

Configuring IPv4 Routing for VOSS

© 2017-2020, Extreme Networks, Inc. All Rights Reserved.

Legal Notice

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

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

Trademarks

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

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

For additional information on Extreme Networks trademarks, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: https://www.extremenetworks.com/support/policies/open-source-declaration/

Contents

Chapter 1: About this Document	9
Purpose	
Conventions	10
Text Conventions	10
Documentation and Training	12
Getting Help	12
Providing Feedback	13
Chapter 2: New in this Document	15
Notice about Feature Support	16
Chapter 3: IP Routing Operations	17
IP routing operations fundamentals	
IP addressing	
Loopback	
Static routes	21
Black hole static routes	23
VLANs and routing	23
Equal Cost Multipath	
Alternative routes	25
IP Source Routing	28
Multihoming	28
Enabling or Disabling IPv4 ICMP Broadcast	29
IP routing configuration using the CLI	30
Enabling routing globally or on a VRF instance	
Enabling routing on an IP interface	31
Deleting a dynamically learned route	32
Configuring IP route preferences	33
Flushing routing tables by VLAN or port	35
Assigning an IP address to a port	35
Assign an IP Address to a VLAN	37
Configure an IP Address for a Segmented Management Instance	38
Configure an IP Address for the Management Port	
Viewing IP addresses for all router interfaces	40
Configure IP Routing Globally or for a VRF	41
Configure ECMP	43
Configuring Static Routes	45
Configuring a black hole static route	48
Configuring a default static route	
Enabling ICMP Router Discovery globally	
Enabling or disabling IPv4 ICMP broadcast globally	52

Enabling or disabling IPv4 ICMP broadcast per VRF	
Configuring Router Discovery on a port or VLAN	
Configuring a CLIP Interface	
Configure BFD on an IPv4 Interface	57
Display BFD Configurations on the Loopback Interface	59
Viewing TCP and UDP information	59
IP routing configuration using Enterprise Device Manager	61
Enabling routing for a router or a VRF instance	62
Deleting a dynamically-learned route	62
Configuring IP route preferences	64
Flushing routing tables by VLAN	65
Flushing routing tables by port	65
Assigning an IP address to a port	66
Assigning an IP address to a VLAN	67
Viewing IP addresses for all router interfaces	68
Configuring IP routing features globally	
Configure ECMP	
Enabling alternative routes globally	72
Configure Static Routes using EDM	
Deleting a static route	
Configuring a default static route	
Configuring a black hole static route	76
Viewing IP routes	77
Configuring ICMP Router Discovery globally	79
Configuring the ICMP Router Discovery table	
Configuring ICMP Router Discovery for a port	
Configuring ICMP Router Discovery on a VLAN	
Configure a Circuitless IPv4 Interface	
Enabling OSPF on a CLIP interface	
Enabling PIM on a CLIP interface	85
Enable BFD on a CLIP interface	85
Viewing TCP global information	87
Viewing TCP connections information	
Viewing TCP listeners information	
Chapter 4: Distributed Virtual Routing	91
Distributed Virtual Routing (DvR) Fundamentals	
DvR Domain	
DvR Controller	
DvR Leaf Node	
Configuration Limitations	
Summary of Controller and Leaf Node Functions	
DvR backbone	
DvR operation	

ARP Learning	97
dvr-leaf-mode boot flag	
In-band management	98
DvR deployment scenarios	98
DvR Route Redistribution	101
DvR Limitations	105
Migrate from VRRP to DvR	107
DvR configuration using the CLI	109
Configuration Limitations on a DvR Controller	109
Configuring a DvR Controller	109
Disabling injection of default routes on a Controller	
Configuring DvR route redistribution	112
Clearing DvR host entries	
Configuring a DvR Leaf	
Configuring vIST on a DvR Leaf node pair	117
Configure a Management VLAN on a DvR Leaf Node	
Delete a Management VLAN on a DvR Leaf Node	120
Moving a vIST Leaf node pair from one domain to another	
Moving a vIST Controller pair from one domain to another	
Configuring a non-DvR BEB to join the DvR backbone	
DvR show commands	
Configuring a DvR solution	
DvR Configuration Using the EDM	
Configure a DvR Controller or a DvR Leaf Globally	
View DvR Routes	
View Members of a DvR Domain	
View DvR Backbone Members	
View DvR Interfaces	
View DvR Host Entries	
Clear DvR Host Entries	
View Layer 3 VSN Information	
View the DvR Database	
View DvR Backbone Entries on a Controller	165
Chapter 5: Address Resolution Protocol	167
Address Resolution Protocol	167
Reverse Address Resolution Protocol	169
ARP configuration using the CLI	
Enabling ARP on a port or a VLAN	
Enabling ARP proxy	
View ARP Information	
Configuring IP ARP static entries	
Clearing ARP entries	
Showing ARP table information	175

Configuring Gratuitous ARP	178
ARP configuration using Enterprise Device Manager	179
Enabling or disabling ARP on the brouter port or a VRF instance	179
Enabling or disabling ARP on a VLAN or a VRF instance	
Viewing and managing ARP	180
Creating static ARP entries	181
Configuring ARP proxy	182
Chapter 6: Dynamic Host Configuration Protocol and User Datagram Protocol	
Configuration	183
DHCP option 82	183
DHCP Suboptions	185
Agent Operations	186
UDP broadcast forwarding	187
DHCP and UDP configuration using the CLI	187
Configure DHCP Parameters Globally	187
Showing DHCP relay information	189
Configuring DHCP option 82	191
Configuring DHCP relay on a port or VLAN	192
Configuring UDP broadcast forwarding	195
Configuring UDP protocols	
Configuring a UDP port forward entry	
Configuring the UDP port forwarding list	197
Showing UDP forward information	
DHCP and UDP configuration using Enterprise Device Manager	
Configuring DHCP on a brouter port or a VRF instance	
Configuring BootP/DHCP on a VLAN or VRF instance	
Configuring DHCP relay	
Viewing DHCP relay configuration information	
Managing UDP forwarding protocols	
Managing UDP forwarding	
Creating the forwarding profile	
Managing the broadcast interface	
Viewing UDP endpoint information	
Chapter 7: Route Filtering and IP Policies	212
Route Filtering and IP Policies	212
Accept Policies	214
Redistribution Filters	214
Announce Policies	
Route Filtering Stages	
Prefix list	
Route Policy Definition	217
IP policy configuration using the CLI	
Configuring prefix lists	222

Configuring IP Route Policies	224
Configuring a policy to accept external routes from a router	231
Applying OSPF accept policy changes	232
Configuring inter-VRF redistribution policies	234
IP Policy Configuration using Enterprise Device Manager	236
Configuring a prefix list	237
Configure a Route Policy	238
Apply a Route Policy	242
Viewing IP routes	243
Configure an OSPF Accept Policy	245
Configuring inbound/outbound filtering policies on a RIP interface	246
Deleting inbound/outbound filtering policies on a RIP interface	247
Chapter 8: Routed Split MultiLink Trunking	248
RSMLT	
SMLT and RSMLT Operation in Layer 3 Environments	
RSMLT configuration using the CLI	
Configuring RSMLT on a VLAN	
Showing IP RSMLT information	
Configuring RSMLT edge support	
RSMLT configuration using Enterprise Device Manager	
Configuring RSMLT on a VLAN	
Viewing and editing RSMLT local information	
Viewing RSMLT peer information	
Enabling RSMLT Edge support	
Viewing RSMLT edge support information	
Chapter 9: Virtual Router Redundancy Protocol	
VRRP Fundamentals	
Critical IP Address	
VRRP and SMLT	
VRRP Fast Hello Timers	
Handling of IPv4 Layer 2 Unicast Packets at VRRP Backup Master	
VRRP guidelines	
VRRPv3	
VRRPv3 guidelines	
VRRP configuration using the CLI	
Configuring VRRP on a port or a VLAN	
Showing VRRP information	
Showing extended VLAN VRRP	
Showing VRRP interface information	
Enabling ping to a virtual IP address	
Configuring VRRP notification control	
Configuring VRRP version on an interface	
Enabling IPv4 VRRP preempt-mode	

Contents

VRRP configuration using EDM	285
Enabling VRRP global variables	286
Modifying VRRP parameters for an interface	286
Configuring VRRP on a V3 interface	289
Configuring VRRPv3 Checksum	291
Configuring Fast Advertisement Interval on a port or a VRF instance	292
Configuring Fast Advertisement Interval on a VLAN or a VRF instance	292
Chapter 10: VRF Lite	294
VRF Lite Fundamentals	
VRF Lite capability and functionality	297
VRF Lite and inter-VRF route redistribution	
Port parameters and VRF Lite management	300
Management VRF	
VRF Lite configuration rules	303
Virtualized Protocols	304
VRF Lite configuration using the CLI	305
Creating a VRF Instance	306
Associating a VLAN or port with a VRF instance	309
Creating an IP VPN instance on a VRF	311
Configure the Maximum Number of VRFs	312
VRF Lite configuration using Enterprise Device Manager	314
Configuring a VRF instance	
Associating a port to a VRF instance	316
Associating an Extreme Integrated Application Hosting Port to a VRF Instance	316
Configuring interVRF route redistribution policies	
Viewing brouter port and VRF associations	
Viewing global VRF status information	
Viewing VRF instance statistics and status information	
Viewing Statistics for a VRF	
Selecting and launching a VRF context view	
Create an IP VPN Instance on a VRF	
Configure the maximum number of VRFs	323
Glossary	325

8

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document provides information on features in VSP Operating System Software (VOSS). VOSS runs on the following product families:

- Extreme Networks VSP 4000 Series (includes VSP 4450 Series)
- Extreme Networks VSP 4900 Series
- Extreme Networks VSP 7200 Series
- Extreme Networks VSP 7400 Series
- Extreme Networks VSP 8000 Series (includes VSP 8200 Series and VSP 8400 Series)
- Extreme Networks VSP 8600 Series
- Extreme Networks XA1400 Series



VOSS is licensed on the XA1400 Series as a Fabric Connect VPN (FCVPN) application, which includes a subset of VOSS features. FCVPN transparently extends Fabric Connect services over third-party provider networks.

This document provides procedures and conceptual information that you can use to configure the general routing operations on the switch. The operations included are:

- Address Resolution Protocol (ARP)
- TCP and UDP
- Dynamic Host Configuration Protocol (DHCP) Relay
- Virtual Router Redundancy Protocol (VRRP)
- VRF-Lite
- Routed Split Multi-Link Trunking (RSMLT)
- · Circuitless IP (CLIP) interfaces

- Static routes
- Point-to-Point Protocol over Ethernet
- Equal Cost Multipath (ECMP)
- · Routing policies

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

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

Table 1: Notice Icons

Icon	Alerts you to
• Important:	A situation that can cause serious inconvenience.
Note:	Important features or instructions.
😷 Tip:	Helpful tips and notices for using the product.
▲ Danger:	Situations that will result in severe bodily injury; up to and including death.
⚠ Warning:	Risk of severe personal injury or critical loss of data.
⚠ Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description	
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.	
	If the command syntax is cfm maintenance-	
	domain maintenance-level <0-7> , you can	

Convention	Description
	enter cfm maintenance-domain
	maintenance-level 4.
Bold text	Bold text indicates the GUI object name you must act upon.
	Examples:
	• Click OK .
	On the Tools menu, choose Options .
Braces ({})	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.
	For example, if the command syntax is ip address {A.B.C.D}, you must enter the IP address in dotted, decimal notation.
Brackets ([])	Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command.
	For example, if the command syntax is show clock [detail], you can enter either show clock or show clock detail.
Ellipses ()	An ellipsis () indicates that you repeat the last element of the command as needed.
	For example, if the command syntax is ethernet/2/1 [<parameter></parameter>
	<pre><value>], you enter ethernet/2/1 and as many parameter-value pairs as you need.</value></pre>
Italic Text	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
Plain Courier Text	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages.
	Examples:
	• show ip route
	• Error: Invalid command syntax [Failed] [2013-03-22 13:37:03.303 -04:00]
Separator (>)	A greater than sign (>) shows separation in menu paths.

Convention	Description	
	For example, in the Navigation tree, expand the Configuration > Edit folders.	
Vertical Line ()	A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command.	
	For example, if the command syntax is access- policy by-mac action { allow deny }, you enter either access-policy by-mac action allow Or access-policy by-mac action deny, but not both.	

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation

Release Notes

Hardware and software compatibility for Extreme Networks products

Extreme Optics Compatibility

Other resources such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

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

Extreme Portal	Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.
The Hub	A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
Call GTAC	For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

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

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

Subscribe to Service Notifications

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

- 1. Go to www.extremenetworks.com/support/service-notification-form.
- 2. Complete the form (all fields are required).
- 3. Select the products for which you would like to receive notifications.
 - Note:

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

4. Select Submit.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

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

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Chapter 2: New in this Document

The following sections detail what is new in this document.

Bidirectional Forwarding Detection over Fabric Extend Tunnels

Bidirectional Forwarding Detection (BFD) provides a fast failure-detection mechanism between peer systems. The peer systems exchange BFD packets, and when one of the systems does not receive a BFD packet after a specific period of time, the system assumes that the link or the other system is not operating, and declares the link down.

