic fails, a loss of BUM traffic can occur. NVC cluster NSX service node replication is not applicable for BUM traffic.

The following diagram shows an example of NVCs managing a distributed virtual network.



OVSDB protocol support for VXLAN Gateway Dependencies and Restrictions

- The VXLAN Gateway must operate in Full Interworking mode.
- You must configure at least one NVC connection to the VXLAN Gateway.
- You can configure up to three NVCs connections to one VXLAN Gateway in this release.
- The VXLAN Gateway must use the Segmented Management Instance IP address to connect with the NVCs. For more information about how to configure Segmented Management Instance, see [Segmented Management](#) on page 106.
- You cannot change or delete the VTEP source IP address when OVSDB is enabled.
- The OVSDB managed interface I-SIDs are communicated to the NVC as physical ports of the switch.
- EDM support for OVSDB is limited in this release. You can use CLI where EDM configuration is unavailable.
- The VXLAN Gateway general considerations and limitations also apply to OVSDB protocol support for VXLAN Gateway.



Note

NVC connectivity issues for SSL communication failures are not logged on the switch. You can use the logs on the NVC to troubleshoot SSL communications with the VXLAN Gateway.

OVS Modules on VXLAN Gateway

OVSDB protocol support for VXLAN Gateway uses OVS modules on the VXLAN Gateway. The switch runs the following OVS modules:

- OVSDB server — Manages the database.
- Database — Manages the OVS schema and HW VTEP schema and JSON content.

OVSDB Replication Considerations and Limitations



Note

OVSDB replication is different than NSX service node replication. Service node replication is configured and supported in NSX. OVSDB replication uses vIST and is configured and supported on the VXLAN Gateway.

- You require at least two NVCs to enable OVSDB replication.
- OVSDB replication operates in an ACTIVE-STANDBY server configuration.
- You must disable OVSDB on the standby vIST device before disabling OVSDB on the active vIST device.
- One active NVC and one standby NVC are supported by the VXLAN Gateway.
- VNID to I-SID binding, Remote VTEPs, VNI to Remote VTEPs for BUM, Remote MACs, and Local MACs are replicated across NVCs.

VXLAN Gateway configuration using the CLI

This section provides procedures to configure VXLAN Gateway using the CLI.

Enable VXLAN Gateway in Full Interworking Mode

Enables VXLAN Gateway in Full Interworking Mode, which supports SPB, SMLT, and vIST.

About This Task



Important

After you change this boot flag, you must save the change to the configuration file and reboot the switch.

By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes by enabling this boot configuration flag.

The **no** operator results in the default Base Interworking Mode. In this mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`

2. Enable Full Interworking Mode:
`boot config flags vxlan-gw-full-interworking-mode`
3. Save the changed configuration.
`save config`
4. Restart the switch.
`reset`
5. Verify that VXLAN Gateway is in Full Interworking Mode:
`show boot config flags`

The **true** parameter in the output indicates that Full Interworking Mode is enabled.

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags vxlan-gw-full-interworking-mode
```

Configure VXLAN Gateway

Use the following procedure to configure VXLAN Gateway parameters to enable SPBM to operate on the switch.

Before You Begin

Use the **show boot config flags** to verify that you are in the mode you want. The **vxlan-gw-full-interworking-mode** should be set to **false** for the Base Interworking Mode (default) and **true** for the Full Interworking Mode.

Procedure

1. Enter Global Configuration mode:
`enable`
`configure terminal`
2. Configure the VTEP source IP address:
`vtep source-ip <A.B.C.D> [vrf WORD<1-16>]`
3. Configure the remote VTEP IP address:
`vtep <1-500> ip <A.B.C.D> [name WORD<1-64>]`
4. Create a VNID instance:
`vnid <1-16777215> i-sid <1-16777215>`



Note

The command prompt changes to **#vxlan** to indicate that you are now in VXLAN Configuration mode for the VNID specified in `vnid <1-16777215>`.

5. Associate VLANs in a port or MLT list to this VNID instance:

```
c-vid <1-4094> port <{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...][slot/all][all]]>
```

or

```
c-vid <1-4094> mlt <1-4094>
```

6. Specify the ports or MLT in this VNID instance that will support untagged traffic:

```
untagged-traffic port <{slot/port[/sub-port]]>
```

or

```
untagged-traffic mlt <1-4094>
```

7. Associate VTEPs to this VNID instance:

```
vtep <1-255>
```

Example

For an example of this procedure, see [VXLAN Gateway configuration example in Base Interworking mode](#) or [VXLAN Gateway configuration example in Full Interworking mode](#).

Variable Definitions

Use the data in the following table to use the **vtep** command.

Variable	Value
<i>source-ip</i> <A.B.C.D> [<i>vrf</i> <i>WORD</i> <1-16>]	Specifies the VXLAN tunnel end point (VTEP) source IP address, which can be on the GRT or a VRF. Note: The VTEP source IP address must be on a loopback interface.
<1-500> <i>ip</i> <A.B.C.D> [<i>name</i> <i>WORD</i> <1-64>]	Specifies an index value and an IP address that uniquely identifies this remote VTEP. Optionally, you can assign a specific name to this tunnel. By default, the switch assigns a name in this format: VTEP-<#ID>-<IP address> Note: The remote VTEP IP address cannot be a local, broadcast, or multicast IP address.

Use the data in the following table to use the **vnid** command.

Variable	Value
<code><1-16777215> i-sid <1-16777215></code>	Uses this VNID and I-SID information to create a VNID instance and enter VXLAN Configuration Mode in the CLI. Note: A VNID must not have the same value as an I-SID.

Use the data in the following table to configure commands in the VXLAN Configuration Mode (`#vxlan`). These commands apply to the VNID instance that you specified in the `vnid <1-16777215> i-sid <1-16777215>` command.

Variable	Value
<code>c-vid <1-4094> port <{slot/port[/sub- port] [-slot/port[/ sub-port]] [,...] [slot/all] [all]}></code>	Specifies a value that uniquely identifies the customer VLAN ID and ports of this ELAN end point. The switch reserves the following VLAN IDs: <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
<code>c-vid <1-4094> mlt <1-4094></code>	Specifies a value that uniquely identifies the customer VLAN ID and MLTs of this ELAN end point. The switch reserves the following VLAN IDs: <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
<code>untagged-traffic port <{slot/port[/sub- port] [-slot/port[/ sub-port]] [,...] [slot/all] [all]}></code>	Specifies the ports that support untagged traffic.
<code>untagged-traffic mlt <1-4094></code>	Specifies the MLTs that support untagged traffic.
<code>vtep <1-255></code>	Lists the remote VTEP destinations to associate with the specified VNID.

Flush the MAC Forwarding Table

Use this procedure to flush all the learned MAC addresses from the forwarding database of the selected VNID.

Procedure

1. Enter Global Configuration mode:
`enable`

`configure terminal`
2. Flush the MAC FDB:
`vnid mac-address-entry <1-16777215> flush`

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# vnid mac-address-entry 100 flush
```

VXLAN Gateway show commands

Use the procedures in this section to display specific information about VXLAN Gateway on the switch.

Display VTEP Source Information

Use the following procedure to display information about the VTEP source.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display the VTEP's source IP address and the name of the VRF:
`show vtep local`

Example

```
Switch:1# show vtep local
=====
                        VTEP General Info
=====
vtep source-ip-address: 198.51.100.1
vtep vrf : GlobalRouter
```

Display Remote VTEP Information

Use the following procedure to display information about the remote VTEP.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. Display information about all of the remote VTEPs or use the option to specify a particular VTEP:
`show vtep remote [<1-500>]`

Example

```
Switch:1# show vtep remote
=====
Remote VTEP Info
=====
ID      NAME      VTEP      L3_VTEP_NEXT_HOP_INFO
  DEST-IP  PORT/MLT  VLAN      VRF
-----
1       vtep-1      198.51.100.6  Mlt78      3851      GlobalRouter
56      vtep-56     198.51.100.22 Mlt130     3855      underlay-vxlan

-----
2 out of 2 Total Num of remote vteps
-----
```

Display the Name of the Remote VTEP

Use the following procedure to display the name of the remote VTEP.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the names of all remote VTEPs or use the option to specify a particular VTEP:
show vtep remote name [<vtepId>]

Example

```
Switch:1# show vtep remote name
=====
Remote VTEP Name
=====
ID      NAME
-----
1       vxlan_tunnel_1
2       --
3       --

-----
3 out of 3 Total Num of remote vteps
-----
```

Display VNID Information

Use the following procedure to display information about the VNID.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about all the VNIDs or use the option to specify a particular VNID:
show vnid i-sid [<1-500>]

Example

```
Switch:1# show vnid i-sid
=====
vnid Info
=====
```

VNID ID	ISID ID	PORT INTERFACES	MLT INTERFACES	VTEP IDS
110	120	c3:1/10	-	2,56
11000060	1060	c60:1/10	-	56
11000065	1065	c65:1/10	-	56
11000066	1066	c66:1/48	-	2
.				
.				
.				

Display the MAC Addresses in the VNID FDB

Use the following procedure to display the MAC addresses in the VNID forwarding database.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display all the VNIDs in the FDB table or use one of the options to specify a particular VNID:
show vnid mac-address-entry [<1-16777215> | port <{slot/
port[/sub-port]}[-slot/port[/sub-port]][,...][slot/all][all]]> | mac
<0x00:0x00:0x00:0x00:0x00:0x00> | remote]

Example

```
Switch:1# show vnid mac-address-entry
```

VNID Fdb Table				
VNID	STATUS	MAC-ADDRESS	INTERFACE	TYPE
11000065	learned	00:00:5e:00:01:01	vtep_56	REMOTE
14000100	learned	a4:25:1b:51:a5:04	vtep_67	REMOTE
14000100	learned	a4:25:1b:52:a5:04	vtep_67	REMOTE
14000100	learned	a4:25:1b:52:bd:03	vtep_2	REMOTE

```
c: customer vid    u: untagged-traffic

All 4 out of 4 Total Num of vnid FDB Entries displayed
```

Display the VXLAN Running Configuration

Use the following procedure to display the currently running VXLAN configuration.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the running configuration:
show running-config module vxlan

Example

This example shows the VXLAN information only; it does not show the hardware-specific information that the system also displays in the output for this command.

The example shows the VXLAN Gateway in Base Interworking Mode. It also shows an external loopback (see **c-vid** commands), which enables communication between the VNID and an I-SID in an SPB network.

```
Switch:1# show running-config module vxlan

#
# VTEP CONFIGURATION
#
vtep source-ip 1.1.1.1 vrf underlay-vxlan

#
# REMOTE VTEP CONFIGURATIONS
#
vtep 2 ip 1.1.1.1 name "vtep-2"
vtep 67 ip 1.1.1.67 name "vtep-67"

#
# VNID CONFIGURATION
#
vnid 12002105 i-sid 102105
c-vid 105 port 1/11
c-vid 2105 port 1/10
vtep 67
exit
```

OVSDB protocol support for VXLAN Gateway Configuration Using CLI

This section provides procedures to configure OVSDB protocol support for VXLAN Gateway using the Command Line Interface (CLI).

Configure OVSDB Managed Interfaces

Use the following procedure to configure OVSDB managed interfaces on the switch.

Before You Begin

- You must enable VXLAN Gateway Full Interworking Mode. For more information, see [Enable VXLAN Gateway in Full Interworking Mode](#) on page 4183.
- You must create the required Flex UNI or CVLAN I-SIDs before you configure the OVSDB managed interface.