BFD functionality on VOSS is extended to support fast failure-detection for Fabric Extend (FE) tunnels.

For information about enabling BFD for FE tunnels for CLIP interfaces, see the following sections:

- Configure BFD on an IPv4 Interface on page 57
- Display BFD Configurations on the Loopback Interface on page 59
- Enable BFD on a CLIP interface on page 85

For more information about BFD, see Administering VOSS.

Rlogin Deprecation

Remote login (rlogin) is no longer supported on the switch.

rlogin parameters and outputs remain in this document, but are only supported on VSP 8600 Series .

Segmented Management Instance IPv4 Routing Changes

The Management Instance now supports CLIP, VLAN, and a new Out-of-Band (OOB) management interface type. The legacy mgmtEthernet management interface automatically migrates to a Segmented Management Instance OOB management interface during the upgrade to this release. You can use the new mgmt OOB interface CLI mode to configure Layer 3 networking OOB management parameters. You can use the legacy interface mgmtEthernet mgmt CLI mode to configure Layer 1 and Layer 2 networking OOB management parameters.

For information, see the following sections:

- Configure an IP Address for a Segmented Management Instance on page 38
- Configure an IP Address for the Management Port on page 39
- Management VRF on page 300

Note:

All the IPv4 routing mgmtEthernet mgmt and mgmtRouter VRF procedures remain in this document to support the VSP 8600 Series.

Notice about Feature Support

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

If a documented command, parameter, tab, or field does not display on your hardware, it is not supported.

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

Chapter 3: IP Routing Operations

This section provides conceptual information and procedures to configure IP Routing using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

IP routing operations fundamentals

Use the information in this section to understand IP routing.

For more information about Border Gateway Protocol (BGP), see Configuring BGP Services for VOSS.

For more information about Open Shortest Path First (OSPF) and Routing Information Protocol (RIP), see Configuring OSPF and RIP for VOSS.

IP addressing

An IP version 4 address consists of 32 bits expressed in dotted-decimal format (x.x.x.x). The IP version 4 address space is divided into classes, with classes A, B, and C reserved for unicast addresses and accounting for 87.5 percent of the 32-bit IP address space. Class D is reserved for multicast addressing. The following table lists the breakdown of IP address space by address range and mask.

Class	Address range	Mask	Number of addresses
Α	1.0.0.0 to 126.0.0.0	255.0.0.0	126
В	128.0.0.0 to 191.0.0.0	255.255.0.0	127 * 255
С	192.0.0.0 to 223.0.0.0	255.255.255.0	31 * 255 * 255
D	224.0.0.0 to 239.0.0.0	_	_

To express an IP address in dotted-decimal notation, you convert each octet of the IP address to a decimal number and separate the numbers by decimal points. For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary, has a different boundary point between the network and host portions of the address as illustrated in the following figure. The network portion is a network number field from 8 through 24 bits. The remaining 8 through 24 bits identify a specific host on the network.

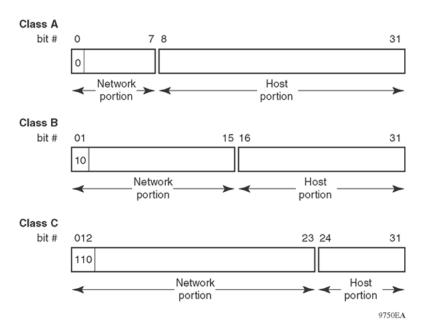


Figure 1: Network and host boundaries in IP address classes

Subnet Addressing

Subnetworks (or subnets) extend the IP addressing scheme an organization uses to one with an IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

You create a subnet address by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is found in the first octet of the host portion (10). A subnet mask is applied to the IP address and identifies the network and host portions of the address.

The following table illustrates how subnet masks used with class B and class C addresses can create differing numbers of subnets and hosts. This example includes the zero subnet, which is permitted on the switch.

Table 3: Subnet masks for class B and class C IP addresses

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
		Class B	
2	255.255.192.0	2	16 382
3	255.255.224.0	6	8 190
4	255.255.240.0	14	4 094
5	255.255.248.0	30	2 046
6	255.255.252.0	62	1 022
7	255.255.254.0	126	510

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts for each subnet
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1 022	62
11	255.255.255.224	2 046	30
12	255.255.255.240	4 094	14
13	255.255.255.248	8 190	6
14	255.255.255.252	16 382	2
	Class C		
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

You use variable-length subnet masking (VLSM) to divide your intranet into pieces that match your requirements. Routing is based on the longest subnet mask or network that matches. Routing Information Protocol version 2 and Open Shortest Path First are routing protocols that support VLSM.

Supernet Addressing and CIDR

A supernet, or classless interdomain routing (CIDR) address, is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed. You can use supernetting to address an entire block of class C addresses and avoid using large routing tables to track the addresses.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the class C addresses shown in the following figure. By adding the mask 255.255.128.0 to IP address 192.32.128.0, you aggregate the addresses 192.32.128.0 through 192.32.255.255 and 128 class C addresses use a single routing advertisement. In the bottom half of the following figure, you use 192.32.0.0/17 to aggregate the 128 addresses (192.32.0.0/24 to 192.32.127.0/24).

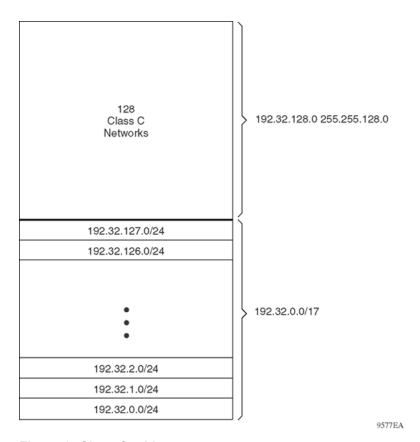


Figure 2: Class C address supernet

Another example is the block of addresses 192.32.0.0 to 192.32.7.0. The supernet address for this block is 11000000 00100000 00000, with the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an address and mask pair:

- The address is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- The mask is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/21.

Although classes prohibit using an address mask with the IP address, you can use CIDR to create networks of various sizes using the address mask. With CIDR, the routers outside the network use the addresses.

Loopback

Circuitless IP (CLIP) is a virtual (or loopback) interface that is not associated with a physical port. You can use the CLIP interface to provide uninterrupted connectivity to your device as long as a path exists to reach the device.

For example, as shown in the following figure, a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Use an interior Border Gateway Protocol (iBGP) session between two additional addresses, 195.39.128.1/30 (CLIP 1) and 195.39.281.2/30 (CLIP 2).

CLIP 1 and CLIP 2 represent the virtual CLIP addresses that you configure between R1 and R2. These virtual interfaces are not associated with the physical link or hardware interface, which permits the BGP session to continue as long as a path exists between R1 and R2. An IGP (such as OSPF) routes addresses that correspond to the CLIP addresses. After the routers learn all the CLIP addresses in the AS, the system establishes iBGP and exchanges routes.

The system advertises loopback routes to other routers in the domain either as external routes using the route-redistribution process, or after you enable OSPF in passive mode to advertise an OSPF internal route.

You can also use CLIP for PIM-SM, typically, as a Rendezvous Point (RP), or as a source IP address for sending SNMP traps and Syslog messages.

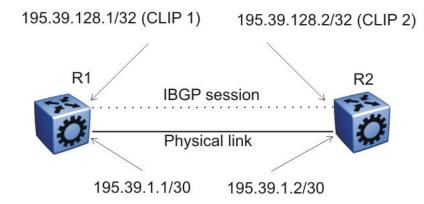


Figure 3: Routers with iBGP connections

The system treats the CLIP interface as an IP interface. The network associated with the CLIP is treated as a local network attached to the device. This route always exists and the circuit is always up because there is no physical attachment.

Static routes

Table 4: Static routing product support

Feature	Product	Release introduced
For configuration details, see Configuring IPv4 Routing for VOSS.		
Static routing	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1

Feature	Product	Release introduced
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50

A static route is a route to a destination IP address that you manually create.

The Layer 3 redundancy feature supports the creation of static routes to enhance network stability. Use the local next hop option to configure a static route with or without local next hop.

You can configure static routes with a next hop that is not directly connected, but that hop must be reachable. Otherwise, the static route is not enabled.

Layer 3 redundancy supports only address resolution protocol (ARP) and static route. Static ARP must configure the nonlocal next-hop of static routes. No other dynamic routing protocols provide nonlocal next-hop.



Note:

Static ARP entries are not supported for NLB Unicast or NLB Multicast operations.

You can use a default static route to specify a route to all networks for which no explicit routes exist in the forwarding information base or the routing table. This route has a prefix length of zero (RFC1812). You can configure the switch with a route through the IP static routing table.

To create a default static route, you must configure the destination address and subnet mask to 0.0.0.0.



Mote:

We recommend that you do not configure static routes on a DvR Leaf node unless the configuration is for reachability to a management network using a Brouter port.

Also, configuring the preference of static routes is not supported on a Leaf node.

Static Route Tables

A router uses the system routing table to make forwarding decisions. In the static route table, you can change static routes directly. Although the two tables are separate, the static route table manager entries are automatically reflected in the system routing table if the next-hop address in the static route is reachable, and if the static route is enabled.

The system routing table displays only active static routes with a best preference. A static route is active only if the route is enabled and the next-hop address is reachable (for example, if a valid ARP entry exists for the next hop).

You can enter multiple routes (for example, multiple default routes) that have different costs, and the routing table uses the lowest cost route that is available. However, if you enter multiple next hops for the same route with the same cost, the software does not replace the existing route. If you enter the same route with the same cost and a different next-hop, the first route is used. If the first route

becomes unreachable, the second route (with a different next-hop) is activated with no connectivity loss.

Black hole static routes

A black hole static route is a route with an invalid next hop, and the device drops data packets destined for this network.

While the router aggregates or injects routes to other routers, the router does not have a path to the aggregated destination. In such cases, the result is a black hole and a routing loop. To avoid routing loops, configure a black hole static route to the destination the router is advertising.

You can configure a preference value for a black hole route. However, you must configure that preference value appropriately so that when you want to use the black hole route, it is elected as the best route.

Before you add a black hole static route, perform a check to ensure that no other static route to that identical destination is enabled. If such a route exists, you cannot add the black hole route and an error message appears.

If you enable a black hole route, you cannot add another static route to that destination. You must first delete or disable the black hole route before you add a regular static route to that destination.

VLANs and routing

When traffic is routed on a virtual local area network (VLAN), an IP address is assigned to the VLAN and is not associated with a particular physical port. Brouter ports are VLANs that route IP packets and bridge nonroutable traffic in a single-port VLAN.

Virtual Routing Between VLANs

The switch supports wire-speed IP routing between VLANs. As shown in the following figure, VLAN 1 and VLAN 2 are on the same device, yet for traffic to flow from VLAN 1 to VLAN 2, the traffic must be routed.

When you configure routing on a VLAN, you assign an IP address to the VLAN, which acts as a virtual router interface address for the VLAN (a virtual router interface is not associated with a particular port). You can reach the VLAN IP address through the VLAN ports, and frames are routed from the VLAN through the gateway IP address. Routed traffic is forwarded to another VLAN within the device.

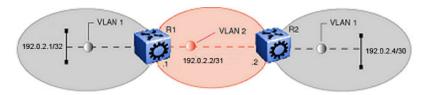


Figure 4: IP routing between VLANs

When Spanning Tree Protocol is enabled in a VLAN, the spanning tree convergence must be stable before the routing protocol begins. This requirement can lead to an additional delay in the IP traffic forwarding.

Because a port can belong to multiple VLANs (some of which are configured for routing on the device and some of which are not), a one-to-one correspondence no longer exists between the physical port and the router interface.

As with an IP address, virtual router interface addresses using Virtual Router Redundancy Protocol (VRRP) are also used for device management. For Simple Network Management Protocol (SNMP) or Telnet management, you can use virtual router interface address to access the device as long as routing is enabled on the VLAN.

Brouter Ports

The switch also supports brouter ports. A brouter port is a single-port VLAN that routes IP packets and bridges all nonroutable traffic. The difference between a brouter port and a standard IP protocol-based VLAN configured to route traffic is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic and still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

Because a brouter port is a single-port VLAN, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

The switch allows IP routing to be enabled on VLANs and brouter ports. For the maximum number of interfaces, see the Software scaling capabilities section of the Release Notes for VOSS.

Equal Cost Multipath

Table 5: Equal Cost Multiple Path for IPv4 product support

Feature	Product	Release introduced	
For configuration details, see Config	For configuration details, see Configuring IPv4 Routing for VOSS.		
Equal Cost Multiple Path (ECMP)	VSP 4450 Series	VSP 4000 4.0	
for IPv4	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VSP 8200 4.0	
	VSP 8400 Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 4.5	
	XA1400 Series	VOSS 8.0.50	

With Equal Cost Multipath (ECMP), the switch can determine up to eight equal-cost paths to the same destination prefix. You can use multiple paths for load sharing of traffic. These multiple paths provide faster convergence to other active paths in case of network failure. By maximizing load

sharing among equal-cost paths, you can use links between routers more efficiently when sending IP traffic. Equal Cost Multipath is formed using routes from the same source or protocol.

All IP ECMP routes that share the same combination of ECMP next hops consume the same ECMP group resource. The following list illustrates how shared next hops affect resource consumption:

- Prefix 3.1.1.0/16 is learned as an ECMP route with next hops A and B, and consumes one entry in the ECMP GROUP table.
- Prefix 4.1.1.0/16 is learned as an ECMP route with the same next hops, A and B. No additional resource is taken in the ECMP GROUP table.
- Prefix 5.1.1.0/16 is learned as an ECMP route with next hops B and C, and consumes one additional entry in the ECMP GROUP table.

ECMP is supported on both the Global Routing Table (GRT) and Virtual Routing and Forwarding (VRF).

The ECMP feature supports and complements the following protocols and route types:

- Border Gateway Protocol (BGP)
- Default route
- Intermediate System-to-Intermediate System (IS-IS)
- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Static route
- VRF

ECMP Pathlist

Use the ECMP Pathlist feature to control how many equal-cost paths to add to the routing manager for the same destination.



Product Notice: Not all products support Equal Cost Multipath Pathlist with Fabric Connect. For more information, see VOSS Feature Support Matrix.

Alternative routes

Table 6: Alternative routes product support

Feature	Product	Release introduced
Alternative routes for IPv4	VSP 4450 Series	VSP 4000 4.0
For configuration details, see	VSP 4900 Series	VOSS 8.1
Configuring IPv4 Routing for VOSS.	VSP 7200 Series	VOSS 4.2.1

Feature	Product	Release introduced
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
Alternative routes for IPv6	VSP 4450 Series	VOSS 5.1
For configuration details, see	VSP 4900 Series	VOSS 8.1
Configuring IPv6 Routing for	VSP 7200 Series	VOSS 5.1
<u>VOSS</u> .	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 5.1
	VSP 8400 Series	VOSS 5.1
	VSP 8600 Series	VSP 8600 6.2
	XA1400 Series	Not Supported

To avoid traffic interruption, you can globally enable the alternative routes feature so the router can use the next-best route, also known as an alternative route, if the best route becomes unavailable.

Routers learn routes to a destination through routing protocols. Routers maintain a routing table of the learned alternative routes sorted in order by route preference, route costs, and route sources. The first route on the list is the best route and the route that the router prefers to use.

The alternative route concept also applies between routing protocols. For example, if an OSPFv3 route becomes unavailable and an alternative RIPng route is available, the system activates the RIPng route without waiting for the update interval to expire.

Route Preference

On the switch, all standard routing protocols have default preference values that determine the routing priority of the protocol. The router uses default preferences to select the best route when a clash exists in preference between the protocols.

You can modify the global preference for a protocol to give the protocol a higher or lower priority than other protocols. If you change the global preference for a static route and all best routes remain best routes, only the local route tables change. However, if the protocol preference change causes best routes to no longer be best routes, the change affects neighboring route tables.

Important:

Changing route preferences is a process-intensive operation that can affect system performance and network reach while you perform route preference procedures. It is recommended that if you want to change preferences for static routes or routing protocols, do so when you configure routes or during a maintenance window.

If a router learns a route with the same network mask and cost values from multiple sources, the router uses the route preferences to select the best route to add to the forwarding database.

Note:

To modify the preference for a route, you do *not* need to disable a route before you edit the configuration.

Preferences for Static Routes

When you configure a static route on the switch, you can specify a global preference for the route. You can also specify an individual route preference that overrides the global static route preference. The preference value can be between 0 and 255, with 0 reserved for local routes and 255 representing an unreachable route.

Preferences for Dynamic Routes

You can modify the preference value for dynamic routes through route filtering and IP policies, and this value overrides the global preference for the protocol.

The following table shows the default preferences for routing protocols and route types. Use this table to help you modify the global preference value.

Table 7: Routing protocol default preferences

Protocol	Default preference
Local	0
Static	5
SPBM_L1	7
OSPF intra-area	20
OSPF inter-area	25
Exterior BGP	45
RIP/RIPng	100
OSPF external type 1	120
OSPF external type 2	125
IBGP	175
Staticv6	5
OSPFv3 intra-area	20
OSPFv3 inter-area	25
OSPFv3 external type 1	120
OSPFv3 external type 2	125

IP Source Routing

Table 8: IP Source Routing product support

Feature	Product	Release introduced	
For configuration details, see the fe	For configuration details, see the following documents:		
Configuring IPv4 Routing for VOSS			
Configuring IPv6 Routing for VO	<u>SS</u>		
IP Source Routing enable or	VSP 4450 Series	VOSS 5.1	
disable	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 5.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VOSS 5.1	
	VSP 8400 Series	VOSS 5.1	
	VSP 8600 Series	VSP 8600 4.5	
	XA1400 Series	Not Supported	

IP Source Routing allows the sender of a packet to specify the route that the packet must travel through the network. When the Source Route option is not enabled, the router chooses the primary routing path to send the packets. If IP Source Routing flag is on, the source host dictates the datapath for the packet to reach the destination using the information contained in the IP header.

The routing behavior in VSP 8600 Series is controlled by the datapath specified, and not by status of the IP Source Route. VSP 8600 Series inspects the packets only if the router itself is specified in the Source Routing. Otherwise, the switch forwards the packets to another router in the network using IP Routing, whether or not IP Source Route is enabled. You can use an ACL filter to block the datapath from forwarding any IP Source Routing packets.

IP Source Routing is considered as a security risk because it allows the users to specify their own path through the network outside of the primary forwarding path. This can cause packets to bypass the security devices. Therefore, the Source Routing is disabled by default.

Multihoming

The switch uses the multihoming feature to support clients or servers that have multiple IP addresses associated with a single MAC address. Multihomed hosts can be connected to port-based and policy-based VLANs.

The IP addresses associated with a single MAC address on a host must be in the same IP subnet.

Enabling or Disabling IPv4 ICMP Broadcast

Table 9: Internet Control Message Protocol product support

Feature	Product	Release introduced
Internet Control Message Protocol (ICMP)	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
For configuration details, see	VSP 7200 Series	VOSS 4.2.1
Configuring IPv4 Routing for VOSS.	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
ICMP broadcast and multicast enable or disable	VSP 4450 Series	VOSS 5.1
	VSP 4900 Series	VOSS 8.1
For configuration details, see the	VSP 7200 Series	VOSS 5.1
following documents:	VSP 7400 Series	VOSS 8.0
 Configuring IPv4 Routing for VOSS Configuring IPv6 Routing for VOSS 	VSP 8200 Series	VOSS 5.1
	VSP 8400 Series	VOSS 5.1
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	Not Supported

On IPv4 networks, a packet can be directed to an individual machine or broadcast to an entire network. When a packet is sent to an IP broadcast address from a machine on the local network, that packet is delivered to all machines on that network.

If a packet that is broadcast is an ICMP echo request packet, the machines on the network receive this ICMP echo request packet and send an ICMP echo reply packet back. When all the machines on a network respond to this ICMP echo request, the result can be severe network congestion or outages.

The switch always responds to IPv4 ICMP packets sent to a broadcast address. You can disable the processing of these IPv4 ICMP packets sent to the broadcast address. On disabling the ICMP broadcast processing, all the packets containing ICMP sent to the broadcast addresses will be dropped when the packets reach the control plane.

You can disable or enable the IPv4 ICMP broadcast processing at the VRF level.

IP routing configuration using the CLI

Configure the IP router interface so that you can configure and use routing protocols and features on the interface. This section contains instructions for both the Global Router and Virtual Router Forwarding (VRF) instances.

Enabling routing globally or on a VRF instance

Use IP forwarding (routing) on a global level so that the device supports routing. You can use the IP address of an interface for IP-based network management.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Activate IP forwarding:

```
ip routing
```

3. View the forwarding configuration:

```
show ip routing [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

Activate IP forwarding and view the forwarding configuration:

```
Switch: 1>enable
Switch: 1#configure terminal
Switch:1(config) #router vrf green
Switch:1(router-vrf) #ip routing
Switch:1(router-vrf) #show ip routing
______
                          IP - GlobalRouter
IP Forwarding is enabled
IP ECMP feature is disabled
Maximum ECMP paths number is 1
ECMP 1 pathlist:
ECMP 2 pathlist :
ECMP 3 pathlist:
ECMP 4 pathlist :
ECMP 5 pathlist :
ECMP 6 pathlist :
ECMP 7 pathlist :
ECMP 8 pathlist :
Gratuitous-Arp : enable
IP Alternative Route feature is enabled
```

```
IP More Specific Non Local Route feature is disabled
IP ICMP Unreachable Message is disabled
IP Supernetting is disabled
IP Icmp-echo-broadcast-request is enabled

IP Default TTL is 255 seconds
IP ARP life time is 360 minutes
IP Source Route Option is disabled
```

Variable definitions

Use the data in the following table to use the show ip routing command.

Table 10: Variable definitions

Variable	Value
vrf WORD<0-16>	Specifies a VRF instance by VRF name.
vrfids WORD<0-512>	Specifies a VRF instance by VRF number.

Enabling routing on an IP interface

About this task

You can enable or disable routing capabilities on a VLAN or brouter port.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```

Note:

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

2. Enable routing:

routing enable

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #interface gigabitethernet 1/2
Switch:1(config-if) #routing enable
```

Deleting a dynamically learned route

About this task

Delete a dynamically learned route from the routing table if you do not want the switch to use the route. Exercise caution when you delete entries from the routing table.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. View IP route information:

```
show ip route [<A.B.C.D>] [-s default|-s <A.B.C.D/X>] [alternative] [count-summary] [spbm-nh-as-mac][preference] [vrf WORD<0-16>] [vrfids WORD<0-512>] [static]
```

3. Delete the dynamically learned route:

```
no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> dynamic
```

Example

Delete a dynamically learned route:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#no ip route 192.0.2.32 255.255.255.0 198.51.100.31 dynamic
```

Variable definitions

Use the data in the following table to use the show ip route commands.

Table 11: Variable definitions

Variable	Value
<a.b.c.d></a.b.c.d>	Specifies the IP address of the route to the network.
alternative	Displays the alternative routes.
count-summary	Displays a summary of the number of routes learned from each routing protocol for each VRF.
preference	Displays the route preference.
-s <a.b.c.d x=""></a.b.c.d>	Indicates the IP address and subnet mask for which to display routes.
-s default	Indicates the default subnet.
static	Displays the static route information.
vrf WORD<0-16>	Displays the route for a particular VRF.
vrfids WORD<0-512>	Displays the route for a particular VRF number.
spbm-nh-as-mac	Displays the spbm route next hop as mac.

Use the data in the following table to use the no ip route command.

Table 12: Variable definitions

Variable	Value
<a.b.c.d> <a.b.c.d> <a.b.c.d></a.b.c.d></a.b.c.d></a.b.c.d>	Specifies the IP address, the subnet mask, and the next-hop IP address, respectively.
dynamic	Specifies that a dynamic route is to be deleted.
enable	Disables the route.
local-next-hop enable	Disables the local-next-hop option.
preference	Deletes the value of the route preference.
next-hop-vrf WORD<0-16>	Specifies the name of the next-hop VRF router.

Configuring IP route preferences

Before you begin

Disable ECMP before you configure route preferences

Important:

Changing route preferences can affect system performance and network accessibility while you perform the procedure. You must therefore change a prefix list or a routing protocol *before* you activate the protocols.

About this task

Configure IP route preferences to give preference to routes learned for a specific protocol. You must disable ECMP before you configure route preferences.

To configure route preferences for a VRF, access VRF Router Configuration mode, rather than Global Configuration mode.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Configure the route preference:

```
ip route preference protocol <static|ospf-intra|ospf-inter|ebgp|
ibgp|rip|ospf-extern1|ospf-extern2|spbm-level1> <0-255>
```

3. Confirm that the configuration is correct:

```
show ip route preference [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Configure the route preference to SPBM Level 1 and confirm the configuration is correct.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #ip route preference protocol spbm-level1 7
Switch:1(config) #show ip route preference

IP Route Preference - GlobalRouter

PROTOCOL DEFAULT CONFIG

LOCAL 0 0 0
STATIC 5 5
SPBM_L1 7 7
OSPF_INTRA 20 20
OSFF_INTER 25 25
EBGP 45 45
RIP 100 100
OSFF_E1 120 120
OSFF_E1 120 120
OSFF_E2 125 125
IBGP 175 175
```

View the route preference configuration for a specific VRF, for example 444.

Variable definitions

Use the data in the following table to use the ip route preference protocol command.

Variable	Value
ebgp	Protocol type eBGP
ibgp	Protocol type iBGP
ospf-extern1	Protocol type ospf-extern1
ospf-extern2	Protocol type ospf-extern2
ospf-intra	Protocol type ospf-intra

Variable	Value
ospf-inter	Protocol type ospf-inter
rip	Protocol type rip
spbm-level1	Protocol type spbm-level1
static	Protocol type static

Flushing routing tables by VLAN or port

About this task

For administrative and troubleshooting purposes, flush the routing tables.

To flush tables on a VRF instance for a port or VLAN, ensure that the VRF is associated with the port or VLAN.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
```