Procedure

1. Enter OVSDB Configuration mode:
enable

configure terminal

ovsdb
2. Configure the OVSDB managed interface:
managed-interface i-sids WORD<1-1024>

Example

The following is an example configuring two I-SIDs as OVSDB management interfaces.

```
SWITCH:1>enable
SWITCH:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SWITCH:1(config)#ovsdb
SWITCH:1(config-ovsdb)#managed-interface i-sids 300
SWITCH:1(config-ovsdb)#managed-interface i-sids 500
```

What to Do Next

Configure the OVSDB protocol support for VXLAN Gateway, see [Configure OVSDB protocol support for VXLAN Gateway](#) on page 4191.

Variable Definitions

Use the data in the following table to use the **managed-interface** command.

Variable	Value
<i>i-sids</i> WORD<1-1024>	Specifies the I-SID information to create an interface to manage OVSDB protocol support for VXLAN Gateway. Note: You must create the required Flex UNI or CVLAN I-SIDs before you configure the managed interface I-SID.

Configure OVSDB protocol support for VXLAN Gateway

Use the following procedure to configure OVSDB protocol support for VXLAN Gateway on the switch.

Before You Begin

- You must enable VXLAN Gateway Full Interworking Mode. You can use **show boot config flags** to verify the current VXLAN Gateway mode. For more information, see [Enable VXLAN Gateway in Full Interworking Mode](#) on page 4183.
- You must configure and use the Segmented Management Instance IP address on the VXLAN Gateway to establish connectivity with the NVC. For more information about Segmented Management Instance, see [Segmented Management Instance Configuration using the CLI](#) on page 123.
- You must configure an OVSDB management interface. For more information, see [Configure OVSDB Managed Interfaces](#) on page 4190.
- You must transfer an OVSDB certificate file and private-key file to the flash storage of the switch. You can use an ovs-pki utility with SSL libraries to generate the private keys and certificates. You can use **boot config flags FTPD** and then an SCP utility to transfer the private key and certificate file to the flash storage of the switch.
- If the switch is an aggregation switch, the IST peer must support OVSDB, have the same source VTEP-IP and OVSDB managed-interface, and must communicate with the NVC management IP.

Procedure

1. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```
2. Configure the source VTEP IP address:

```
vtep source-ip <A.B.C.D> [vrf WORD<1-16>]
```
3. Enter OVSDDB Configuration mode:

```
ovsdb
```
4. Install the OVSDDB certificate file:

```
install-cert-file WORD<1-128>
```
5. Install the OVSDDB private key:

```
private-key WORD<1-128>
```
6. Enable OVSDDB protocol support for VXLAN Gateway:

```
enable
```
7. Configure the NVC:

```
controller <1-100> ip address <A.B.C.D> protocol <ssl|tcp> [port <1-65535>]
```

Example

The following is an example of configuring a VTEP source IP, installing an OVSDDB certificate file, installing an OVSDDB private key, enabling OVSDDB protocol support for VXLAN Gateway, and configuring the IP address, protocol, and port for one NVC.

```
SWITCH:1>enable
SWITCH:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SWITCH:1(config)#vtep source-ip 192.0.2.1 vrf vxlan-underlay
SWITCH:1(config)#ovsdb
SWITCH:1(config-ovsdb)#install-cert-file /intflash/tom/sc-cert.pem
SWITCH:1(config-ovsdb)#private-key /intflash/tom/sc-privkey.pem
SWITCH:1(config-ovsdb)#enable
SWITCH:1(config-ovsdb)#controller 1 ip address 192.0.2.2 protocol ssl port 6640
```



Note

You can configure multiple controllers for high availability. One VXLAN Gateway can support a maximum of three controllers.



Important

If you add or delete a controller, or modify the OVSDDB managed interface when a controller is configured, the existing controller connections reset. Log messages generate to indicate the status changes as the controllers disconnect and reconnect.

If you change a previously configured VTEP source-ip and re-enable OVSDDB, the controller sees a new VXLAN tunnel instead of updating the existing VXLAN. You must configure the VNID to I-SID binding on the controller for the new VXLAN tunnel associated with the new VTEP IP address.

Variable Definitions

Use the data in the following table to use the **vtep source-ip** command.

Variable	Value
<code><A.B.C.D> [vrf WORD<1-16>]</code>	Specifies the VXLAN tunnel end point (VTEP) source IP address in IPv4 format. Optionally you can specify a VRF. Note: The VTEP source IP address must be on a loopback interface.

Use the data in the following table to use the **install-cert-file** command.

Variable	Value
<code>WORD<1-128></code>	Specifies the path and file name of the OVSDB certificate file.

Use the data in the following table to use the **private-key** command.

Variable	Value
<code>WORD<1-128></code>	Specifies the path and file name of the OVSDB private key.

Use the data in the following table to use the **controller** command.

Variable	Value
<code><1-100></code>	Specifies the ID of the controller.
<code>ip address <A.B.C.D></code>	Specifies the IP address of the controller in IPv4 format.
<code>protocol <ssl tcp></code>	Specifies the networking protocol as SSL or TCP for controller communications.
<code>port <1-65535></code>	Specifies the networking port of the controller.

Configure OVSDB Replication

Use the following procedure to configure OVSDB replication to support vIST.

Before You Begin

- You must enable VXLAN Gateway Full Interworking Mode on the switch.
- You must configure an OVSDB managed interface on the switch.
- You must configure and enable OVSDB protocol support for VXLAN Gateway on the switch.
- You must have at least two NVCs configured to communicate on the managed interface network.

Procedure

1. Enter OVSDb Configuration mode:

```
enable
```

```
configure terminal
```

```
ovsdb
```
2. Configure the OVSDb replication:

```
replication peer-ip <A.B.C.D> local-ip <A.B.C.D>
```

Example

The following is an example of configuring OVSDb replication to support vIST.

```
SWITCH:1>enable
SWITCH:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SWITCH:1 (config)#ovsdb
SWITCH:1 (config-ovsdb)#replication peer-ip 192.0.2.3 local-ip 192.0.2.2
```

Variable definitions

Use the data in the following table to use the **replication** command.

Variable	Value
<i>peer-ip</i> <A.B.C.D>	Specifies the IP address of the secondary controller for OVSDb replication support.
<i>local-ip</i> <A.B.C.D>	Specifies the IP address of the primary controller for OVSDb replication support.

Display the OVSDb Configuration

Use the following procedure to display the OVSDb protocol support for VXLAN Gateway configuration.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the OVSDb configuration:

```
show ovsdb config
```

Example

The following is an example of displaying the OVSDb configuration.

```
SWITCH:1>show ovsdb config
=====
                        Ovldb Config Info
=====
                        Status - Enabled
                        Certificate - Installed
                        Private Key - Installed
                        Replication Peer NLS-IP - 192.0.2.3
                        Replication Local NLS-IP - 192.0.2.2
```


Ovsdb Controller Config Info			
ID	Protocol	IP	Port
1	ssl	192.0.2.10	6640

Display the OVSDb Controller Status

Use the following procedure to display the OVSDb controller status.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the OVSDb controller status:

```
show ovsdb controller status
```

Example

The following is an example of displaying the OVSDb controller status.

```
SWITCH:1>show ovsdb controller status
```

OVSDb Controller Status				
Protocol	IP	Port	status	source
ssl	192.0.2.2	6640	Up	Learnt
ssl	192.0.2.3	6640	Down	Learnt
2 out of 2 Total Num of Controllers displayed				

Display the OVSDb Replication State

Use the following procedure to display the OVSDb replication state.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the OVSDb controller status:

```
show ovsdb replication state
```

Example

The following is an example of displaying the OVSDb replication state.

```
SWITCH:1>show ovsdb replication state
```

Ovsdb Replication Status		
Local NLS IP	Peer NLS IP	State
192.0.2.2	192.0.2.3	ENABLED

Display the OVSDb Managed Interfaces

Use the following procedure to display the OVSDb managed interfaces.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display the OVSDb managed interfaces:

```
show ovssdb managed-interface
```

Example

The following is an example of displaying the OVSDb managed interfaces.

```
SWITCH:1>show ovssdb managed-interface
=====
OVSDb Managed Interfaces - ISID
=====
50001
1 out of 1 Total Num of Managed Interfaces displayed
```

VXLAN Gateway configuration using EDM

This section provides procedures to configure VXLAN Gateway using the Enterprise Device Manager (EDM).

Configure the VXLAN Gateway Boot Flag

There are two configuration modes: Base Interworking Mode and Full Interworking Mode. The Base Interworking Mode is the default mode. To switch to the Full Interworking Mode, use this procedure.

**Important**

After you change this boot flag, you must save the change to the configuration file and reboot the switch.

Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Boot Config** tab.
3. Select **EnableVxlanGwFullInterworkingMode** to enable Full Interworking Mode.
4. Select **Apply**.

Configure a VTEP Source IP Address

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Globals** tab.
4. In **SourceIp**, type the VTEP source IP address.
5. In **Vrf**, select `GlobalRouter` or a VRF ID.
6. Select **Apply**.

Globals *Field Descriptions*

Use the data in the following table to use the **Globals** tab.

Name	Description
SourceIp	Specifies the VXLAN tunnel end point (VTEP) source IP address in an A.B.C.D format.
Vrf	<p>Specifies the VRF name used for the source IP address, which you select from the drop down menu. This IP address can be on the GRT or a VRF.</p> <p>Note: The VTEP source IP address must be on a loopback interface.</p>

Configure a Remote VTEP

Use the following procedure to configure the remote VTEP destination.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Vtep** tab.
4. In **Id**, assign a unique index value to this remote VTEP.
5. In **IpAddr**, enter the IP address of this remote VTEP.
6. In **Name**, assign a name to this remote VTEP.
7. Select **Apply**.

Vtep *Field Descriptions*

Use the data in the following table to use the **Vtep** tab.

Name	Description
Id	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
IpAddr	<p>Specifies an IP address that uniquely identifies this remote VTEP.</p> <p>Note: The remote VTEP IP address cannot be a local, broadcast, or multicast IP address.</p>
Name	Assigns a name to this remote VTEP tunnel.

Name	Description
NextHopVrfName	Displays the VRF name of the next hop to reach the remote VTEP.
OvsdbConfigured	Specifies if the VTEP is configured with OVSDB protocol support for VXLAN Gateway.

Configure a VNID

Use the following procedure to configure a VNID, which identifies one of the potential 16 million VXLAN segments that can coexist within the same administrative domain.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Vnid** tab.
4. In **Identifier**, assign a unique value to this VNID.
5. In **Isid**, specify the I-SID associated with this VNID.
6. In **Action**, select one of the available actions: none or flush.
7. Select **Apply**.

Vnid Field Descriptions

Use the data in the following table to use the **Vnid** tab.

Name	Description
Identifier	Specifies a value that uniquely identifies the VNID of this VXLAN tunnel end point.
Isid	<p>Specifies a value that uniquely identifies the I-SID associated with this VXLAN tunnel end point.</p> <p>Note: A VNID must not have the same value as an I-SID.</p>
Action	<p>Choose one of the following actions to take with the MAC FDB:</p> <ul style="list-style-type: none"> • none: Don't take any action. This is the default. • flushMacFdb: Flush all the learned MAC addresses from the forwarding database of the selected VNID.
OvsdbConfigured	Specifies if the VNID is configured with OVSDB protocol support for VXLAN Gateway.

Configure VNID Endpoints

Use the following procedure to configure a VNID and the VXLAN tunnel end point.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Vnid End Points** tab.
4. In **Vnid**, type the VNID ID.
5. In **Vtepld**, type the ID of the remote VTEP.
6. Select **Apply**.

Vnid End Points Field Descriptions

Use the data in the following table to use the **Vnid End Points** tab.

Name	Description
Vnid	Specifies a value that uniquely identifies the VNID of this VXLAN tunnel end point.
Vtepld	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
Isid	Displays a value that uniquely identifies the I-SID associated with this VXLAN tunnel end point.

Configure ELAN Endpoints

Use the following procedure to configure information associated with ELAN end points.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Elan End Points** tab.
4. In **Vnid**, specify the VNID associated with this ELAN.
5. In **Cvid**, specify the customer VLAN ID.
6. In **Ifindex**, specify the interface index for this ELAN endpoint.
7. Select **Apply**.

Elan End Points Field Descriptions

Use the data in the following table to use the **Elan End Points** tab.

Name	Description
Vnid	Specifies a value that uniquely identifies the VNID of this ELAN tunnel end point.
Cvid	<p>Specifies a value that uniquely identifies the customer VLAN ID of this ELAN endpoint.</p> <p>The switch reserves the following VLAN IDs:</p> <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
IfIndex	Specifies the interface index for the MLT or port for this ELAN endpoint.
Isid	Specifies the I-SID associated with this VNID.

Display the VTEP Next Hop

Use the following procedure to display information about the remote VTEP.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Next Hop** tab.

Next Hop Field Descriptions

Use the data in the following table to use the **Next Hop** tab.

Name	Description
VtepId	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
Ip	Specifies an IP address that uniquely identifies this remote VTEP tunnel.
IfIndex	Specifies the interface index of the next hop to reach the remote VTEP.
Vid	Specifies the VLAN ID of the next hop to reach remote VTEP.

Display the VNID Forwarding Database

Use the following procedure to display information about the VNID forwarding database (FDB).

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **VTEP**.
3. Select the **Vnid FDB** tab.

Vnid FDB Field Descriptions

Use the data in the following table to use the **Vnid FDB** tab.

Name	Description
Vnid	Displays the VNID IDs of the VXLAN tunnel endpoints.
Address	Displays all the learned MAC addresses in the forwarding database of the selected VNID.
Status	Displays whether or not the MAC address was learned.
InterfaceIndex	Displays the interface index name for the VNID.
Type	Displays whether the VNID is local or remote.

OVSDB protocol support for VXLAN Gateway configuration using EDM

This section provides procedures to configure OVSDB protocol support for VXLAN Gateway using the Enterprise Device Manager (EDM).

Configure OVSDB Globally

Use the following procedure to configure the OVSDB certificates, private key, and replication global settings.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **OVSD**.
3. Select the **Globals** tab.
4. Select **Enable**.
5. In **CertificateFilename**, type the file path/file name.
6. In **CertFileInstallAction**, select **install**.
7. Select **Apply**.
8. In **PrivateKeyFilename**, type the file path/file name.
9. In **PrivateKeyInstallAction**, select **install**.

10. Select **Apply**.
11. Select **ReplicationEnable**.
12. In **ReplicationPeerIpAddr**, type the IP address.
13. In **ReplicationLocalIpAddr**, type the IP address.
14. Select **Apply**.
15. Verify the **ReplicationState**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
Enable	Enables or disables the feature. Default is disabled (cleared).
CertificateFilename	Specifies the file name of the certificate.
CertFileInstallAction	Specifies to install or uninstall the certificate. You must have a certificate file on the switch.
PrivateKeyFilename	Specifies the file name of the private key.
PrivateKeyInstallAction	Specifies to install or uninstall the private key. You must have a private key file on the switch.
ReplicationEnable	Enables or disables replication support. Default is disabled (cleared).
ReplicationPeerIpAddr	Specifies the replication server peer IP address in A.B.C.D format.
ReplicationLocalIpAddr	Specifies the replication local IP address in A.B.C.D format.
ReplicationState	Displays the current replication status.

Configure an OVSDB Controller

Use the following procedure to configure an OVSDB controller ID, IP address, protocol, and port.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **OVSDB**.
3. Select the **Controller** tab.
4. Select **Insert**.
5. In **Id**, assign a unique index value to this controller.
6. In **IpAddr**, enter the IP address of this controller.
7. In **Protocol**, select **tcp** or **ssl**.
8. In **Port**, enter a port number.
9. Select **Insert**.

Controller Field Descriptions

Use the data in the following table to use the **Controller** tab.

Name	Description
Id	Specifies an ID number for a controller. Value range of 1 to 100.
IpAddr	Specifies the IP address of a controller. Type an IP address in A.B.C.D format.
Protocol	Specifies the communications protocol for the controller. Select tcp or ssl .
Port	Specifies the communications port for the controller. Value range of 1 to 65535.

Display the OVSDb Controller Status

Use the following procedure to display status information about the OVSDb controller.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **OSVDB**.
3. Select the **Status** tab.

Status Field Descriptions

Use the data in the following table to use the **Status** tab.

Name	Description
IpAddress	Displays the IP address of a controller.
Protocol	Displays the communications protocol of a controller as TCP or SSL.
Port	Displays the communications port for the controller. Value range of 1 to 65535.
Source	Displays the source of the controller as learned or configured.
OperStatus	Displays the operational status of the controller.

Configure OVSDb Managed Interface

Use the following procedure to configure a I-SID interface for the OVSDb.

Procedure

1. In the navigation pane, expand **Configuration > Edit**.
2. Select **OSVDB**.
3. Select the **Interface** tab.

4. Select **Insert**.
5. In **Isid**, type the I-SID identifier.
6. Select **Insert**.

Interface Field Descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Isid	Specifies an I-SID to use as the managed interface.

Configuration Examples

This section provides network configuration examples for VXLAN Gateway and OVSDB protocol support for VXLAN Gateway.

VXLAN Gateway Configuration Example

This example shows how to configure the VXLAN Gateway. Remember that there are two modes:

- Base Interworking Mode is the default mode so there is no need to set the mode boot flag. However, you should use the **show boot config flags** command to verify that the **vxlan-gw-full-interworking-mode** is set to *false*.
- Full Interworking Mode supports all SPB functionality including vIST and SMLT using an internal loopback. Full Interworking Mode is not the default mode. To configure Full Interworking mode, you must enable the **vxlan-gw-full-interworking-mode** boot flag.



Important

Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.




Figure 260: VXLAN Gateway configuration example

```
#
#SPBM CONFIGURATION
#
spbm
spbm ethertype 0x8100

#
#VLAN CONFIGURATION
#
vlan members remove 1 1/1,1/30,2/1/1,2/1/4
vlan create 3000 type port-mstprstp 0
vlan members 3000 1/30 portmember
interface Vlan 3000
ip address 192.168.30.1 255.255.255.0 0
ip ospf enable
exit

#
#PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 0
flex-uni enable

#
#CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#
interface loopback 1
ip address 1 203.0.113.3/255.255.255.255
ip ospf 1
exit

#
#VTEP CONFIGURATION
#
vtep source-ip 198.51.100.1

#
#REMOTE VTEP CONFIGURATIONS
#
vtep 1 ip 203.0.113.3
```

```
#
#OSPF CONFIGURATION - GlobalRouter
#
router ospf enable
router ospf
exit

#
#IP REDISTRIBUTION CONFIGURATION - GlobalRouter
#
router ospf
redistribute direct
redistribute direct enable
exit

#
#I-SID CONFIGURATION
#
i-sid 300 elan
c-vid 10 port 1/1
exit

#
#VNID CONFIGURATION
#
vnid 100 i-sid 300
vtep 1
exit

#
#IP REDISTRIBUTE APPLY CONFIGURATIONS
#
ip ospf apply redistribute direct
```

OVSDB protocol support for VXLAN Gateway Configuration Example

This example shows how to configure OVSDB protocol support for VXLAN Gateway.



Important

OVSDB protocol support for VXLAN Gateway also requires an ESXi host running VMware NSX configured with a HW-VTEP, and a Logical Switch configured with VNID to HW-VTEP bindings to function. See VMware NSX documentation for more information.



Important

OVSDB requires full interworking mode on the VXLAN Gateway. Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.

```
#
#VXLAN GATEWAY UNDERLAY IP NETWORK CONFIGURATION HW-VTEP 1
#
enable
config terminal
ip vrf underlay-vxlan vrfid 1
router vrf underlay-vxlan
ip ospf
ip ospf admin-state
ip ospf router-id 203.0.113.1
```

```
vlan create 10 type port-mstprstp 1
vlan mlt 10 10
interface vlan 10
vrf underlay-vxlan
ip address 198.51.100.122 255.255.255.0
ip ospf enable
exit

#
#VXLAN GATEWAY UNDERLAY IP NETWORK CONFIGURATION HW-VTEP 2
#
ip vrf underlay-vxlan vrfid 1
router vrf underlay-vxlan
ip ospf
ip ospf admin-state
ip ospf router-id 203.0.113.2

vlan create 12 type port-mstprstp 1
vlan members add 12 1/1 portmember
interface vlan 12
vrf underlay-vxlan
ip address 198.51.100.124 255.255.255.0
ip ospf enable
exit
#
#NLS and VTEP CONFIGURATIONS HW-VTEP 1
#
vlan create 4059 type port-mstprstp 0
vlan members 4059 1/24 portmember
mgmt vlan 4059
ip address 192.0.2.101/24
ip route 192.0.2.0/16 next-hop 192.0.2.1 weight 1
enable
exit

interface loopback 10
ip address 10 122.122.122.1/255.255.255.255 vrf underlay-vxlan
ip ospf 10 vrf underlay-vxlan

vtep source-ip 122.122.122.1 vrf underlay-vxlan

#
#NLS and VTEP CONFIGURATIONS HW-VTEP 2
#
vlan create 4059 type port-mstprstp 0
vlan members 4059 1/24 portmember
mgmt vlan 4059
ip address 192.0.2.102/24
ip route 192.0.2.0/16 next-hop 192.0.2.1 weight 1
enable
exit

interface loopback 10
ip address 10 124.124.124.124/255.255.255.255 vrf underlay-vxlan
ip ospf 10 vrf underlay-vxlan

vtep source-ip 124.124.124.1 vrf underlay-vxlan

#
#LAYER2 VNI SERVICE AND MANAGED-INTERFACE CONFIGURATIONS HW-VTEP 1
#
vlan create 1001 type port-mstprstp 0
vlan members 1001 1/10 portmember
```

```
vlan i-sid 1001 1001

ovsdb
managed-interface i-sids 1001
exit

#
##LAYER2 VNI SERVICE AND MANAGED-INTERFACE CONFIGURATIONS HW-VTEP 2
#
vlan create 1001 type port-mstprstp 0
vlan members 1001 1/10 portmember
vlan i-sid 1001 1001

ovsdb
managed-interface i-sids 1001
exit

#
#OVSDb CERTIFICATE AND PRIVATE KEY CONFIGURATION HW-VTEP 1
#
ovsdb
install-cert-file /intflash/tom/vtep1-cert.pem
private-key /intflash/tom/vtep1-privkey.pem
enable
exit

#
#OVSDb NETWORK VIRTUALIZATION CONTROLLER CONFIGURATION HW-VTEP 1
#
ovsdb
controller 1 ip address 192.0.2.2 protocol ssl port 6640
exit

#
#OVSDb CERTIFICATE AND PRIVATE KEY CONFIGURATION HW-VTEP 2
#
#
ovsdb
install-cert-file /intflash/tom/vtep2-cert.pem
private-key /intflash/tom/vtep2-privkey.pem
enable
exit

#
#OVSDb NETWORK VIRTUALIZATION CONTROLLER CONFIGURATION HW-VTEP 2
#
ovsdb
controller 1 ip address 192.0.2.2 protocol ssl port 6640
exit

#If HW-VTEP 1 needs VXLAN Gateway redundancy, the redundant gateway must be a VIST
peer(HW-VTEP 11).
#For VXLAN Gateway redundancy, the configs at HW-VTEP-1 must also have ovsdb replication
config:
#
#OVSDb NETWORK VIRTUALIZATION CONTROLLER REPLICATION CONFIGURATION
#
ovsdb
replication peer-ip 192.0.2.101 local-ip 192.0.2.111
exit
#
```

```
#Configs at HW-VTEP-11 for reduncancy:

#
#VXLAN GATEWAY UNDERLAY IP NETWORK CONFIGURATION HW-VTEP 11
#
enable
config terminal
ip vrf underlay-vxlan vrfid 1
router vrf underlay-vxlan
ip ospf
ip ospf admin-state
ip ospf router-id 203.0.113.11

vlan create 10 type port-mstprstp 1
vlan mlt 10 10
interface vlan 10
vrf underlay-vxlan
ip address 198.51.100.111 255.255.255.0
ip ospf enable
exit

#
#NLS and VTEP CONFIGURATIONS HW-VTEP 11
#NOTE: HW-VTEP 1 and HW-VTEP 11's VTEP IP must be the same.
#
vlan create 4059 type port-mstprstp 0
vlan members 4059 1/24 portmember
mgmt vlan 4059
ip address 192.0.2.111/24
ip route 192.0.2.0/16 next-hop 192.0.2.1 weight 1
enable
exit

interface loopback 10
ip address 10.122.122.1/255.255.255.255 vrf underlay-vxlan
ip ospf 10 vrf underlay-vxlan

vtep source-ip 122.122.122.1 vrf underlay-vxlan

#
#LAYER2 VNI SERVICE AND MANAGED-INTERFACE CONFIGURATIONS HW-VTEP 11
#NOTE: HW-VTEP 1 and HW-VTEP 11 must have same managed-interface configured under 'ovsdb'
#
vlan create 1001 type port-mstprstp 0
vlan members 1001 1/10 portmember
vlan i-sid 1001 1001

ovsdb
managed-interface i-sids 1001
exit

#
#OVSDb CERTIFICATE AND PRIVATE KEY CONFIGURATION HW-VTEP 11
#NOTE: HW-VTEP 1 and HW-VTEP 11 must have same certificate and private keys configured.
#
ovsdb
install-cert-file /intflash/tom/vtep1-cert.pem
private-key /intflash/tom/vtep1-privkey.pem
enable
exit

#
#OVSDb NETWORK VIRTUALIZATION CONTROLLER CONFIGURATION HW-VTEP 11
#
```

```
ovsdb
controller 1 ip address 192.0.2.2 protocol ssl port 6640
exit

#
#OVSDb NETWORK VIRTUALIZATION CONTROLLER REPLICATION CONFIGURATION
#
ovsdb
replication peer-ip 192.0.2.101 local-ip 192.0.2.111
exit
#
```