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

2. Flush the routing tables:

```
action flushIp
```

Example

Flush the routing tables:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabitethernet 3/6
Switch:1(config-if)#action flushIp
```

Assigning an IP address to a port

Assign an IP address to a port so that it supports routing operations.

About this task

Use a brouter port to route IP packets and to bridge all nonroutable traffic. The routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the

blocking state for nonroutable traffic and still route IP traffic. This feature removes interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

If an IP interface is configured without specifying the VRF instance, it maps to VRF 0 by default.

Use the vrf parameter to associate the port or VLAN with a VRF instance.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
```

Note:

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

2. Assign an IP address to the port:

```
brouter port \{slot/port \ [-slot/port] \ [,...]\} vlan <2-4059> subnet <A.B.C.D/X> [mac-offset \ <MAC-offset> | [name \ WORD \ <0-64>]
```

3. If required, associate the port with a VRF:

```
vrf WORD<1-16>
```

4. Confirm that the configuration is correct:

```
show brouter [<1-4084>]
```

Example

Assign an IP address to a port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #interface gigabitethernet 1/11
Switch:1(config-if) #brouter port 1/11 vlan 2202 subnet 47.17.10.31/255.255.255.0
```

Variable Definitions

Use the data in the following table to use the **brouter port** command.

Variable	Value
mac-offset <mac-offset></mac-offset>	Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated. Different hardware platforms support different ranges. To see which range is available on the switch, use the CLI command completion Help.
name WORD <0-64>	Specifies the IP address name in the range of 0 to 64 characters.

Variable	Value
{slot/port[/sub-port]}	Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
subnet <a.b.c.d x=""></a.b.c.d>	Specifies the IP address and subnet mask (0–32).
<2-4059>	Specifies the VLAN ID that is used if the port is tagged (802.1q encapsulation). The VLAN ID is unique to the switch and is not used if the port is untagged.

Use the data in the following table to use the **show brouter** command.

Variable	Value
<1-4084>	Specifies the VLAN ID that is used if the port is tagged (802.1q encapsulation). The VLAN ID is unique to the switch and is not used if
	the port is untagged.

Assign an IP Address to a VLAN

Assign an IP address to a VLAN so that it supports routing operations.

Before you begin

- · You must create the VLAN.
- Activate IP forwarding globally.

About this task

If an IP interface is configured without specifying the VRF instance, it maps to VRF 0 by default. Use the vrf parameter to associate the VLAN with a VRF instance.

Important:

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4059>
```

2. Assign an IP address to a VLAN:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D> dvr-one-ip name WORD < 0-64>
```

3. (Optional) If required, associate the VLAN with a VRF:

```
vrf WORD<1-16>
```

Example

```
Switch:1> enable

Switch:1# configure terminal

Switch:1(config)# interface vlan 2

Switch:1(config-if)# ip address 192.0.2.5 255.255.255.0 dvr-one-ip name Boston
```

Variable Definitions

The following table defines parameters for the ip address command.

Variable	Value
<a.b.c.d x=""> <a.b.c.d> <a.b.c.d></a.b.c.d></a.b.c.d></a.b.c.d>	Specifies the IP address and subnet mask in the format A.B.C.D/X or A.B.C.D A.B.C.D.
dvr-one-ip	Specifies that the IP address will be used as the DvR gateway IP address and will be used by all other DvR Controllers for the DvR VLAN subnet.
name WORD <0-64>	Specifies the name associated with the IP address on a VLAN.
	This parameter does not apply to all hardware platforms.

The following table defines parameters for the vrf command.

Variable	Value
WORD<0-16>	Specifies the VRF of the VLAN.

Configure an IP Address for a Segmented Management Instance

Use this task to add an IPv4 or IPv6 address to a Management Instance.

Before you begin

• Ensure the IP address you plan to assign is not in use by an existing VLAN or CLIP IP subnet configured on the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enter the configuration mode for the Management Instance:

```
mgmt {clip | oob | vlan}
```

3. Add an IPv4 address:

```
ip address {A.B.C.D [A.B.C.D] | A.B.C.D/X}
```

4. Add an IPv6 address:

```
ipv6 address WORD<0-255>
```

Example

Add an IPv4 address to the VLAN Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #mgmt vlan
Switch:1(mgmt:vlan) #ip address 192.0.2.12/24
```

Add an IPv4 address to the OOB Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #mgmt oob
Switch:1(mgmt:oob) #ip address 192.0.2.12 255.255.255.0
```

Add an IPv6 address to the CLIP Management Instance:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #mgmt clip
Switch:1(mgmt:clip) #ipv6 address 2001:DB8::1/128
```

Configure an IP Address for the Management Port

Note:

This procedure only applies to VSP 8600 Series.

About this task

Configure the IP address for the management port so that you can remotely access the device using the management port. The management port runs on a dedicated VRF and it is recommended that you redirect all commands that are run on the management port to its VRF.

The configured IP subnet has to be globally unique because the management protocols can go through in-band or out-of-band ports.

Procedure

1. Enter mgmtEthernet Interface Configuration mode:

```
enable
configure terminal
interface mgmtEthernet <mgmt | mgmt2>
```

2. Configure the IP address and mask for the management port:

```
ip address <A.B.C.D> <A.B.C.D>
```

3. Show the complete network management information:

```
show interface mgmtEthernet
```

4. Show the management IP interface information:

show ip interface vrf mgmtrouter

Example

Configure the IP address for the management port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #interface mgmtethernet mgmt
Switch:1(config-if) #ip address 192.0.2.31 255.255.255.0
Switch:1(config-if) #show interface mgmtethernet

Port Interface

PORT LINK PORT PHYSICAL STATUS
NUM INDEX DESCRIPTION TRAP LOCK MTU ADDRESS ADMIN OP
ERATE

mgmt 64 mgmtethernet true false 1522 192.0.2.31 up up
```

Variable definitions

Use the data in the following table to use the ip address command.

Table 13: Variable definitions

Variable	Value
<a.b.c.d> <a.b.c.d></a.b.c.d></a.b.c.d>	Specifies the IP address followed by the subnet mask.

Viewing IP addresses for all router interfaces

About this task

Perform the following procedure to display information about all IP interfaces configured on the device.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Show the IP interfaces and addresses on the device:

show ip interface

Example

Show the IP interfaces and addresses on the device:

```
Switch:1>enable
Switch:1#show ip interface
```

IP Interface - GlobalRouter						
INTERFACE	IP	NET	BCASTADDR	REASM	VLAN	BROUTER
	ADDRESS	MASK	FORMAT	MAXSIZE	ID	PORT
Port1/6	192.0.2.6	255.255.255.	0 ones	1500	200	true
Vlan100	192.0.2.5	255.255.255.		1500	100	false
Vlan4000	198.51.100.21	255.255.255.		1500	400	O false

Variable definitions

Use the data in the following table to show ip interface command.

Table 14: Variable definitions

Variable	Value
gigabitethernet	Displays IP interface information for Gigabit Ethernet ports.
vrf	Displays interface information for a particular VRF.
vrfids	Displays interface information for particular VRF IDs.

Configure IP Routing Globally or for a VRF

Configure the IP routing protocol stack to specify which routing features the device can use. You can configure global parameters before or after you configure the routing protocols.

About this task

To configure IP routing globally for a VRF instance, use VRF Router Configuration mode rather than Global Configuration mode.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Configure the default TTL for all routing protocols to use:

```
ip ttl <1-255>
```

This value is placed into routed packets that have no TTL specified.

3. Activate the alternative route feature globally:

```
ip alternative-route
```

4. Configure the remaining global parameters as required.

Variable Definitions

The following table defines parameters for the ip command.

Variable	Value	
alternative-route	Enables or disables the alternative route feature. The default value is enabled.	
	If the alternative-route parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are readded.	
	The default form of this command is default ip alternative-route. The no form of this command is no ip alternative-route.	
max-routes-trap enable	Enables the device to send a trap after the maximum number of routes is exceeded.	
	The no form of this command is no max-routes-trap enable. The default form of this command is default max-routes-trap enable.	
more-specific-non-local-route	Enables the more-specific-non-local-route feature. If enabled, the device can enter a more-specific nonlocal route into the routing table. The default is disabled.	
	The default form of this command is default ip more-specific-non-local-route. The no form of this command is no ip more-specific-non-local-route.	
routing	Enables routing.	
	The no form of this command is no ip routing.	
supernet	Enables or disables supernetting.	
	If you globally enable supernetting, the device can learn routes with a route mask of less then eight bits. Routes with a mask length less than eight bits cannot have ECMP paths, even if the ECMP feature is globally enabled. The default is disabled.	
	The default form of this command is default ip supernet. The no form of this command is no ip supernet.	
ttl <1-255>	Configures the default time-to-live (TTL) value for a routed packet. The TTL is the maximum number of seconds before a packet is discarded. The default value of 255 is used whenever a time is not supplied in the datagram header.	
	The default form of this command is default ip ttl.	

The following table defines parameters for the ip icmp commands.

Variable	Value	
unreachable	Enables the device to send ICMP unreachable messages. When	
	enabled, this variable generates Internet Control Message Protocol	

Variable	Value	
	(ICMP) network unreachable messages if the destination network is not reachable from this router. These messages help determine if the device is reachable over the network. The default is disabled.	
	The default form of this command is default ip icmp unreachable.	

Configure ECMP

Enable Equal Cost MultiPath (ECMP) to permit routers to determine up to eight equal-cost paths to the same destination prefix. You can use the multiple paths for load-sharing of traffic, which provides fast convergence to alternative paths. By maximizing load sharing among equal-cost paths, you can maximize the efficiency of links between routers.

About this task

To configure ECMP for a VRF instance, after you enable ECMP globally, use VRF Router Configuration mode rather than Global Configuration mode.

Different hardware platforms can support a different number of ECMP paths. For more information, see Release Notes for VOSS.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable ECMP globally:

ip ecmp

3. **(Optional)** Configure the maximum number of ECMP paths:

```
ip ecmp max-path <ECMP-Paths>
```

4. **(Optional)** Configure a prefix-list for the target destination:

```
ip prefix-list WORD<1-64> <A.B.C.D/X> [ge <0-32>] [le <0-32>]
```

5. **(Optional)** Configure an ECMP pathlist to specify routes with the required number of paths:

```
ip ecmp pathlist-<1-8> WORD<1-64>
```

6. (Optional) Return to Privileged EXEC mode:

end

7. **(Optional)** Apply changes to all ECMP pathlist configurations:

```
ip ecmp pathlist-apply [vrf WORD<1-16>]
```

Example

Define which IP prefixes use ECMP and which do not.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip ecmp
Switch:1(config)#ip prefix-list ecmpAllowed 192.0.2.0/24 ge 24 le 24
Switch:1(config)#ip prefix-list ecmpDenied 0.0.0.0/0 ge 0 le 32
Switch:1(config)#ip ecmp pathlist-2 ecmpAllowed
Switch:1(config)#ip ecmp pathlist-1 ecmpDenied
Switch:1(config)#end
Switch:1#ip ecmp pathlist-apply
```

Variable Definitions

The following table defines parameters for the ip ecmp command.

Variable	Value	
max-path <ecmp-paths></ecmp-paths>	Specifies the maximum number of ECMP paths. Different hardware platforms can support a different number of ECMP paths. For more information on the maximum number of ECMP paths supported on the switch, see the scaling information in Release Notes for VOSS.	
pathlist-1 WORD<0-64>	Specifies one equal-cost path to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-2 WORD<0-64>	Specifies up to two equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-3 WORD<0-64>	Specifies up to three equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-4 WORD<0-64>	Specifies up to four equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-5 WORD<0-64>	Specifies up to five equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-6 WORD<0-64>	Specifies up to six equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-7 WORD<0-64>	Specifies up to seven equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-8 WORD<0-64>	Specifies up to eight equal-cost paths to the same destination prefix. To remove the policy, enter a blank string.	
pathlist-apply [vrf WORD<1-16>]	Applies the pathlist configuration changes. You can optionally specify a VRF name.	

The following table defines parameters for the ip prefix-list command.

Variable	Value
WORD<0-64>	Specifies the prefix list name.

Variable	Value
<a.b.c.d x=""></a.b.c.d>	Specifies the IP address and network mask in one of the following formats:
	• a.b.c.d/x
	• a.b.c.d/x.x.x.x
	default
ge <0-32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
le <0-32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.

Configuring Static Routes

Before you begin

· Ensure no black hole static route exists.

About this task

Configure a static route when you want to manually create a route to a destination IP address.



When you configure a static route with a next-hop-vrf context, you can specify a next-hop IP address that is a locally owned VRRP IP address of the system itself. However, this is not a supported configuration. The best practice is to implement an alternative method of inter-vrf route sharing, such as route redistribution or ISIS accept polices.