Port Numbering and MAC Address Assignment Reference

[Port Numbering](#) on page 4211

[Interface Indexes](#) on page 4242

[MAC Address Assignment](#) on page 4244

This section provides information about the port numbering and Media Access Control (MAC) address assignment used on the switch.

Port Numbering

A port number includes the slot location of the port in the chassis, as well as the port position, using the format slot/port. For example, the second port in the first slot is structured as 1/2. The number of slots and ports varies depending on the hardware platform. For more information about hardware, see the hardware documentation for your platform.

4220 Series

The following diagrams illustrate the components on the front panels of the 4220 Series switches.

4220-12P-4X and 4220-12T-4X

The following diagram illustrates the components of the front panel of the 4220-12P-4X and 4220-12T-4X. The only difference between these two models is PoE support. All ports are in slot 1.




Figure 261: 4220-12P-4X and 4220-12T-4X front view

Table 358: Figure call outs

Callout	4220-12P-4X	4220-12T-4X
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).
2	System LEDs	System LEDs
3	USB Type A port	USB Type A port
4	10/100/1000BASE-T out-of-band management port	10/100/1000BASE-T out-of-band management port
5	12 x 10/100/1000Base-T 802.3at (30W) PoE ports	12 x 10/100/1000Base-T ports
7	4 x 1/10Gb SFP+ uplink ports	4 x 1/10Gb SFP+ uplink ports
8	USB Micro-B console port	USB Micro-B console port
9	Serial console port (RJ-45)	Serial console port (RJ-45)

4220-24P-4X and 4220-24T-4X

The following diagram illustrates the components of the front panel of the 4220-24P-4X and 4220-24T-4X. The only difference between these two models is PoE support. All ports are in slot 1.




Figure 262: 4220-24P-4X and 4220-24T-4X front view

Table 359: Figure call outs

Callout	4220-24P-4X	4220-24T-4X
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).
2	System LEDs	System LEDs
3	USB Type A port	USB Type A port
4	10/100/1000BASE-T out-of-band management port	10/100/1000BASE-T out-of-band management port
5	24 x 10/100/1000Base-T 802.3at (30W) PoE ports	24 x 10/100/1000Base-T ports
7	4 x 1/10Gb SFP+ uplink ports	4 x 1/10Gb SFP+ uplink ports
8	USB Micro-B console port	USB Micro-B console port
9	Serial console port (RJ-45)	Serial console port (RJ-45)

4220-4MW-8P-4X

The following diagram illustrates the components of the front panel of the 4220-4MW-8P-4X. All ports are in slot 1.




Figure 263: 4220-4MW-8P-4X front view

Table 360: Figure call outs


Callout	4220-4MW-8P-4X
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).
2	System LEDs

Table 360: Figure call outs (continued)

Callout	4220-4MW-8P-4X
3	USB Type A port
4	10/100/1000BASE-T out-of-band management port
5	8 x 10/100/1000Base-T 802.3at (30W) PoE ports
6	4 x 1G/2.5G/5GBaseT 802.3bt (90W) PoE ports
7	4 x 1/10Gb SFP+ uplink ports
8	USB Micro-B console port
9	Serial console port (RJ-45)

4220-4MW-20P-4X

The following diagram illustrates the components of the front panel of the 4220-4MW-20P-4X. All ports are in slot 1.

**Figure 264: 4220-4MW-20P-4X front view****Table 361: Figure call outs**

Callout	4220-4MW-20P-4X
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).
2	System LEDs
3	USB Type A port
4	10/100/1000BASE-T out-of-band management port
5	20 x 10/100/1000Base-T 802.3at (30W) PoE ports
6	4 x 1G/2.5G/5GBaseT 802.3bt (90W) PoE ports
7	4 x 1/10Gb SFP+ uplink ports
8	USB Micro-B console port
9	Serial console port (RJ-45)

4220-8X

The following diagram illustrates the components of the front panel of the 4220-8X. All ports are in slot 1.




Figure 265: 4220-8X front view

Table 362: Figure call outs

Callout	4220-8X
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).
2	System LEDs
3	USB Type A port
4	10/100/1000BASE-T out-of-band management port
7	8 x 1/10Gb SFP+ uplink ports
8	USB Micro-B console port
9	Serial console port (RJ-45)

5320 Series

The following diagrams illustrate the components on the front panels of the 5320 Series switches.

5320-16P-4XE and 5320-16P-4XE-DC

The following diagram illustrates the components of the front panel of the 5320-16P-4XE and 5320-16P-4XE-DC switches. The only difference between these switches is the power source. All ports are in slot 1.




Figure 266: 5320-16P-4XE Front Panel




Figure 267: 5320-16P-4XE-DC Front Panel

Table 363: Figure call outs

Callout	5320-16P-4XE and 5320-16P-4XE-DC
1	RJ-45 serial console port
2	USB Type-A port
3	USB Micro-B console port
4	10/100/1000BASE-T ports — These ports support 802.3af Type 1 PoE and 802.3at Type 2 PoE
5	1/10Gb SFP+ uplink ports
6	1/10Gb SFP+ Universal Ports

The back panel (not shown) includes:

- Grounding lug
- 1 AC or DC power inlet connector

5320-16P-2MXT-2X

The following diagram illustrates the components of the front panel of the 5320-16P-2MXT-2X switch. All ports are in slot 1.




Figure 268: 5320-16P-2MXT-2X

Table 364: Figure call outs

Callout	5320-16P-2MXT-2X
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD)
2	System LEDs
3	10/100/1000BASE-T copper RJ45 full/half duplex ports

Table 364: Figure call outs (continued)

Callout	5320-16P-2MXT-2X
4	1G/2.5G/5G/10GBASE-T copper RJ45 uplink ports
5	1G/10G SFP+ uplink ports
6	RJ-45 Serial console port
7	Type A USB port
8	Type C USB console port

The back panel (not shown) includes:

- Grounding lug
- 1 AC power inlet connector

5320-24P-8XE and 5320-24T-8XE

The following diagram illustrates the components of the front panel of the 5320-24P-8XE and 5320-24T-8XE switches. All ports are in slot 1.




Figure 269: 5320-24P-8XE Front Panel




Figure 270: 5320-24T-8XE Front Panel

Table 365: Figure call outs

Callout	5320-24P-8XE	5320-24T-8XE
1	RJ-45 serial console port	RJ-45 serial console port
2	USB Type-A port	USB Type-A port
3	USB Micro-B console port	USB Micro-B console port
4	10/100/1000BASE-T ports — These ports support 802.3af Type 1 PoE and 802.3at Type 2 PoE+	10/100/1000BASE-T ports

Table 365: Figure call outs (continued)


Callout	5320-24P-8XE	5320-24T-8XE
5	1/10Gb SFP+ uplink ports	1/10Gb SFP+ uplink ports
6	1/10Gb SFP+ Universal Ports	1/10Gb SFP+ Universal Ports

The back panel (not shown) includes:

- Grounding lug
- 1 AC power inlet connector

5320-24T-4X-XT

The following diagram illustrates the components of the front panel of the 5320-24T-4X-XTswitch. All ports are in slot 1.

**Figure 271: 5320-24T-4X-XT Front Panel****Table 366: Figure call outs**

Callout	5320-24T-4X-XT
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD).
2	System LEDs
3	10/100/1000BASE-T out-of-band management port (RJ-45)
4	24 x 10/100/1000BASE-T (autosensing) full/half duplex ports (not MACsec capable) Note: The 10/100 ports support half duplex; the 1000 port supports full duplex. Half-duplex is not supported on these ports when operating at 1Gbps.
5	4 x 1/10Gb SFP+ Universal ports (unpopulated) (not MACsec capable)
7	USB Micro-B console port
8	USB Type-A port for management or external USB flash
9	Serial console port (RJ-45)

The back panel (not shown) includes:

- AC power inlet connector
- Redundant power inlet connector (behind cover)
- Grounding lug

5320-24T-24S-4XE-XT

The following diagram illustrates the components of the front panel of the 5320-24T-24S-4XE-XT switch. All ports are in slot 1.




Figure 272: 5320-24T-24S-4XE-XT Front Panel

Table 367: Figure call outs

Callout	5320-24T-24S-4XE-XT
1	Mode button for port LED control. Port LED mode indicators: system (SYS) and speed (SPD)
2	System LEDs
3	10/100/1000BASE-T out-of-band management port (RJ-45)
4	24 x 10/100/1000BASE-T (autosensing) full/half duplex MACsec capable ports Note: The 10/100 ports support half duplex; the 1000 port supports full duplex. Half-duplex is not supported on these ports when operating at 1Gbps.
5	4 x 1/10Gb SFP+ MACsec capable Universal ports (unpopulated)
6	24 x 100Mb/1G SFP MACsec capable ports
7	USB Micro-B console port
8	USB Type-A port for management or external USB flash
9	Serial console port (RJ-45)

The back panel (not shown) includes:

- AC power inlet connector
- Redundant power inlet connector (behind cover)
- Grounding lug

5320-48P-8XE and 5320-48T-8XE

The following diagram illustrates the components of the front panel of the 5320-48P-8XE and 5320-48T-8XE switches. All ports are in slot 1.




Figure 273: 5320-48P-8XE Front Panel




Figure 274: 5320-48T-8XE Front Panel

Table 368: Figure call outs

Callout	5320-48P-8XE	5320-48T-8XE
1	RJ-45 serial console port	RJ-45 serial console port
2	USB Type-A port	USB Type-A port
3	USB Micro-B console port	USB Micro-B console port
4	10/100/1000BASE-T ports — These ports support 802.3af Type 1 PoE and 802.3at Type 2 PoE+	10/100/1000BASE-T ports
5	1/10Gb SFP+ uplink ports	1/10Gb SFP+ uplink ports
6	1/10Gb SFP+ Universal Ports	1/10Gb SFP+ Universal Ports

The back panel (not shown) includes:

- Grounding lug
- 1 AC power inlet connector

5420 Series

In addition to the regular fixed ports, all 5420 Series switches have two SFP-DD Universal Ethernet ports that are reserved for advanced features (factory default). You can also use these ports as regular ports if you do not require the advanced features. For more information about using the advanced-feature-bandwidth-reservation boot flag to configure the functionality of these ports, see [Advanced Feature Bandwidth Reservation](#) on page 275.

When advanced features are disabled on these SFP-DD Universal Ethernet ports, they can function as regular Ethernet ports, supporting data rates of either 10Gb using SFP+ optics or 20Gb using SFP-DD optics. 5420M switch models support two 10Gb channels on each SFP-DD port. 5420F switch models support one 10Gb channel on each SFP-DD port when the ports are used as Ethernet ports.



Note

SFP28 direct attach cables (DAC) handling in SFP28 and SFP+ ports can be used for 25Gb and 10Gb by setting speed or Custom Auto-Negotiation Advertisement (CANA) at both ends.

Front panel ports 1/1-1/28, 1/29/1, 1/29/2, 1/30/1, and 1/30/2 can be used for traffic on the following switches:

- 5420F-8W-16P-4XE
- 5420F-24P-4XE
- 5420F-24S-4XE
- 5420F-24T-4XE
- 5420M-24W-4YE
- 5420M-24T-4YE

Front panel ports 1/1-1/52, 1/53/2, 1/53/1, 1/53/2, 1/54/1, and 1/54/2 can be used for traffic on the following switches:

- 5420F-16W-32P-4XE
- 5420F-48P-4XE
- 5420F-48P-4XL
- 5420F-48T-4XE
- 5420M-16MW-32P-4YE
- 5420M-48T-4YE
- 5420M-48W-4YE
- 5420M-24W-24S-4YE

On the following switches, you can create two 10Gb channels on each SFP-DD port:

- 5420M-24W-4YE
- 5420M-24T-4YE
- 5420M-16MW-32P-4YE
- 5420M-48T-4YE
- 5420M-48T-4YE
- 5420M-24W-24S-4YE

5420F-8W-16P-4XE

The following diagram illustrates the components of the front panel of the 5420F-8W-16P-4XE switch.




Figure 275: 5420F-8W-16P-4XE Front Panel

The front panel of the 5420F-8W-16P-4XE switch includes:

1. 8 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) 802.3bt Type 4 PoEMACsec-capable ports, 16 1000BASE-T FDX/HDX Type 2 PoE+ MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK, and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes two fixed fan modules, an AC power inlet connector, a power supply slot, and a grounding lug.

5420F-24P-4XE

The following diagram illustrates the components of the front panel of the 5420F-24P-4XE switch.




Figure 276: 5420F-24P-4XE Front Panel

The front panel of the 5420F-24P-4XE switch includes:

1. 24 10/100/1000BASE-T FDX/HDX 802.3bt Type 2 PoE+ MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes 2 fixed fan modules, an AC power inlet connector, a power supply slot, and a grounding lug.

5420F-24S-4XE

The following diagram illustrates the components of the front panel of the 5420F-24S-4XE switch.




Figure 277: 5420F-24S-4XE Front Panel

The front panel of the 5420F-24S-4XE switch includes:

1. 24 1000BASE-X SFP MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 10 Gbps SFP+ ports

The back panel (not shown) includes 2 fixed fan modules, an AC power inlet connector, a power supply slot, and a grounding lug.

5420F-24T-4XE

The following diagram illustrates the components of the front panel of the 5420F-24T-4XE switch.




Figure 278: 5420F-24T-4XE Front Panel

The front panel of the 5420F-24T-4XE switch includes:

1. 24 10/100/1000BASE-T FDX/HDX MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes 2 fixed fan modules, an AC power inlet connector, a power supply slot, and a grounding lug.

5420F-16MW-32P-4XE

The following diagram illustrates the components of the front panel of the 5420F-16MW-32P-4XE switch.




Figure 279: 5420F-16MW-32P-4XE Front Panel

The front panel of the 5420F-16MW-32P-4XE switch includes:

1. 16 100Mb/1GB/2.5GB 802.3bt Type 4 PoE ports.
2. 32 10/100/1000BASE-T FDX/HDX Type 2 PoE+ MACsec-capable ports.

3. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

4. 1 Micro-B USB console port
5. 1 RJ-45 out of band (OOB) management port
6. 1 RJ-45 serial console port
7. 2 Type-A USB ports
8. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
9. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes three fixed fan modules, a power supply slot, an AC power inlet connector, and a grounding lug.

5420F-16W-32P-4XE

The following diagram illustrates the components of the front panel of the 5420F-16W-32P-4XE switch.




Figure 280: 5420F-16W-32P-4XE Front Panel

The front panel of the 5420F-16W-32P-4XE switch includes:

1. 16 10/100/1000BASE-T FDX/HDX 802.3bt Type 4 PoE MACsec-capable ports.
2. 32 10/100/1000BASE-T FDX/HDX Type 2 PoE+ MACsec-capable ports.
3. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

4. 1 Micro-B USB console port
5. 1 RJ-45 out of band (OOB) management port
6. 1 RJ-45 serial console port
7. 2 Type-A USB ports
8. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
9. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes three fixed fan modules, a power supply slot, an AC power inlet connector, and a grounding lug.

5420F-48P-4XE

The following diagram illustrates the components of the front panel of the 5420F-48P-4XE switch.




Figure 281: 5420F-48P-4XE Front Panel

The front panel of the 5420F-48P-4XE switch includes:

1. 48 10/100/1000BASE-T FDX/HDX Type 2 PoE+ MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes three fixed fan modules, a power supply slot, an AC power inlet connector, and a grounding lug.

5420F-48P-4XL

The following diagram illustrates the components of the front panel of the 5420F-48P-4XL switch.




Figure 282: 5420F-48P-4XL Front Panel

The front panel of the 5420F-48P-4XL switch includes:

1. 48 10/100/1000BASE-T FDX/HDX Type 2 PoE+ MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10 Gbps LRM/MACsec-capable SFP+ ports

The back panel (not shown) includes three fixed fan modules, a power supply slot, an AC power inlet connector, and a grounding lug.

5420F-48T-4XE

The following diagram illustrates the components of the front panel of the 5420F-48T-4XE switch.




Figure 283: 5420F-48T-4XE Front Panel

The front panel of the 5420F-48T-4XE switch includes:

1. 48 10/100/1000BASE-T FDX/HDX MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10 Gbps SFP+ ports

The back panel (not shown) includes two fixed fan modules, a power supply slot, an AC power inlet connector, and a grounding lug.

5420M-24T-4YE

The following diagram illustrates the components of the front panel of the 5420M-24T-4YE switch.




Figure 284: 5420M-24T-4YE Front Panel

The front panel of the 5420M-24T-4YE switch includes:

1. 24 10/100/1000BASE-T FDX/HDX Type 2 MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10/25 Gbps SFP28 ports

The back panel (not shown) includes a removable fan module, two power supply slots, and a grounding lug.

5420M-24W-4YE

The following diagram illustrates the components of the front panel of the 5420M-24W-4YE switch.




Figure 285: 5420M-24W-4YE Front Panel

The front panel of the 5420M-24W-4YE switch includes:

1. 24 10/100/1000BASE-T FDX/HDX 802.3bt Type 4 PoE MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10/25 Gbps SFP28 ports

The back panel (not shown) includes a removable fan module, two power supply slots, and a grounding lug.

5420M-16MW-32P-4YE

The following diagram illustrates the components of the front panel of the 5420M-16MW-32P-4YE switch.




Figure 286: 5420M-16MW-32P-4YE Front Panel

The front panel of the 5420M-16MW-32P-4YE switch includes:

1. 16 100Mb/1Gb/2/5Gb 802.3bt Type 4 PoE ports.
2. 32 10/100/1000BASE-T FDX/HDX Type 2 PoE+ MACsec-capable ports.
3. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

4. 1 Micro-B USB console port
5. 1 RJ-45 out of band (OOB) management port
6. 1 RJ-45 serial console port
7. 2 Type-A USB ports
8. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
9. 4 1/10/25 Gbps SFP28 ports

The back panel (not shown) includes a removable fan module, two power supply slots, and a grounding lug.

5420M-48T-4YE

The following diagram illustrates the components of the front panel of the 5420M-48T-4YE switch.




Figure 287: 5420M-48T-4YE Front Panel

The front panel of the 5420M-48T-4YE switch includes:

1. 48 10/100/1000BASE-T FDX/HDX MACsec-capable ports.
2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

3. 1 Micro-B USB console port
4. 1 RJ-45 out of band (OOB) management port
5. 1 RJ-45 serial console port
6. 2 Type-A USB ports
7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
8. 4 1/10/25 Gbps SFP28 ports

The back panel (not shown) includes a removable fan module, two power supply slots, and a grounding lug.

5420M-48W-4YE

The following diagram illustrates the components of the front panel of the 5420M-48W-4YE switch.




Figure 288: 5420M-48W-4YE Front Panel

The front panel of the 5420M-48W-4YE switch includes:

- 1. 48 10/100/1000BASE-T FDX/HDX 802.3bt Type 4 PoE MACsec-capable ports.
- 2. Mode button for port LED control. Port LED mode indicators (SYS, SPD, STK and BT). Switch LEDs for system power (P1, P2), fan (FAN), and System Locator LED (LOC).



Note

The software only supports SYS and SPD LED mode indicators. STK and BT LED indicators remain off.

- 3. 1 Micro-B USB console port
- 4. 1 RJ-45 out of band (OOB) management port
- 5. 1 RJ-45 serial console port
- 6. 2 Type-A USB ports
- 7. 2 SFP-DD Universal Ethernet ports (labeled U1 and U2)
- 8. 4 1/10/25 Gbps SFP28 ports

The back panel (not shown) includes a removable fan module, two power supply slots, and a grounding lug.