If a black hole route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

For route scaling information and for information on the maximum number of static routes supported on your hardware platform, see <u>Release Notes for VOSS</u>.



As a best practice, do not configure static routes on a DvR Leaf node unless the configuration is for reachability to a management network using a Brouter port.

You cannot configure the preference of static routes on a Leaf node.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Create an IP static route:

```
ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> weight <1-65535>
```

3. Enable an IP static route:

```
ip route <A.B.C.D> <A.B.C.D> <A.B.C.D> enable
```

- 4. Use the following variable definitions table to configure other static route parameters as required.
- 5. View existing IP static routes for the device, or for a specific network or subnet:

```
show ip route static
```

6. Delete a static route:

```
no ip route <A.B.C.D> <A.B.C.D> <A.B.C.D>
```

Example

Create an IP static route, enable a static route, and view the existing IP static routes for the device, or for a specific network or subnet.

Variable definitions

Use the data in the following table to use the ip route command.

Table 15: Variable definitions

Variable	Value
<a.b.c.d> <a.b.c.d> <a.b.c.d></a.b.c.d></a.b.c.d></a.b.c.d>	The first and second <a.b.c.d> specify the IP address and mask for the route destination. The third <a.b.c.d> specifies the IP address of the next-hop router (the next router at which packets must arrive on this route). When you create a black hole static route, configure this parameter to 255.255.255.255 as the IP address of the router through which the specified route is accessible.</a.b.c.d></a.b.c.d>
disable	Disables a route to the router or VRF.
enable	Adds a static route to the router or VRF.
	The no form of this command is no ip route <a.b.c.d></a.b.c.d>
	<a.b.c.d> <a.b.c.d> enable.</a.b.c.d></a.b.c.d>

Variable	Value
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> enable.</a.b.c.d></a.b.c.d>
local-next-hop enable	Enables the local next hop for this static route. The default form of this command is default ip route <a.b.c.d> <a.b.c.d> <a.b.c.d></a.b.c.d></a.b.c.d></a.b.c.d>
	The no form of this command is no ip route <a.b.c.d> <a.b.c.d> <a.b.c.d> local-next-hop enable.</a.b.c.d></a.b.c.d></a.b.c.d>
next-hop-vrf <word 0-16=""></word>	Specifies the next-hop VRF instance by name.
	After you configure the next-hop-vrf parameter, the static route is created in the local VRF, and the next-hop route is resolved in the next-hop VRF instance (next-hop-vrf).
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> next-hop-vrf <word 0-16="">.</word></a.b.c.d></a.b.c.d>
	The no form of this command is no ip route <a.b.c.d> <a.b.c.d> <a.b.c.d> next-hop-vrf <word 0-16="">.</word></a.b.c.d></a.b.c.d></a.b.c.d>
weight <1-65535>	Specifies the static route cost.
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> weight.</a.b.c.d></a.b.c.d>
name <1-64>	Specifies the name of the static route. You can name the route before or after it is created.
	Only 32 characters display. The tilde (~) symbol indicates that the name is truncated.
preference <1-255>	Specifies the route preference.
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> oreference.</a.b.c.d></a.b.c.d>

Use the data in the following table to use the show ip route static command.

Table 16: Variable definitions

Variable	Value
<a.b.c.d></a.b.c.d>	Specifies the route by IP address.
-s { <a.b.c.d> <a.b.c.d> default}</a.b.c.d></a.b.c.d>	Specifies the route by IP address and subnet mask.
vrf WORD<0-16>	Specifies a VRF by name.
vrfids WORD<0-512>	Specifies a range of VRF IDs.
name <1-64>	Specifies the name of the static route. You can name the route before or after it is created.
	Only 32 characters display. The tilde (~) symbol indicates that the name is truncated.

Configuring a black hole static route

About this task

Configure a black hole static route to the destination a router advertises to avoid routing loops after the router aggregates or injects routes to other routers.

If a black hole route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Create a black hole static route:

```
ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 weight <1-65535>
```

3. Enable a black hole static route:

```
ip route <A.B.C.D> <A.B.C.D> 255.255.255.255 enable [next-hop-vrf WORD < 0-16 > 1
```

4. Configure other black hole static route parameters as required.

When you specify a route preference, appropriately configure the preference so that when the black-hole route is used, it is elected as the best route.

Example

Create a black hole static route and enable the black hole static route.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip route 192.0.2.0 255.255.0.0 255.255.255 weight 200
Switch:1(config)#ip route 192.0.2.0 255.255.0.0 255.255.255 enable
```

Variable definitions

Use the data in the following table to use the ip route command.

Table 17: Variable definitions

Variable	Value
<a.b.c.d></a.b.c.d>	The first and second < <i>A.B.C.D</i> > specify the IP address and mask for the route destination. 255.255.255.255 is the destination of the black hole route.

Variable	Value
enable	Adds a static route to the router or VRF.
	The no form of this command is no ip route <a.b.c.d> <a.b.c.d> 255.255.255.255 enable.</a.b.c.d></a.b.c.d>
local-next-hop enable	Enables the local next hop for this static route.
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> <a.b.c.d> enable.</a.b.c.d></a.b.c.d></a.b.c.d>
	The no form of this command is no ip route <a.b.c.d> <a.b.c.d> 255.255.255.255 local-next-hop enable.</a.b.c.d></a.b.c.d>
next-hop-vrf WORD<0-16>	Specifies the next-hop VRF instance by name.
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> 255.255.255 next-hop-vrf <word 0-16="">.</word></a.b.c.d></a.b.c.d>
	The no form of this command is no ip route <a.b.c.d> <a.b.c.d> 255.255.255.255 next-hop-vrf <word 0-16="">.</word></a.b.c.d></a.b.c.d>
weight <1-65535>	Specifies the static route cost.
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> 255.255.255 weight.</a.b.c.d></a.b.c.d>
preference <1-255>	Specifies the route preference.
	The default form of this command is default ip route <a.b.c.d> <a.b.c.d> 255.255.255 preference.</a.b.c.d></a.b.c.d>

Configuring a default static route

About this task

The default route specifies a route to all networks for which there are no explicit routes in the forwarding information base or the routing table. This route has a prefix length of zero (RFC 1812). You can configure the switch with a static default route, or they can learn it through a dynamic routing protocol.

To create a default static route, configure the destination address and subnet mask to 0.0.0.0.

Note:

When you configure a static route with a next-hop-vrf context, you can specify a next-hop IP address that is a locally owned VRRP IP address of the system itself. However, this is not a supported configuration. The best practice is to implement an alternative method of inter-vrf route sharing, such as route redistribution or ISIS accept polices.

Note:

As a best practice, do not configure static routes on a DvR Leaf node unless the configuration is for reachability to a management network using a Brouter port.

You cannot configure the preference of static routes on a Leaf node.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Create a default static route:

```
ip route 0.0.0.0 0.0.0.0 <A.B.C.D> weight <1-65535>
```

3. Enable a default static route:

```
ip route 0.0.0.0 0.0.0.0 <A.B.C.D> enable [next-hop-vrf WORD<0-16>]
```

4. Configure other default static route parameters as required.

Example

Create a default static route and enable the default static route.

```
Switch:1>enable
Switch:1>configure terminal
Switch:1(config) #ip route 0.0.0.0 0.0.0.0 192.0.2.128 weight 100
Switch:1(config) #ip route 0.0.0.0 0.0.0.0 192.0.2.128 enable
```

Variable definitions

Use the data in the following table to use the ip route command.

Table 18: Variable definitions

Variable	Value
<a.b.c.d></a.b.c.d>	<a.b.c.d></a.b.c.d> specifies the IP address of the next-hop router (the next router at which packets must arrive on this route).
enable	Adds a static or default route to the router or VRF.
	The no form of this command is no ip route 0.0.0.0 0.0.0.0 <a.b.c.d> enable.</a.b.c.d>
local-next-hop enable	Enables the local next hop for this static route.
	The default form of this command is default ip route 0.0.0.0 0.0.0.0 <a.b.c.d> local-next-hop enable.</a.b.c.d>
	The no form of this command is no ip route 0.0.0.0 0.0.0.0 <a.b.c.d> local-next-hop enable.</a.b.c.d>
next-hop-vrf WORD<0-16>	Specifies the next-hop VRF instance by name.
	The default form of this command is default ip route 0.0.0.0 0.0.0.0 <a.b.c.d> next-hop-vrf WORD<0-16>.</a.b.c.d>

Variable	Value
	The no form of this command is no ip route 0.0.0.0 0.0.0.0 <a.b.c.d> next-hop-vrf WORD<0-16>.</a.b.c.d>
weight <1-65535>	Specifies the static route cost.
	The default form of this command is default ip route 0.0.0.0 0.0.0.0 <a.b.c.d> weight.</a.b.c.d>
preference <1-255>	Specifies the route preference.
	The default form of this command is default ip route 0.0.0.0 0.0.0.0 <a.b.c.d> preference.</a.b.c.d>

Enabling ICMP Router Discovery globally

About this task

Enable Router Discovery globally so that the device supports Router Discovery. Use ICMP Router Discovery to enable hosts attached to the broadcast network to discover the IP addresses of their neighboring routers.

If you enable ICMP Router Discovery globally, you automatically enable it for all VLANs. If you do not require ICMP Router Discovery on a specific VLAN, you must manually disable the feature.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Enable ICMP Router Discovery on the device:

```
ip irdp
```

3. Confirm that Router Discovery is enabled:

```
show ip irdp [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

Enable ICMP router discovery on the device and confirm that router discovery is enabled.

Variable definitions

Use the data in the following table to show ip irdp command.

Table 19: Variable definitions

Variable	Value
interface	Displays route discovery interface information.
vrf WORD<0-16>	Displays route discovery for particular VRF.
vrfids WORD<0-512>	Displays route discovery for particular VRF IDs.

Enabling or disabling IPv4 ICMP broadcast globally

On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when they reach the control plane.

About this task

Use these commands to enable or disable the IPv4 ICMP broadcast feature on the global router.

Procedure

1. Enter Global Configuration mode:

enable
configure terminal

2. Enable IPv4 ICMP broadcast feature, enter:

ip icmp echo-broadcast-request

3. Disable IPv4 ICMP broadcast feature, enter:

no ip icmp echo-broadcast-request

4. Set the IPv4 ICMP broadcast feature to default state, enter:

default ip icmp echo-broadcast-request



By default, the IPv4 ICMP broadcast feature is enabled.

5. View the IPv4 ICMP broadcast feature state:

show ip routing

Enabling or disabling IPv4 ICMP broadcast per VRF

On disabling the ICMP broadcast processing, all the packets containing ICMP sent to broadcast addresses, will be dropped when they reach the control plane.

About this task

Use these commands to enable or disable the IPv4 ICMP broadcast feature on the VRF router.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Enable IPv4 ICMP broadcast feature, enter:

```
ip icmp echo-broadcast-request
```

3. Disable IPv4 ICMP broadcast feature, enter:

```
no ip icmp echo-broadcast-request
```

4. Set the IPv4 ICMP broadcast feature to default state, enter:

```
default ip icmp echo-broadcast-request
```



By default, the IPv4 ICMP broadcast feature is enabled.

5. View the IPv4 ICMP broadcast feature state:

```
show ip routing
```

Configuring Router Discovery on a port or VLAN

Enable Router Discovery so that the device forwards Router Discovery Advertisement packets to the VLAN or port.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```



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

2. Specify the address placed in advertisement packets:

```
ip irdp address <A.B.C.D>
```

3. Enable the interface to send the advertisement packets:

ip irdp multicast

4. Configure other Router Discovery parameters for the interface as required.

Example

Log on to the GigabitEthernet Interface mode:

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/16
```

Specify the address placed in advertisement packets to the all-systems multicast address:

```
Switch:1(config-if) # ip irdp address 244.0.0.1
```

Enable the interface to send the advertisement packets:

```
Switch:1(config-if) # ip irdp multicast
```

Configure the lifetime for advertisements:

```
Switch:1(config-if) # ip irdp holdtime 180
```

Variable definitions

Use the data in the following table to use the ip irdp command.

Table 20: Variable definitions

Variable	Value
address <a.b.c.d></a.b.c.d>	Specifies the IP destination address use for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.
	The default address is 255.255.255.
	The default form of this command is default ip irdp address.
holdtime <4-9000>	Configures the lifetime for advertisements. The default form of this command is default ip irdp holdtime.
maxadvertinterval <4-1800>	Specifies the maximum time (in seconds) that elapses between unsolicited router advertisement transmissions from the router interface. The default is 600 seconds.
	The default form of this command is default ip irdp maxadvertinterval.
minadvertinterval <3-1800>	Specifies the minimum time (in seconds) that elapses between unsolicited router advertisement transmissions from the interface. The range is 3 seconds to maxadvertinterval.
	The default is 450 seconds.
	The default form of this command is default ip irdp minadvertinterval.

Variable	Value
multicast	Specifies if multicast advertisements are sent. The no form of this command is no ip irdp multicast.
preference <-2147483648-2147483647>	Specifies the preference (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet The default is 0.
	The default form of this command is default ip irdp preference.

Configuring a CLIP Interface

About this task

Configure a circuitless IP (CLIP) interface to provide a virtual interface that is not associated with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to your device.

For scaling information and for information on the maximum number of CLIP interfaces you can configure on your device, see <u>Release Notes for VOSS</u>.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create or access a CLIP interface:

```
interface loopback <1-256>
```

<1-256> indicates the identification number for the CLIP.

The command prompt changes to indicate you now access the Loopback Interface Configuration mode.

3. Configure an IP address and name for the interface:

```
ip address [<1-256>] <A.B.C.D/X> [vrf WORD<0-16>] [name WORD<0-64>]
```

4. Enable OSPF on the CLIP interface:

```
ip ospf [<1-256>] [vrf WORD<1-16>]
```

You can configure other protocols on the CLIP interface; OSPF is the most common. See the following variable definitions table for other options.

5. View the IP address on the CLIP interface:

```
show ip interface
```

Example

Create a CLIP interface, and enable OSPF on the CLIP interface.