5420M-24W-24S-4YE front panel

The following figure illustrates the components on the front panel of the 5420M-24W-24S-4YE.




Figure 289: 5420M-24W-24S-4YE Front Panel

The front panel of the switch includes:

Table 369: Figure call outs

Callout	5420M-24W-24S-4YE
1	10/100/1000BASE-T full/half duplex, MACsec capable ports
2	24x100M (FX/LX)/1G SFP ports
3	100MB/1Gb/2.5Gb ports with 802.3bt Type 4 PoE (90W)
4	Mode button
5	Type- C USB console port
6	System LEDs
7	RJ-45 Management and Console ports
8	Type A USB ports

Table 369: Figure call outs (continued)

Callout	5420M-24W-24S-4YE
9	SFP-DD Universal Ethernet ports
10	10G SFP+ or 10/25G SFP28 uplink ports

The back panel (not shown) includes a removable fan module, two power supply slots, and a grounding lug.

5520 Series

The following diagrams illustrate the components on the front panels of the 5520 Series switches.

For information about supported VIM modules, see *5520 Series Installation Guide*.


Additional Ports

In addition to the regular fixed ports, all 5520 Series switches have two QSFP28/QSFP+ Universal Ethernet ports and a Versatile Interface Module (VIM) slot. If you install a VIM, the switch reserves the Universal Ethernet ports for advanced features. If you do not install a VIM, the switch reserves two internal VIM ports for advanced features and you can use the Universal Ethernet ports as regular ports. If you do not require advanced features, you can use the Universal Ethernet ports as regular ports. For more information about using the advanced-feature-bandwidth-reservation boot flag to configure the functionality of these ports, see [Advanced Feature Bandwidth Reservation](#) on page 275.

When used as regular ports, the port speed is 40 Gbps as a single channel port. Although the maximum supported single channel port speed is 40 Gbps, the ports can be channelized to operate as four 10 or 25 Gbps channels. You can use 100 Gbps LR4 optics in Universal Ethernet ports to have a 100 Gbps connection with a 4x25 Gbps MLT. When the switch detects the optic, it automatically channelizes the port.

5520-24T

The following diagram illustrates the components of the front panel of the 5520-24T switch. Slot 1 is used for the 24 copper ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.

**Figure 290: 5520-24T**

1. 24 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX), MACsec-capable ports
2. 1 USB A port
3. 1 USB micro B management port

4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5520-24W

The following diagram illustrates the components of the front panel of the 5520-24W switch. Slot 1 is used for the 24 copper ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.




Figure 291: 5520-24W

1. 24 10/100/1000BASE-T FDX/HDX 802.3bt Type 4 PoE MACsec-capable ports
2. 1 USB A port
3. 1 USB micro B management port
4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5520-48T

The following diagram illustrates the components of the front panel of the 5520-48T switch. Slot 1 is used for the 48 copper ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.




Figure 292: 5520-48T

1. 48 10/100/1000BASE-T FDX/HDX MACsec-capable ports
2. 1 USB A port
3. 1 USB micro B management port

4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5520-48W

The following diagram illustrates the components of the front panel of the 5520-48W switch. Slot 1 is used for the 48 copper ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.




Figure 293: 5520-48W

1. 48 10/100/1000BASE-T FDX/HDX 802.3bt Type 4 PoE MACsec-capable ports
2. 1 USB A port
3. 1 USB micro B management port
4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5520-12MW-36W

The following diagram illustrates the components of the front panel of the 5520-12MW-36W switch. Slot 1 is used for the 48 copper ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.



Figure 294: 5520-12MW-36W

1. 36 10/100/1000BASE-T FDX/HDX 802.3bt Type 4 PoE MACsec-capable ports
2. 1 USB A port
3. 1 USB micro B management port
4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.
6. 12 100 Mbps/1 Gbps/2.5 Gbps/5 Gbps 802.3bt Type 4 PoE MACsec-capable ports

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5520-24X

The following diagram illustrates the components of the front panel of the 5520-24X switch. Slot 1 is used for the 24 SFP+ ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.




Figure 295: 5520-24X

1. 24 100/1000BASE-X/10GBASE-X SFP+ ports
2. 1 USB A port
3. 1 USB micro B management port
4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5520-48SE

The following diagram illustrates the components of the front panel of the 5520-48SE switch. Slot 1 is used for the 48 SFP ports and the two QSFP28/QSFP+ ports, and slot 2 is used for Versatile Interface Module (VIM) ports, if a VIM is installed in the VIM slot.




Figure 296: 5520-48SE

1. 48 100/1000BASE-X MACsec-capable SFP ports
2. 1 USB A port
3. 1 USB micro B management port
4. 2 Universal Ethernet QSFP28 ports
5. VIM slot (shown with VIM covered). Port numbering depends on the type of VIM installed in the slot.

The back panel (not shown) includes an RJ-45 out of band (OOB) management port, an RJ-45 console port, a USB A port, locator LED, two PSU slots, and hot-swappable fan units.

5720 Series

The following diagrams illustrate the components on the front panels of the 5720 Series switches.



Note

In addition to the regular fixed ports, all 5720 Series switches have two QSFP28/QSFP+ Universal Ethernet ports and a Versatile Interface Module (VIM) slot. If you install a VIM, the 48-port models reserve the Universal Ethernet ports for advanced features. If you do not install a VIM, the switch reserves two internal VIM ports for advanced features and you can use the Universal Ethernet ports as regular ports. If you do not require advanced features, you can use the Universal Ethernet ports as regular ports. For more information about using the advanced-feature-bandwidth-reservation boot flag to configure the functionality of these ports, see [Advanced Feature Bandwidth Reservation](#) on page 275.

When used as regular ports, the port speed is 100 Gbps as a single channel port. Although the maximum supported single channel port speed is 100 Gbps, the ports can be channelized to operate as four 10 or 25 Gbps channels.

5720-24MW

The following diagram illustrates the components of the front panel of the 5720-24MW switch.





Figure 297: View of the 5720-24MW front panel

1. USB Micro-B console port
2. USB Type-A port
3. System LEDs
4. Mode button
5. 2 Universal Ethernet/QSFP28 ports (unpopulated)
6. VIM Slot. Port numbering depends on the type of VIM installed in the slot.
7. 24 100M/1/2.5/5GbaseT 802.3bt PoE (90W) full-duplex (FDX), MACsec-capable ports

5720-48MW


The following diagram illustrates the components of the front panel of the 5720-48MW switch.



1. USB Micro-B console port
2. USB Type-A port
3. System LEDs
4. Mode button
5. 2 Universal Ethernet/QSFP28 ports (unpopulated)
6. VIM Slot. Port numbering depends on the type of VIM installed in the slot.
7. 48 100M/1/2.5/5GbaseT 802.3bt PoE (90W) full-duplex (FDX), MACsec-capable ports

5720-24MXW

The following diagram illustrates the components of the front panel of the 5720-24MXW switch.



1. USB Micro-B console port
2. USB Type-A port
3. System LEDs
4. Mode button
5. 24 100M/1/2.5/5/10GbaseT full-duplex (FDX), MACsec capable ports with 802.3bt Type 4 PoE (90W)
6. VIM Slot. Port numbering depends on the type of VIM installed in the slot.
7. 2 Universal Ethernet/QSFP28 ports (unpopulated)

5720-48MXW

The following diagram illustrates the components of the front panel of the 5720-48MXW switch.




Figure 298: 5720-48MXW Front Panel

1. USB Micro-B console port
2. USB Type-A port
3. System LEDs
4. Mode button
5. 2 Universal Ethernet/QSFP28 ports (unpopulated)
6. VIM Slot. Port numbering depends on the type of VIM installed in the slot.
7. 48 100M/1/2.5/5/10GbaseT full-duplex (FDX), MACsec capable ports with 802.3bt Type 4 PoE (90W)

VIMs

- 5720-VIM-6YE

Versatile Interface Module that provides six 25GbE SFP28 MACsec-capable ports


- 5720-VIM-2CE

Versatile Interface Module that provides two 100GbE QSFP28 MACsec-capable ports

For more information about VIM modules, see *5720 Series Installation Guide*.

5720 Series Back Panel View

The following diagram illustrates the components of the back panel of the 5720 Series



1. SSD slot
2. Fan modules
3. 10/100/1000 BASE-T out-of-band management port
4. Power supply slots
5. Grounding lug
6. Serial console port (RJ-45)
7. Type-A USB port for management or external USB flash

7520 Series

7520-48Y-8C

The following figure illustrates the components on the front panel of the 7520-48Y-8C.




Figure 299: 7520-48Y-8C front panel view

1. 48 1/10/25 Gbps SFP28 ports
2. 8 40/100 Gbps QSFP28 ports
3. 1 RJ-45 management port for CPU and BMC network access
4. 1 USB Type-A port
5. 1 RJ-45 serial console port
6. System LEDs

7520-48YE-8CE

The following figure illustrates the components on the front panel of the 7520-48YE-8CE.




Figure 300: 7520-48YE-8CE front panel view

1. 48 10/25 Gbps SFP28 MACsec-capable ports
2. 8 40/100 Gbps QSFP28 MACsec-capable ports
3. 1 RJ-45 management port for CPU and BMC network access
4. 1 USB Type-A port
5. 1 RJ-45 serial console port
6. System LEDs

7520-48XT-6C

The following figure illustrates the components on the front panel of the 7520-48XT-6C.




Figure 301: 7520-48XT-6C front panel view

1. 48 1G/10GBase-T ports
2. 6 40/100 Gbps QSFP28 ports
3. 1 RJ-45 management port for CPU and BMC network access
4. 1 USB Type-A port
5. 1 RJ-45 serial console port
6. System LEDs

The back panel (not shown) includes dual redundant hot-swappable power supplies and six hot-swappable fan units.

The 7520 Series supports one internal IAH port, 1/s1, used by IAH virtual machines. For conceptual information see [Extreme Integrated Application Hosting](#) on page 1097. For configuration instructions to use internal IAH ports, see [Virtual Services Configuration using CLI](#) on page 1112 and [Virtual Services Configuration using EDM](#) on page 1126.

7720 Series

7720-32C

The following figure illustrates the components on the front panel of the 7720-32C.




Table 370: Figure call outs

Callout	7720-32C
1	32 40/100 Gbps QSFP28 ports
2	1 RJ-45 serial console port
3	1 USB Type-A port
4	Two RJ-45 management ports: <ul style="list-style-type: none">• MGMT2 - CPU and BMC (secondary)• MGMT1 - Direct to CPU (primary)

The back panel (not shown) includes dual redundant hot-swappable power supplies and six hot-swappable fan units.


The 7720-32C supports two internal Extreme Integrated Application Hosting (IAH) ports labeled as Insight ports, 1/s1 and 1/s2, used by IAH virtual machines. For conceptual information see [Extreme Integrated Application Hosting](#) on page 1097. For configuration instructions to use internal IAH ports, see [Virtual Services Configuration using CLI](#) on page 1112 and [Virtual Services Configuration using EDM](#) on page 1126.

7830 Series

The following diagrams illustrate the components on the front panels of the 7830 Series switches.

7830-32CE-8DE

The following diagram illustrates the components of the front panel of the 7830-32CE-8DE switch.