```
Switch:1>enable
Switch: 1#configure terminal
Switch:1(config) #interface loopback 3
Switch:1(config-if) #ip address 23.23.23.23 255.255.255.0 name Canada
Switch:1>show ip interface
IP Interface - GlobalRouter
______
      IPNETBCASTADDRREASMVLANBROUTERIPSECADDRESSMASKFORMATMAXSIZEIDPORTSTATE
INTERFACE IP
______
Port1/2
      10.3.4.2
              255.255.255.0 ones
                            1500
                                    true
Boston
All 6 out of 6 Total Num of IP interfaces displayed
```

Variable Definitions

Use the data in the following table to use the ip commands.

Table 21: Variable definitions

Variable	Value
address [<1-256>] <a.b.c.d x=""> [vrf WORD<0-16>][name WORD<0-64]</a.b.c.d>	Specifies the IP address for the CLIP interface.
	<1-256> specifies the interface.
-	<a.b.c.d x=""> specifies the IP address and mask (0–32).</a.b.c.d>
	vrf WORD<0-16> specifies an associated VRF by name.
	The no form of this command is no ip address [<1-32>] <a.b.c.d> [vrf WORD<0-16>].</a.b.c.d>
	name WORD<0-16> specifies a name for the IP address.
area <1-256> <a.b.c.d> [vrf</a.b.c.d>	Designates an area for the CLIP interface.
WORD<0-16>]	vrf WORD<0-16> specifies an associated VRF by name
	The default form of this command is default ip area <1-256> <a.b.c.d> [vrf WORD<0-16>]. The no form of this command is no ip area <1-256> vrf WORD<0-16>].</a.b.c.d>

Configure BFD on an IPv4 Interface

BFD for IPv6 interfaces is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

About this task

Use the following procedure to enable and to configure Bidirectional Forwarding Detection (BFD) on an IPv4 interface. All interface configuration is performed at the VLAN, GigabitEthernet, or Loopback level.



Enabling BFD on an interface does not establish a BFD session. To establish a BFD session. you must enable BFD globally and at the application level.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
```

followed by one of the following:

- interface GigabitEthernet {slot/port[/sub-port][-slot/port[/subport]][,...]}
- interface loopback <1-256>
- interface vlan <1-4059>



■ Note:

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

2. Enable BFD on an interface:

```
ip bfd enable
```

3. **(Optional)** Configure the transmit interval:

```
ip bfd interval <100-65335>
```

4. (Optional) Configure the minimum receive interval:

```
ip bfd min-rx <100-65335>
```

5. (Optional) Configure the multiplier:

```
ip bfd multiplier <1-20>
```

6. (Optional) In GigabitEthernet Interface Configuration mode, you can configure a value for port:

```
ip bfd port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

7. (Optional) In VLAN Interface Configuration mode, you can configure a value for VLAN:

ip bfd vlan <1-4094>

8. **(Optional)** In Loopback Interface Configuration mode, you can configure a value for loopback:

ip bfd loopback <1-256>

Variable Definitions

The following table defines parameters for the ip bfd command.

Variable	Value	
{slot/port[/sub-port] [-slot/ port[/sub-port]] [,]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.	
enable	Enable BFD on a port, VLAN, or loopback.	
interval <100-65335>	Specifies the transmit interval in milliseconds. The default is 200 ms.	
	Note:	
	For XA1400 Series, the default is 1000 ms.	
	Note:	
	The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.	
min-rx <100-65535>	Specifies the receive interval in milliseconds. The default is 200 ms.	
	Note:	
	For XA1400 Series, the default is 1000 ms.	
	Note:	
	The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.	
multiplier <1-20>	Specifies the multiplier used to calculate the amount of time BFD waits before declaring a receive timeout. The default is 3.	
	Note:	
	If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.	

Variable	Value
port {slot/port[/sub-port] [- slot/port[/sub-port]] [,]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
vlan <1-4094>	Specifies the VLAN ID in the range of 1 to 4094.
loopback <1-256>	Specifies the Loopback ID in the range of 1 to 256.

Display BFD Configurations on the Loopback Interface

About this task

Use the following procedure to display all BFD configurations on the loopback interface.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Display the BFD configurations:

show ip bfd interfaces loopback

Example

The following example displays BFD configurations on the Loopback Interface:

Viewing TCP and UDP information

Use this procedure to view TCP and UDP configuration information for IPv4.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the IPv4 TCP connection information:

show ip tcp connections

3. View the IPv4 TCP connection information for a specific vrf or vrfids:

```
show ip tcp connections vrf WORD<0-16>
show ip tcp connections vrfids WORD<0-512>
```

4. View IPv4 TCP properties:

show ip tcp properties

5. View IPv4 TCP statistics

show ip tcp statistics

6. View IPv4 udp endpoints

show ip udp endpoints

7. View IPv4 udp statistics

show ip udp statistics

Example

Switch:1#sh	ow ip tcp connec	tions			
		TCP connection t	able info		
LOCALPORT	LOCALADDR	REMOTEPOR	T REMOTEADDR	STATE	VRF ID
21 22 23 80 443 23	0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 192.0.2.146	0 0 0 0 0 0 52583	0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 198.51.100.30	listen listen listen listen listen established	0 0 0 0 0

Switch: 1#show ip tcp properties

show ip tcp global properties command:

RtoAlgorithm constant
RtoMin 5002 milliseconds
RtoMax 60128 milliseconds

MaxConn 127

Switch: 1#show ip tcp statistics show ip tcp global statistics:

ActiveOpens: 0
PassiveOpens: 240
AttemptFails: 0
EstabResets: 239
CurrEstab: 1
InSegs: 5807
OutSegs: 6819
RetransSegs: 24
InErrs: 0
OutBets: 29 29 OutRsts:

Switch: 1#show ip udp endpoints

______ UDP endpoint table info ______

Variable definitions

HCInDatagrams:

HCOutDatagrams:

Use the data in the following table to use the show ip tcp command.

887

887

Variable	Value
connections	Specifies the TCP connection information.
	Use the following parameters:
	• vrf WORD<0–16>
	Specifies a virtual routing and forwarding (VRF) by name.
	• vrfids WORD<0–512>
	Specifies the IDs of a VRF path as an integer from 1 to 512.
	Example: show ip tcp connections vrf 0
properties	Specifies the TCP global properties information.
statistics	Specifies the TCP global statistics.

Use the data in the following table to use the show ip udp command.

Variable	Value
endpoints	Specifies the IP UDP endpoint information.
statistics	Specifies IP UDP statistics information.

IP routing configuration using Enterprise Device Manager

Configure the IP router interface so that you can use routing protocols and features on the interface. This section contains instructions for both the Global Router and Virtual Router Forwarding (VRF) instances.

Enabling routing for a router or a VRF instance

About this task

Enable IP forwarding (routing) on a router or a Virtual Router Forwarding (VRF) instance so that they support routing. You can use the IP address of any physical or virtual router interface for an IP-based network management.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Globals tab.
- 4. To enable routing, select Forwarding.
- 5. Click Apply.

Deleting a dynamically-learned route

About this task

Use the Routes tab to view and manage the contents of the system routing table. You can also delete a dynamically learned route using this table. Exercise caution if you delete entries from the route table.

To delete a static route, use the **StaticRoute** tab.

To delete dynamic routes from the table for a VRF instance, first select the appropriate instance.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Routes tab.
- 4. To delete a route, select the route and click **Delete**.

Routes field descriptions

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

Name	Description
Dest	Specifies the destination IP network of this route. An entry with a value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to multiple entries depends on the table access mechanisms defined by the network management protocol in use.

Name	Description
Mask	Indicates the network mask to logically add with the destination address before comparison to the destination IP network.
NextHop	Specifies the IP address of the next hop of this route.
AltSequence	Indicates the alternative route sequence. The value of 0 denotes the best route.
NextHopId	Specifies the identifier of the next-hop, hostname or MAC address.
HopOrMetric	Specifies the primary routing metric for this route. The semantics of this metric are specific to various routing protocols.
Interface	Specifies the router interface for this route.
	 Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation.
	 Brouter interfaces are identified by the slot and port number of the brouter port.
Proto	Specifies the routing mechanism through which this route was learned:
	 local—nonprotocol information, for example, manually configured entries
	• static
	• isis
	inter-vrf redistributed route
Age	Specifies the number of seconds since this route was last updated or otherwise determined correct.
PathType	Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths.
	iA indicates Indirect Alternative route without an ECMP path
	iAE indicates Indirect Alternative ECMP path
	iB indicates Indirect Best route without ECMP path
	iBE indicates Indirect Best ECMP path
	dB indicates Direct Best route
	iAN indicates Indirect Alternative route not in hardware
	iAEN indicates Indirect Alternative ECMP route not in hardware
	iBN indicates Indirect Best route not in hardware
	iBEN indicates Indirect Best ECMP route not in hardware
	dBN indicates Direct Best route not in hardware
	iAU indicates Indirect Alternative Route Unresolved
	iAEU indicates Indirect Alternative ECMP Unresolved

Name	Description
	iBU indicates Indirect Best Route Unresolved
	iBEU indicates Indirect Best ECMP Unresolved
	dBU indicates Direct Best Route Unresolved
	iBF indicates Indirect Best route replaced by FTN
	iBEF indicates Indirect Best ECMP route replaced by FTN
	iBV indicates Indirect best IPVPN route
	iBEV indicates Indirect best ECMP IP VPN route
	iBVN indicates Indirect best IP VPN route not in hardware
	iBEVN indicates Indirect best ECMP IP VPN route not in hardware
Pref	Specifies the preference.
NextHopVrfld	Specifies the VRF ID of the next-hop address.

Configuring IP route preferences

Before you begin

• Disable ECMP before you configure route preferences.

About this task

Change IP route preferences to force the routing protocols to prefer a route over another. Configure IP route preferences to override default route preferences and give preference to routes learned for a specific protocol.

Important:

Changing route preferences is a process-oriented operation that can affect system performance and network reachability while you perform the procedures. Therefore, it is recommended that if you want to change default preferences for routing protocols, you do so before you enable the protocols.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the RoutePref tab.
- 4. In the **ConfiguredValue** column, change the preference for the given protocol.
- 5. Click Apply.

RoutePref field descriptions

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

Name	Description
DefaultValue	Specifies the default preference value for the specified protocol.
Protocol	Specifies the protocol name.
ConfiguredValue	Configures the preference value for the specified protocol.

Flushing routing tables by VLAN

About this task

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use Enterprise Device Manager (EDM) to flush the routing tables by VLAN or by port. Use this procedure to flush the IP routing table for a VLAN.

To flush routing tables by VLAN for a VRF instance, first select the appropriate instance.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANS.
- 3. Click the Advanced tab.
- 4. In the **Vian Operation Action** column, select a flush option.

In a VLAN context, all entries associated with the VLAN are flushed. You can flush the ARP entries and IP routes for the VLAN.

5. Click Apply.

Flushing routing tables by port

For administrative and troubleshooting purposes, sometimes you must flush the routing tables. You can use EDM to flush the routing tables by VLAN or flush them by port. Use this procedure to flush the IP routing table for a port.

About this task

To flush routing tables by port for a VRF instance, first select the appropriate instance.

Procedure

- 1. In the Device Physical View, select a port.
- 2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
- 3. Click General.
- 4. Click the Interface tab.
- 5. In the Action section, select flushAll.

In a port context, all entries associated with the port are flushed. You can flush the ARP entries and IP routes for a port.

After you flush a routing table, it is not automatically repopulated. The repopulation time delay depends on the routing protocols in use.

6. Click Apply.

Assigning an IP address to a port

Assign an IP address to a port so that it acts as a routable VLAN (a brouter port) and supports IP routing.

To configure a brouter port, assign an IP address to an IP policy-based single-port VLAN.

Before you begin

- Ensure routing (forwarding) is globally enabled.
- Ensure the VLAN is configured.
- If required, ensure the VRF instance exists.

About this task

! Important:

After you configure the IP address, you cannot edit the IP address, and you can assign only one IP address to any router interface (brouter or virtual).

You cannot assign an IP address to a brouter port that is a member of a routed VLAN. To assign an IP address to the brouter port, you must first remove the port from the routed VLAN.

If you want to assign a new IP address to a VLAN or brouter port that already has an IP address, first delete the existing IP address and then insert the new IP address.

Procedure

- 1. In Device Physical View, select the port.
- 2. In the navigation tree, expand the **Configuration > Edit > Port** folders.
- 3. Click IP.
- 4. Click the IP Address tab.
- 5. Click Insert.
- 6. In the Insert IP Address dialog box, type the IP address, network mask, and VLAN ID.
- 7. Click Insert.

IP Address Field Descriptions

Use the data in the following table to help use the **IP Address** tab.

Name	Description
Interface	Specifies the router interface.
	The name of the VLAN followed by the VLAN designation identifies virtual router interfaces.
	The slot and port number of the brouter port identifies brouter interfaces.
Ip Address	Specifies the IP address of the brouter interface on this port. You can define only one IP address on a given port interface.
Net Mask	Specifies the subnet mask of the brouter interface on this port. The mask is an IP address with all the network bits set to 1 and all the hosts bits set to 0.
Name	Specifies the name of the brouter interface. The value ranges from 0 to 64 characters.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface.
ReasmMaxSize	Specifies the size of the largest IP packet which the interface can reassemble from fragmented incoming IP packets.
Vlanid	Specifies the ID of the VLAN associated with the brouter port. This parameter is used to tag ports.
BrouterPort	Indicates whether this is a brouter port.
MacOffset	Specifies a number by which to offset the MAC address of the VLAN from the chassis MAC address. This ensures that each IP address has a different MACaddress.
Vrfld	Specifies the associated VRF interface. The Vrfld associates VLANs or brouter ports to a VRF after the creation of VLANs or brouter ports. VRF ID 0 is reserved for the Global Router.

Assigning an IP address to a VLAN

Before you begin

- Ensure routing (forwarding) is globally enabled.
- Ensure VLAN is configured.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.

About this task

Specify an IP address for a VLAN so that the VLAN can perform IP routing.

Important:

You can assign only one IP address to any router interface (brouter or VLAN).

You cannot assign an IP address to a VLAN if a brouter port is a member of the VLAN. To assign an IP address to the VLAN, you must first remove the brouter port member.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANs > Basic.
- 3. Select a VLAN.
- 4. Click IP.
- 5. Click Insert.
- 6. In the **Insert IP Address** dialog box, type the IP address and network mask.
- 7. Click Insert.

Viewing IP addresses for all router interfaces

About this task

Use the Addresses tab to view IP addresses (and their associated router interfaces) from one central location.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Addresses tab.

Addresses field descriptions

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

Name	Description
Interface	Specifies the router interface.
	The name of the VLAN followed by the VLAN designation identifies virtual router interfaces.
	The slot and port number of the brouter port identifies brouter interfaces.
Ip Address	Specifies the IP address of the router interface.
Net Mask	Specifies the subnet mask of the router interface.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface; that is, whether 0 (zero) or one is used for the broadcast address. The switch uses 1.
ReasmMaxSize	Specifies the size of the largest IP packet that this interface can reassemble from incoming fragmented IP packets.
VlanId	Identifies the VLAN associated with this entry. This value corresponds to the lower 12 bits in the IEEE 802.1Q VLAN tag.

Name	Description
BrouterPort	Indicates whether this is a brouter port (as opposed to a routable VLAN).
MacOffset	Specifies a number by which to offset the MAC address of the VLAN from the chassis MAC address. This ensures that each IP address has a different MAC address.

Configuring IP routing features globally

About this task

Configure the IP routing protocol stack to determine which routing features the switch can use.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the **Globals** tab.
- 4. To globally enable routing, select **Forwarding**.
- 5. To globally configure the default TTL parameter type a value in the **DefaultTTL** field. This value is placed into routed packets that have no TTL specified.
- 6. To globally enable IPv4 ICMP broadcast, select IcmpEchoBroadcastRequestEnable.
- 7. To globally enable the Alternative Route feature, select **AlternativeEnable**.
- 8. To globally enable ICMP Router Discovery, select RouteDiscoveryEnable.
- 9. To globally enable IP Sorce Routing, select SourceRouteEnable.
- 10. To globally enable ECMP, select **EcmpEnable**.
- 11. Configure the remaining parameters as required.
- 12. Click Apply.

Globals Field Descriptions

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

Name	Description
Forwarding	Configures the system for forwarding (routing) or nonforwarding. The default value is forwarding.
DefaultTTL	Configures the default time-to-live (TTL) value for a routed packet. TTL indicates the maximum number of seconds elapsed before a

Name	Description
	packet is discarded. Enter an integer from 1 to 255. The default value of 255 is used if a value is not supplied in the datagram header.
ReasmTimeout	Specifies the maximum number of seconds that received fragments are held while they wait for reassembly. The default value is 30 seconds.
ICMPUnreachableMsgEnable	Enables the generation of Internet Control Message Protocol (ICMP) network unreachable messages if the destination network is not reachable from this system. These messages help determine if the system is reachable over the network. The default is disabled.
	Important:
	It is recommended that you enable icmp- unreach-msg only if it is absolutely required. If icmp-unreach-msg is enabled and a packet is received for which there is no route in the routing table, CPU utilization can dramatically increase.
ICMPRedirectMsgEnable	Enables or disables the system sending ICMP destination redirect messages.
IcmpEchoBroadcastRequestEnable	Enables or disables IP ICMP echo broadcast request feature. The default is enabled.
AlternativeEnable	Globally enables or disables the Alternative Route feature.
	If the alternative-route parameter is disabled, all existing alternative routes are removed. After the parameter is enabled, all alternative routes are re-added. The default is enabled.
RouteDiscoveryEnable	Enables the ICMP Router Discovery feature. The default is disabled (not selected). Use ICMP Router Discovery to enable hosts attached to multicast or broadcast networks to discover the IP addresses of neighboring routers.
AllowMoreSpecificNonLocalRouteEnable	Enables or disables a more-specific nonlocal route. If enabled, the system can enter a more-specific nonlocal route into the routing table. The default is disabled.
SuperNetEnable	Enables or disables supernetting.

Name	Description
	If supernetting is globally enabled, the system can learn routes with a route mask less than 8 bits. Routes with a mask length less than 8 bits cannot have ECMP paths, even if you globally enable the ECMP feature. The default is disabled.
UdpCheckSumEnable	Enables or disables the UDP checksum calculation. The default is enable.
SourceRouteEnable	Enables or disables IP Source Routing globally. It is disabled by default.
ARPLifeTime	Specifies the lifetime of an ARP entry within the system, global to the switch. The default value is 360 minutes.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled.
	After ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.
EcmpMaxPath	Globally configures the maximum number of ECMP paths.
	You cannot configure this feature unless ECMP is enabled globally.
	Different hardware platforms can support a different number of ECMP paths. For more information, see Release Notes for VOSS.
Ecmp1PathList	Selects a preconfigured ECMP path.
Ecmp2PathList	Selects a preconfigured ECMP path.
Ecmp3PathList	Selects a preconfigured ECMP path.
Ecmp4PathList	Selects a preconfigured ECMP path.
Ecmp5PathList	Selects a preconfigured ECMP path.
Ecmp6PathList	Selects a preconfigured ECMP path.
Ecmp7PathList	Selects a preconfigured ECMP path.
Ecmp8PathList	Selects a preconfigured ECMP path.
EcmpPathListApply	Applies changes in the ECMP pathlist configuration, or in the prefix lists configured as the pathlists.

Configure ECMP

Enable Equal Cost MultiPath (ECMP) to permit routers to determine up to eight equal-cost paths to the same destination prefix. You can use the multiple paths for load-sharing of traffic, which allows

fast convergence to alternative paths. By maximizing load sharing among equal-cost paths, you can maximize the efficiency of links between routers.

Before you begin

- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.
- To configure an ECMP pathlist, you must first configure a prefix list that you reference in the pathlist configuration.

About this task

Different hardware platforms can support a different number of ECMP paths. For more information, see Release Notes for VOSS.

Procedure

- 1. In the navigation pane, expand **Configuration > IP**.
- 2. Select IP.
- Select the Globals tab.
- 4. Select EcmpEnable.
- 5. **(Optional)** In **EcmpMaxPath**, type the preferred maximum number of equal-cost paths.
- 6. (Optional) Configure an ECMP pathlist to specify routes with the required number of paths.
- 7. (Optional) If you modified the ECMP pathlist configuration, select EcmpPathListApply.
- 8. Select Apply.

Enabling alternative routes globally

Before you begin

• Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see Selecting and launching a VRF context view on page 321.

About this task

Globally enable alternative routes so that you can subsequently enable it on interfaces.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Globals tab.
- 4. Select Alternative Enable.

If the **AlternativeEnable** parameter is disabled, all existing alternative routes are removed. After you enable the parameter, all alternative routes are re-added.

5. Click Apply.

Configure Static Routes using EDM

About this task

Use static routes to force the router to make certain forwarding decisions. Create IP static routes to manually provide a path to destination IP address prefixes.

Note:

It is recommended that you do not configure static routes on a DvR Leaf node unless the configuration is for reachability to a management network using a Brouter port.

Also, configuring the preference of static routes is not supported on a Leaf node.

For route scaling information, see Release Notes for VOSS.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the **Static Routes** tab.
- 4. Click Insert.
- 5. If required, in the **OwnerVrfld** check box, select the appropriate VRF ID. By default, the VRF is the GlobalRouter VRF 0.
- 6. In the **Dest** field, type the IP address.
- 7. In the **Mask** field, type the subnet mask.