**Table 371: Figure call outs**

Callout	7830-32CE-8DE
1	100M/1G/10G RJ-45 out-of-band management port
2	System LEDs
3	RJ-45 Serial console ports
4	Global Navigation Satellite System for GPS signal (not supported)
5	Mode button and channel selection LEDs
6	Advanced services port for the processor to offload to external appliance (not supported)
7	Time of Day (TOD), 1PPS signal inputs
8	HD-BNC connector for Time synchronization, 1PPS and 10MHz
9	USB Type A port
10	VIM slots (unpopulated)
11	QSPF56-DD 10G/25G/100G/400G MACsec capable ports
12	QSFP28 10G/25G/40G/100G MACsec capable ports
13	1G/10G SFP+ out-of-band management port

The back panel (not shown) includes a grounding lug, 4 fan modules, and 2 modular power supplies.

Interface Indexes

The Simple Network Management Protocol (SNMP) uses interface indexes to identify ports, Virtual Local Area Networks (VLAN), and Multilink Trunking (MLT).

Port Interface Index

To determine the interface index (IfIndex), you can calculate it, or use the CLI command provided in this section.

As a result of channelization support, the IfIndex of each channelization-capable port increases by 4. The number is reserved for the 3 sub-ports when channelization is enabled.

For switches that include channelization-capable ports, use the following equations to determine the IfIndex of a port:

- If the port does not support channelization, use $(64 \times \text{slot number}) + 128 + (\text{port number} - 1)$.
- If the port supports channelization, use the following equations:
 - for the port in question: $(64 \times \text{slot number}) + 128$
 - for subsequent ports: $(64 \times \text{slot number}) + 128 + ((\text{port number} - 1) * 4)$

This equation reserves space for the creation of the 3 sub-ports on the previous port, if or when you enable channelization.

Table 372: Number of slots per platform

Platform	Number of slots
4220 Series	1
5320 Series	1
5420 Series	1
5520 Series	2
5720 Series	2
7520 Series	1
7720 Series	1
7830 Series	3

CLI command

To determine the port interface index through the CLI, use the following command:

```
show interfaces gigabitEthernet
```

The following example shows output for this command:

```
Switch:1(config)#show interfaces gigabitEthernet
```

Port Interface								
PORT NUM	INDEX	DESCRIPTION	LINK TRAP	PORT LOCK	MTU	PHYSICAL ADDRESS	STATUS	
							ADMIN	OPERATE
1/1	192	1000BaseTX	true	false	1950	f0:64:26:70:f8:00	up	up
1/2	193	1000BaseTX	true	false	1950	f0:64:26:70:f8:01	up	down
1/3	194	1000BaseTX	true	false	1950	f0:64:26:70:f8:02	up	down
1/4	195	1000BaseTX	true	false	1950	f0:64:26:70:f8:03	up	down
1/5	196	1000BaseTX	true	false	1950	f0:64:26:70:f8:04	up	down
1/6	197	1000BaseTX	true	false	1950	f0:64:26:70:f8:05	up	down
1/7	198	1000BaseTX	true	false	1950	f0:64:26:70:f8:06	up	down
1/8	199	1000BaseTX	true	false	1950	f0:64:26:70:f8:07	up	down
1/9	200	1000BaseTX	true	false	1950	f0:64:26:70:f8:08	up	down

```

1/10    201    1000BaseTX      true  false    1950  f0:64:26:70:f8:09 up    down
1/11    202    1000BaseTX      true  false    1950  f0:64:26:70:f8:0a up    down
1/12    203    1000BaseTX      true  false    1950  f0:64:26:70:f8:0b up    down
1/13    204    1000BaseTX      true  false    1950  f0:64:26:70:f8:0c up    down
--More-- (q = quit)

```

VLAN interface index

The interface index of a VLAN is computed using the following formula:

Because the default VLAN always uses an MGID value of 1, its interface index is always 2049.

$\text{ifIndex} = 2048 + \text{VLAN multicast group ID (MGID)}$

MLT interface index

The interface index of a multilink trunk (MLT) is computed using the following formula:

$\text{ifIndex} = 6143 + \text{MLT ID number}$

MAC Address Assignment

You must understand MAC addresses assignment if you perform one of the following actions:

- Define static Address Resolution Protocol (ARP) entries for IP addresses in the switch
- Use a network analyzer to decode network traffic

Each chassis is assigned a base number of MAC addresses with a number reserved for ports and other internal purposes, and the remainder assigned to routable VLANs. The following table identifies the numbers provided by the switch.

Base assignment	Reserved	Assigned to routable VLANs
1,024	First 256	remaining 768

Virtual MAC Addresses

Virtual MAC addresses are the addresses assigned to VLANs. The system assigns a virtual MAC address to a VLAN when it creates the VLAN. The MAC address for a VLAN IP address is the virtual MAC address assigned to the VLAN.



Supported Standards, RFCs, and MIBs

[Supported IEEE Standards](#) on page 4245

[Supported RFCs](#) on page 4246

[Standard MIBs](#) on page 4253

[Proprietary MIBs](#) on page 4256

This section provides information about supported IEEE standards, RFCs, and MIBs.

Supported IEEE Standards

Table 373: Supported IEEE Standards

IEEE standard	Description
802.1AB	LLDP
802.1ag	Connectivity Fault Management
802.1ah	Provider Backbone Bridging
802.1aq	Shortest Path Bridging (SPB)
802.1AX	Link Aggregation
802.1D	MAC Bridges
P802.1p	Traffic Class Expediting and Dynamic Multicast Filtering
802.1Q	Virtual LANs
802.1s	Multiple Spanning Trees
802.1t	802.1D Technical & Editorial Corrections
802.1w	Rapid Spanning Tree Protocol (RSTP)
802.1X-2010	Port-based Network Access Control (NAC)
802.3 CSMA/CD Ethernet ISO/IEC 8802	International Organization for Standardization (ISO) /International Eletrotechnical Commission (IEC) 8802-3
802.3ab	1000 Mbps Operation, implemented as 1000BASE-T Copper
802.3ac	Carrier sense multiple access with collision detection (CSMA/CD) frame extensions for Virtual Bridged Local Area Networks (VLAN) tagging on 802.3 networks

Table 373: Supported IEEE Standards (continued)

IEEE standard	Description
802.1AE	MAC Security
802.3ae	10 Gbps Operation, implemented as 10GBASE-X SFP+
802.3af 802.3at 802.3bt (Type 3 and Type 4)	Power over Ethernet (PoE)
802.3az	Energy Efficient Ethernet (EEE)
802.3ba	40 Gbps and 100 Gbps Operation, implemented as 40GBASE-QSFP+ and 100GBASE-QSFP28
802.3x	Full Duplex & Flow Control
802.3z	1000 Mbps Operation, implemented as 1000BASE-X SFP
ANSI/TIA-1057	LLDP-MED

Supported RFCs

The following table and sections list the RFCs that the switch supports.

Table 374: Supported request for comments

Request for comment	Description
draft-grant-tacacs-02.txt	TACACS+ Protocol
RFC 768	UDP Protocol
RFC 783	Trivial File Transfer Protocol (TFTP)
RFC 791	Internet Protocol (IP)
RFC 792	Internet Control Message Protocol (ICMP)
RFC 793	Transmission Control Protocol (TCP)
RFC 815	IP Datagram Reassembly Algorithm
RFC 826	Address Resolution Protocol (ARP)
RFC 854	Telnet protocol
RFC 894	A standard for the Transmission of IP Datagrams over Ethernet Networks
RFC 896	Congestion control in IP/TCP internetworks
RFC 906	Bootstrap loading using TFTP
RFC 950	Internet Standard Subnetting Procedure
RFC 951	BootP
RFC 959, RFC 1350, and RFC 2428	FTP and TFTP client and server

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 1027	Using ARP to implement transparent subnet gateways/Nortel Subnet based VLAN
RFC 1042	Encapsulating the Internet Protocol (IP) datagrams and Address Resolution Protocol (ARP) Note: Applicable for 802.3 networks only
RFC 1058	RIPv1 Protocol
RFC 1112	Host Extensions for IP Multicasting (IGMPv1)
RFC 1122	Requirements for Internet Hosts
RFC 1155	SMI
RFC 1156	MIB for network management of TCP/IP
RFC 1157	SNMP
RFC 1195	Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
RFC 1212	Concise MIB definitions
RFC 1213	TCP/IP Management Information Base
RFC 1215	Convention for defining traps for use with the SNMP
RFC 1253	OSPF MIB
RFC 1256	ICMP Router Discovery
RFC 1305	Network Time Protocol v3 Specification, Implementation and Analysis ³
RFC 1323	TCP Timestamp (The switch is only compliant if the TCP timestamp is enabled.)
RFC 1340	Assigned Numbers
RFC 1350	TFTP Protocol (Revision 2)
RFC 1389	Ethernet MIB
RFC 1442	Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)
RFC 1450	Management Information Base for v2 of the Simple Network Management Protocol (SNMPv2)
RFC 1519	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 1541	Dynamic Host Configuration Protocol
RFC 1542	Clarifications and Extensions for the Bootstrap Protocol
RFC 1573	Interface MIB
RFC 1587	The OSPF NSSA Option
RFC 1591	DNS Client
RFC 1650	Definitions of Managed Objects for the Ethernet-like Interface Types
RFC 1657	BGP-4 MIB using SMIv2
RFC 1723	RIP v2 — Carrying Additional Information
RFC 1812	Router requirements
RFC 1866	Hypertext Markup Language version 2 (HTMLv2) protocol
RFC 1981	Path MTU discovery
RFC 1907	Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)
RFC 1930	Guidelines for creation, selection, and registration of an Autonomous System (AS)
RFC 2021	Remote Network Monitoring (RMON) MIB using SMIv2
RFC 2068	Hypertext Transfer Protocol
RFC 2080	RIP
RFC 2131	Dynamic Host Control Protocol (DHCP)
RFC 2132	DHCP Options and BOOTP Vendor Extensions
RFC 2138	RADIUS Authentication
RFC 2139	RADIUS Accounting
RFC 2178	OSPF MD5 cryptographic authentication / OSPFv2
RFC 2233	The Interfaces Group MIB using SMIv2
RFC 2236	IGMPv2 Snooping
RFC 2358	Definitions of Managed Objects for the Ethernet-like Interface Types
RFC 2284	PPP Extensible Authentication Protocol
RFC 2328	OSPFv2

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 2338	VRRP: Virtual Redundancy Router Protocol
RFC 2362	PIM-SM
RFC 2407	IP Security Domain Interpretation of Internet Security Association and Key Management Protocol (ISAKMP)
RFC 2408	Internet Security Associations and Key Management Protocol (ISAKMP)
RFC 2428	FTP Extensions for IPv6
RFC 2452	IPv6 Management Information Base for the Transmission Control Protocol
RFC 2454	IPv6 Management Information Base for the User Datagram Protocol
RFC 2466	IPv6 Management Information Base for ICMPv6 Group
RFC 2541	DNS Security Operational Considerations
RFC 2453	RIPv2 Protocol
RFC 2460	IPv6 base stack
RFC 2462	IPv6 Stateless Address Autoconfiguration
RFC 2463	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
RFC 2464	Transmission of IPv6 packets over Ethernet networks
RFC 2474	Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
RFC 2475	An Architecture for Differentiated Services
RFC 2545	Use of BGP-4 multi-protocol extensions for IPv6 inter-domain routing
RFC 2548	Microsoft vendor specific RADIUS attributes
RFC 2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC 2573	SNMP Applications
RFC 2574	User-based Security Model (USM) for v3 of the Simple Network Management Protocol (SNMPv3)

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
RFC 2576	Coexistence between v1, v2, & v3 of the Internet standard Network Management Framework
RFC 2578	Structure of Management Information v2 (SMIv2)
RFC 2579	Textual Conventions for SMI v2
RFC 2580	Conformance Statements for SMI v2
RFC 2598	Expedited Forwarding PHB
RFC 2616	Hypertext Transfer Protocol 1.1
RFC 2710	Multicast Listener Discovery (MLD) for IPv6
RFC 2716	PPP EAP Transport Level Security (TLS) Authentication Protocol
RFC 2737	Entity MIB (Version 2)
RFC 2787	Definitions of Managed Objects for the Virtual Router Redundancy Protocol
RFC 2819	Remote Network Monitoring (RMON) MIB
RFC 2863	Interface Group MIB
RFC 2865	RADIUS
RFC 2874	DNS Extensions for IPv6
RFC 2918	Route Refresh Capability for BGP-4
RFC 2925	Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations
RFC 2992	Analysis of an Equal-Cost Multi-Path Algorithm
RFC 3046	DHCP Option 82
RFC 3162	IPv6 RADIUS client
RFC 3246	An Expedited Forwarding PHB (Per-Hop Behavior)
RFC 3315	IPv6 DHCP Relay
RFC 3376	IGMPv3
RFC 3411 and RFC 2418	SNMP over IPv6 networks
RFC 3416	v2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)
RFC 3417	Transport Mappings for SNMP

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 3484	Default Address Selection for IPv6
RFC 3513	Internet Protocol Version 6 (IPv6) Addressing Architecture
RFC 3569	An overview of Source-Specific Multicast (SSM)
RFC 3579	RADIUS Support For Extensible Authentication Protocol (EAP)
RFC 3580	IEEE 802.1X Remote Authentication Dial In User Service
RFC 3587	IPv6 Global Unicast Address Format
RFC 3596	DNS Extensions for IPv6
RFC 3621	Power Ethernet MIB
RFC 3748	Extensible Authentication Protocol
RFC 3810	Multicast Listener Discovery Version 2 (MLDv2) for IPv6
RFC 3825	Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information
RFC 3879	Deprecating Site Local Addresses
RFC 3986	Uniform Resource Identifiers (URI)
RFC 4007	IPv6 Scoped Address Architecture
RFC 4022	Management Information Base for the Transmission Control Protocol (TCP)
RFC 4113	Management Information Base for the User Datagram Protocol (UDP)
RFC 4133	Entity MIB (version 3)
RFC 4193	Unique Local IPv6 Unicast Address
RFC 4213	IPv6 configured tunnel
RFC 4250–RFC 4256	SSH server and client support
RFC 4291	IPv6 Addressing Architecture
RFC 4292	IP Forwarding Table MIB
RFC 4293	MIB for IP
RFC 4301	Security Architecture for IPv6
RFC 4302	IP Authentication Header (AH)
RFC 4303	IP Encapsulated Security Payload (ESP)
RFC 4305	Cryptographic algorithm implementation requirements for ESP and AH

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 4308	Cryptographic suites for Internet Protocol Security (IPsec)
RFC 4363	Bridges with Traffic MIB
RFC 4443	ICMP for IPv6
RFC 4541	Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping
RFC 4552	OSPFv3 Authentication and confidentiality for OSPFv3
RFC 4601	Protocol Independent Multicast - Sparse Mode (PIM-SM)
RFC 4604	Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast
RFC 4607	Source-Specific Multicast (SSM)
RFC 4649	Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option
RFC 4673	RADIUS Dynamic Authorization Server MIB
RFC 4675	Egress VLAN
RFC 4750	OSPFv2 MIB
RFC 4760	Multiprotocol Extensions for BGP-4
RFC 4835	Cryptographic algorithm implementation for ESP and AH
RFC 4861	IPv6 Neighbor discovery
RFC 4862	IPv6 stateless address autoconfiguration
RFC 4893	BGP Support for Four-octet AS Number Space
RFC 5095	Deprecation of Type 0 Routing headers in IPv6
RFC 5176	Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)
RFC 5187	OSPFv3 Graceful Restart (helper-mode only)
RFC 5242	The Syslog Protocol
RFC 5321	Simple Mail Transfer Protocol
RFC 5340	OSPF for IPv6

Table 374: Supported request for comments (continued)

Request for comment	Description
RFC 5746	Secure SSL Renegotiation
RFC 5798	Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6
RFC 5905	Network Time Protocol Version 4: Protocol and Algorithms Specification
RFC 5997	Use of Status-Server Packets in the Remote Authentication Dial In User Service (RADIUS) Protocol
RFC 6105	IPv6 Router Advertisement Guard
RFC 6329	IS-IS Extensions supporting Shortest Path Bridging
RFC 7047	The Open vSwitch Database Management Protocol
RFC 7348	Virtual Extensible LAN (VXLAN)
RFC 7610	DHCPv6 Shield
RFC 7627	Transport Layer Security (TLS) Session Hash and Extended Master Secret Extension

Standard MIBs

The following table details the standard MIBs that the switch supports.

Table 375: Supported MIBs

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STD MIB2— Link Aggregation Control Protocol (LACP) (802.3ad)	802.3ad	ieee802-lag.mib
STD MIB3—Extensible Authentication Protocol Over Local Area Networks (EAPoL) (802.1x)	802.1x	ieee8021x.mib
STD MIB4—Internet Assigned Numbers Authority (IANA) Interface Type		iana_if_type.mib
STD MIB5—Structure of Management Information (SMI)	RFC1155	rfc1155.mib

Table 375: Supported MIBs (continued)

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STD MIB6—Simple Network Management Protocol (SNMP)	RFC1157	rfc1157.mib
STD MIB7—MIB for network management of Transfer Control Protocol/Internet Protocol (TCP/IP) based Internet MIB2	RFC1213	rfc1213.mib
STD MIB8—A convention for defining traps for use with SNMP	RFC1215	rfc1215.mib
STD MIB10—Definitions of Managed Objects for Bridges	RFC1493	rfc1493.mib
STD MIB11—Evolution of the Interface Groups for MIB2	RFC2863	rfc2863.mib
STD MIB12—Definitions of Managed Objects for the Ethernet-like Interface Types	RFC1643	rfc1643.mib
STD MIB15—Remote Network Monitoring (RMON)	RFC2819	rfc2819.mib
STD MIB17—Management Information Base of the Simple Network Management Protocol version 2 (SNMPv2)	RFC1907	rfc1907.mib
STD MIB21—Interfaces Group MIB using SMIv2	RFC2233	rfc2233.mib
STD MIB26b—Message Processing and Dispatching for the SNMP	RFC2572	rfc2572.mib
STD MIB26c—SNMP Applications	RFC2573	rfc2573.mib
STD MIB26d—User-based Security Model (USM) for version 3 of the SNMP	RFC2574	rfc2574.mib
STD MIB26e—View-based Access Control Model (VACM) for the SNMP	RFC2575	rfc2575.mib

Table 375: Supported MIBs (continued)

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB26f—Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework	RFC2576	rfc2576.mib
STDMIB29—Definitions of Managed Objects for the Virtual Router Redundancy Protocol	RFC2787	rfc2787.mib
STDMIB31—Textual Conventions for Internet Network Addresses	RFC2851	rfc2851.mib
STDMIB32—The Interface Group MIB	RFC2863	rfc2863.mib
STDMIB33—Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations	RFC2925	rfc2925.mib
STDMIB35—Internet Group Management Protocol MIB	RFC2933	rfc2933.mib
STDMIB36—Protocol Independent Multicast MIB for IPv4	RFC2934	rfc2934.mib
STDMIB38—SNMPv3 These Request For Comments (RFC) make some previously named RFCs obsolete	RFC3411, RFC3412, RFC3413, RFC3414, RFC3415	rfc2572.mib, rfc2573.mib, rfc2574.mib, rfc2575.mib
STDMIB39—Entity Sensor Management Information Base	RFC3433	
STDMIB40—The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model	RFC3826	rfc3826.mib
STDMIB41—Management Information Base for the Transmission Control protocol (TCP)	RFC4022	rfc4022.mib
STDMIB43—Management Information Base for the User Datagram Protocol (UDP)	RFC4113	rfc4113.mib

Table 375: Supported MIBs (continued)

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
Q-BRIDGE-MIB — Management Information Base for managing Virtual Bridged LANs	RFC4363	q-bridge.mib
LLDP-EXT-MED-MIB — LLDP-MED	ANSI/TIA-1057	lldpExtMed.mib

Proprietary MIBs

The following table details the proprietary MIBs that the switch supports.

Table 376: Proprietary MIBs

Proprietary MIB name	File name
Extreme Networks Energy Saver MIB	bayStackNes.mib
Extreme Networks Link-state tracking (LST) MIB	bayStackLinkStateTracking.mib
Extreme Networks IGMP MIB	rfc_igmp.mib
Extreme Networks IP Multicast MIB	ipmroute_rcc.mib
Extreme Networks MIB definitions	wf_com.mib
Extreme Networks PIM MIB	pim-rcc.mib
Extreme Networks RSTP/MSTP proprietary MIBs	nnrst000.mib, nnmst000.mib
Other SynOptics definitions	s5114roo.mib
Other SynOptics definitions	s5emt103.mib
Other SynOptics definitions	s5tcs112.mib
Other SynOptics definition for Combo Ports	s5ifx.mib
Other SynOptics definition for PoE	bayStackPethExt.mib
Rapid City MIB Note: The MACsec tables, namely, rcMACSecCATable and rcMACSecIfConfigTable are a part of the Rapid City MIB.	rapid_city.mib
SynOptics Root MIB	synro.mib



ICMPv6 Type and Code

The Internet Control Message Protocol (ICMPv6) uses many messages identified by a type and code field (see RFC 4443). Error messages use message types 0 to 127. Informational messages use message types 128 to 255. The following table provides the type and code reference.

Table 377: ICMPv6 type and code details

Type	Name	Code	Reference
1	Destination Unreachable	0—no route to destination 1—communication with destination administratively prohibited 2—(not assigned) 3—address unreachable 4—port unreachable	RFC 4443
2	Packet Too Big		RFC 4443
3	Time Exceeded	0—hop limit exceeded in transit 1—fragment reassembly time exceeded	RFC 4443
4	Parameter Problem	0—erroneous header field encountered 1—unrecognized Next Header type encountered 2—unrecognized IPv6 option encountered	RFC 4443
128	Echo Request		RFC 4443
129	Echo Reply		RFC 4443
130	Multicast Listener Query		
131	Multicast Listener Report		

Table 377: ICMPv6 type and code details (continued)

Type	Name	Code	Reference
132	Multicast Listener Done		
133	Router Solicitation		RFC 4861
134	Router Advertisement		RFC 4861
135	Neighbor Solicitation		RFC 4861
136	Neighbor Advertisement		RFC 4861
137	Redirect Message		RFC 4861
138	Router Renumbering	0—router renumbering command 1—router renumbering result 255—sequence number reset	
139	ICMP Node Information Query		
140	ICMP Node Information Response		
141	Inverse neighbor discovery Solicitation Message		RFC 3122
142	Inverse neighbor discovery Advertisement Message		RFC 3122
143	Version 2 Multicast Listener Report		RFC 3810
144	Home Agent Address Discovery Request Message		RFC 3775
145	Home Agent Address Discovery Reply Message		RFC 3775
146	Mobile Prefix Solicitation		RFC 3775
147	Mobile Prefix Advertisement		RFC 3775