- 8. In the **NextHop** field, type the IP address of the router through which the specified route is accessible.
- 9. (Optional) In the NextHopVrfld field, select the appropriate value.
- 10. (Optional) To enable the static route, select the **Enable** check box.
- 11. (Optional) In the Metric field, type the metric.
- 12. **(Optional)** In the **Preference** field, type the route preference.
- 13. (Optional) If required, select the LocalNextHop check box.

Use this option to create Layer 3 static routes.

14. Click Insert.

The new route appears in the **IP** dialog box, **Static Routes** tab.

Static Routes Field Descriptions

Use the data in the following table to use the **Static Routes** tab.

Name	Description
OwnerVrfld	Specifies the VRF ID for the static route.
Dest	Specifies the destination IP address of this route. A value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to such multiple entries is dependent on the table-access mechanisms defined by the network management protocol in use.
Mask	Indicates the mask that the system operates a logically AND function on, with the destination address, to compare the result to the Route Destination. For systems that do not support arbitrary subnet masks, an agent constructs the Route Mask by determining whether it belongs to a class A, B, or C network, and then uses one of:
	255.0.0.0—Class A
	255.255.0.0—Class B
	255.255.255.0—Class C
	If the Route Destination is 0.0.0.0 (a default route) then the mask value is also 0.0.0.0.
NextHop	Specifies the IP address of the next hop of this route. In the case of a route bound to an interface which is realized through a broadcast media, the Next Hop is the IP address of the agent on that interface.
	When you create a black hole static route, configure this parameter to 255.255.255.255.
NextHopVrfld	Specifies the next-hop VRF ID in interVRF static route configurations. Identifies the VRF in which the ARP entry resides.
Name	Specifies the name for the static route.
Note:	
This field does not apply to all hardware platforms.	
Enable	Determines whether the static route is available on the port. The default is enable.
	If a static route is disabled, it must be enabled before it can be added to the system routing table.
Status	Specifies the status of the route. The default is enabled.
Metric	Specifies the primary routing metric for this route. The semantics of this metric are determined by the routing protocol specified in the route RouteProto value. If this metric is not used, configure the value to 1. The default is 1.
IfIndex	Specifies the route index of the Next Hop. The interface index identifies the local interface through which the next hop of this route is reached.

Name	Description
Preference	Specifies the routing preference of the destination IP address. If more than one route can be used to forward IP traffic, the route that has the highest preference is used. The higher the number, the higher the preference.
LocalNextHop	Enables and disables LocalNextHop. If enabled, the static route becomes active only if the system has a local route to the network. If disabled, the static route becomes active if the system has a local route or a dynamic route.

Deleting a static route

About this task

Delete static routes that are no longer needed to prevent routing errors.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- Click the Static Routes tab.
- 4. Select the static route you want to delete.
- Click **Delete**.

Configuring a default static route

Before you begin

• Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.

About this task

The default route specifies a route to all networks for which there no explicit routes exist in the Forwarding Information Base or in the routing table. This route has a prefix length of zero (RFC 1812). You can configure the switch with a static default route, or they can learn it through a dynamic routing protocol.

To create a default static route, you configure the destination address and subnet mask to 0.0.0.0.

Note:

It is recommended that you do not configure static routes on a DvR Leaf node unless the configuration is for reachability to a management network using a Brouter port.

Also, configuring the preference of static routes is not supported on a Leaf node.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Static Routes tab.
- 4. Click Insert.
- 5. In the **OwnerVrfld** check box, select the appropriate VRF ID.
- 6. In the **Dest** field, type 0.0.0.0.
- 7. In the Mask field, type 0.0.0.0.
- 8. In the **NextHop** field, type the IP address of the router through which the specified route is accessible.
- 9. In the **Metric** field, type the HopOrMetric value.
- 10. Click Insert.

Configuring a black hole static route

Before you begin

 Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.

About this task

Create a black hole static route to the destination that a router advertises to avoid routing loops when aggregating or injecting routes to other routers.

If an existing black hole route is enabled, you must first delete or disable it before you can add a regular static route to that destination.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
- 2. Click IP.
- 3. Click the Static Routes tab.
- 4. Click Insert.
- 5. In the **OwnerVrfId** check box, select the appropriate VRF ID.
- 6. In the **Dest** field, enter the IP address.
- 7. In the **Mask** field, enter the network mask.
- 8. In the **NextHop** field, type 255.255.255.255.

To create a black hole static route, you must configure the NextHop address to 255.255.255.255.

- 9. Select the **enable** option.
- 10. In the **Metric** box, type the HopOrMetric value.
- 11. In the **Preference** check box, select the route preference.

When you specify a route preference, be sure to appropriately configure the preference so that when the black hole route is used, it is elected as the best route.

12. Click Insert.

Viewing IP routes

View IP routes learned on the device.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the **Routes** tab to view IP routes learned on the device.
- 4. If you want to limit the routes displayed, click **Filter** to show a smaller subset of the learned routes.
- 5. In the Filter dialog box, select an option, or options, and enter information to limit the routes to display in the Routes table.
- 6. Click **Filter** and the Routes table displays only the routes that match the options and information that you enter.

Routes field descriptions

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

Name	Description
Dest	Specifies the destination IP network of this route. An entry with a value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to multiple entries depends on the table access mechanisms defined by the network management protocol in use.
Mask	Indicates the network mask to logically add with the destination address before comparison to the destination IP network.
NextHop	Specifies the IP address of the next hop of this route.
AltSequence	Indicates the alternative route sequence. The value of 0 denotes the best route.
NextHopId	Displays the MAC address or hostname of the next hop.
HopOrMetric	Displays the primary routing metric for this route. The semantics of this metric are specific to different routing protocols.

Name	Description
Interface	Specifies the router interface for this route.
	Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation.
	Brouter interfaces are identified by the slot and port number of the brouter port.
Proto	Specifies the routing mechanism through which this route was learned:
	other—none of the following
	local—nonprotocol information, for example, manually configured entries
	• static
	• ICMP
	• EGP
	• GGP
	• Hello
	• RIP
	• IS-IS
	• ES-IS
	Cisco IGRP
	bbnSpflgp
	• OSPF
	• BGP
	Inter-VRF Redistributed Route
Age	Displays the number of seconds since this route was last updated or otherwise determined to be correct.
PathType	Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths.
	iA indicates Indirect Alternative route without an ECMP path
	iAE indicates Indirect Alternative ECMP path
	iB indicates Indirect Best route without ECMP path
	iBE indicates Indirect Best ECMP path
	dB indicates Direct Best route
	iAN indicates Indirect Alternative route not in hardware
	iAEN indicates Indirect Alternative ECMP route not in hardware
	iBN indicates Indirect Best route not in hardware

Name	Description
	iBEN indicates Indirect Best ECMP route not in hardware
	dBN indicates Direct Best route not in hardware
	iAU indicates Indirect Alternative Route Unresolved
	iAEU indicates Indirect Alternative ECMP Unresolved
	iBU indicates Indirect Best Route Unresolved
	iBEU indicates Indirect Best ECMP Unresolved
	dBU indicates Direct Best Route Unresolved
	iBF indicates Indirect Best route replaced by FTN
	iBEF indicates Indirect Best ECMP route replaced by FTN
	iBV indicates Indirect best IPVPN route
	iBEV indicates Indirect best ECMP IP VPN route
	iBVN indicates Indirect best IP VPN route not in hardware
	iBEVN indicates Indirect best ECMP IP VPN route not in hardware
Pref	Displays the preference.
NextHopVrfld	Specifies the VRF ID of the next-hop address.

Configuring ICMP Router Discovery globally

About this task

Enable ICMP Router Discovery so that it can operate on the system.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Globals tab.
- 4. Select RouteDiscoveryEnable.
- 5. To select a preconfigured ECMP path, click the **EcmpPathList** ellipsis button.
- 6. Click OK.
- 7. Click Apply.

Configuring the ICMP Router Discovery table

Before you begin

• ICMP Router Discovery must be globally enabled.

 Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.

About this task

Configure the ICMP Router Discovery table to ensure correct ICMP operation for all interfaces that use Router Discovery.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Router Discovery tab.
- 4. Configure the Router Discovery parameters to suit your network.
- 5. Click Apply.

Router Discovery field descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
Interface	Indicates the VLAN ID or the port.
AdvAddress	Specifies the IP destination address used for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address 224.0.0.1, or the limited-broadcast address 255.255.255.
	The default value is 255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface.
	The default value is true (advertise address).
AdvLifetime	Specifies the time to-live-value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds.
	The default value is 1800 seconds.
MaxAdvInterval	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 4 to 1800 seconds.
	The default value is 600 seconds.
MinAdvInterfal	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to MaxAdvInterval.
	The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address relative to other

Name	Description
	router addresses on the same subnet. The range is –2147483648 to 2147483647.
	The default value is 0.

Configuring ICMP Router Discovery for a port

Before you begin

- · You must globally enable ICMP Router Discovery.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.

About this task

Use this procedure to configure Router Discovery on a port. When enabled, the port sends Router Discovery advertisement packets.

Procedure

- 1. In the Device Physical View tab, select a port.
- 2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
- 3. Click IP.
- 4. Click the Router Discovery tab.
- 5. To enable Router Discovery, select **AdvFlag**.
- 6. Configure other parameters as required for proper operation.
- 7. Click Apply.

Router Discovery field descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
AdvAddress	Specifies the destination IP address used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address 224.0.0.1, or the limited-broadcast address 255.255.255.
	The default value is 255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface.
	The default value is True (advertise address).

Name	Description
AdvLifetime	Specifies the time to live value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds.
	The default value is 1800 seconds.
MaxAdvinterval	Specifies the maximum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 4 seconds to 1800 seconds.
	The default value is 600 seconds.
MinAdvInterval	Specifies the minimum time (in seconds) that elapses between unsolicited broadcast or multicast router advertisement transmissions from the interface. The range is 3 seconds to MaxAdvInterval.
	The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address relative to other router addresses on the same subnet. The accepted values are – 2147483648 to 2147483647.
	The default value is 0.

Configuring ICMP Router Discovery on a VLAN

Before you begin

- · You must globally enable ICMP Router Discovery.
- Change the VRF instance as required. For information about how to use EDM for a non-default VRF, see <u>Selecting and launching a VRF context view</u> on page 321.

About this task

Configure Router Discovery on a VLAN so that the ICMP Router Discovery feature can run over the VLAN. When enabled, the system sends Router Discovery advertisement packets to the VLAN.

Procedure

- 1. In the navigation tree, expand the following folders: Configuration > VLAN.
- 2. Click VLANs.
- 3. Select the VLAN ID that you want to configure to participate in Router Discovery.
- 4. Click IP.
- 5. Click the Router Discovery tab.
- 6. To enable Router Discovery for the VLAN, select AdvFlag.
- 7. Configure other parameters as required for proper operation.
- 8. Click Apply.

Router Discovery field descriptions

Use the data in the following table to use the **Router Discovery** tab.

Name	Description
AdvAddress	Specifies the destination IP address used for broadcast or multicast router advertisements sent from the interface. The address is the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255.
	The default value is 255.255.255.
AdvFlag	Indicates whether (true) or not (false) the address is advertised on the interface.
	The default value is true (advertise address).
AdvLifetime	Specifies the time to-live-value (TTL) of router advertisements (in seconds) sent from the interface. The range is MaxAdvInterval to 9000 seconds.
	The default value is 1800 seconds.
MaxAdvinterval	Specifies the maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The range is 4 seconds to 1800 seconds.
	The default value is 600 seconds.
MinAdvInterval	The minimum time (in seconds) allowed between unsolicited broadcast or multicast router advertisements sent from the interface. The range is 3 seconds to MaxAdvInterval.
	The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The range is –2147483648 to 2147483647.
	The default value is 0.

Configure a Circuitless IPv4 Interface

About this task

You can use a circuitless IPv4 (CLIPv4) interface to provide uninterrupted connectivity to your system.

Procedure

- 1. In the navigation pane, expand **Configuration > IP**.
- 2. Select IP.
- 3. Select the Circuitless IP tab.
- 4. Select Insert.

- 5. In the **Interface** field, assign a CLIP interface number.
- 6. Enter the IP address.
- 7. Enter the network mask.
- 8. Select Insert.
- 9. To delete a CLIP interface, select the interface and select **Delete**.

Circuitless IP Field Descriptions

Use the data in the following table to use the Circuitless IP tab.

Name	Description
Interface	Specifies the number assigned to the interface.
Ip Address	Specifies the IP address of the CLIP.
Net Mask	Specifies the network mask.
Name	Specifies the name assigned to the IPv4 CLIP address.

Enabling OSPF on a CLIP interface

Before you begin

- · You must globally enable OSPF.
- The OSPF area must already exist.

About this task

Enable Open Shortest Path First (OSPF) on a CLIP interface so that it can participate in OSPF routing.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Circuitless IP tab.
- 4. Select the required CLIP interface.
- 5. Click OSPF.
- 6. Select the **Enable** check box.

You must enable OSPF on the CLIP interface for CLIP to function.

- 7. In the current **Areald** field, enter the IP address of the OSPF backbone area.
- 8. Click Apply.

Circuitless OSPF field descriptions

Use the data in the following table to use the **Circuitless OSPF** tab.

Name	Description
Enable	Enables OSPF on the CLIP interface.
Areald	Specifies the OSPF area ID.

Enabling PIM on a CLIP interface

Enable Protocol Independent Multicasting (PIM) on a CLIP interface so that it can participate in PIM routing.

Before you begin

You must globally enable PIM.

About this task



Note:

The PIM button does not appear for all hardware platforms.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the Circuitless IP tab.
- 4. Select the required CLIP interface.
- 5. Click PIM.
- 6. Select the **Enable** check box.

You must enable PIM on the CLIP interface for PIM to function. The mode is indicated on this tab.

7. Click Apply.

Circuitless PIM field descriptions

Use the data in the following table to use the Circuitless PIM tab.

Name	Description
Enable	Enables PIM on the CLIP interface.
Mode	Specifies the PIM mode.

Enable BFD on a CLIP interface

Before you begin

• The CLIP Interface must already exist.

About this task

Enable Bidirectional Forwarding Detection (BFD) over Fabric Extend tunnels on a CLIP interface.

Note:

The BFD button does not appear for all hardware platforms.

Procedure

- 1. In the navigation pane, expand the **Configuration > IP** folders.
- 2. Click IP.
- 3. Click the Circuitless IP tab.
- 4. Select the required CLIP interface.
- 5. Click BFD.
- 6. Select the **Enable** check box.
- 7. **(Optional)** In the **MinRxInterval** field, specify the minimum receive interval.
- 8. **(Optional)** In the **TxInterval** field, specify the transmit interval.
- 9. (Optional) In the Multiplier field, specify the multiplier used to calculate a receive timeout.

BFD Field Descriptions

Use the data in the following table to use the BFD tab.

Name	Description	
Enable	Enable BFD on the CLIP interface.	
MinRxInterval	Specifies the minimum interval, in milliseconds, between received BFD control packets that the local system is capable of supporting. The default is 200 ms.	
	Note: For XA1400 Series, the default is 1000 ms.	
	Note:	
	The minimum value you can configure for the receive interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the receive interval. You can configure any remaining BFD sessions with a receive interval that is greater than or equal to the 200 ms default value.	
TxInterval	Specifies the transmit interval in milliseconds. The default is 200 ms.	

Name	Description	
	Note: For XA1400 Series, the default is 1000 ms.	
	Note:	
	The minimum value you can configure for the transmit interval is 100 ms. You can configure a maximum of 4 BFD sessions with the minimum value for the transmit interval. You can configure any remaining BFD sessions with a transmit interval that is greater than or equal to the 200 ms default value.	
Multiplier	Specifies a value for the multiplier used to calculate a receive timeout. The default is 3.	
	Note:	
	If you configure the transmit interval or the receive interval as 100 ms, you must configure a value of 4 or greater for the multiplier.	

Viewing TCP global information

View TCP and UDP information to view the current configuration.

About this task

The fields on the TCP global tab provide information about the handshake (SYN) configuration and the maximum number of TCP connections you can create on your system.

When you initiate a TCP connection, both end points send handshake information to create the channel.

The retransmission algorithm and fields display the configured timeout value and minimum and maximum retransmission times that your system uses to terminate a connection attempt that falls outside your specified parameters.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **IP** folders.
- 2. Click TCP/UDP.
- 3. Click the TCP Globals tab.

TCP Global field descriptions

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

Name	Description
RtoAlgorithm	Determines the timeout value used for retransmitting unacknowledged octets.
RtoMin	Displays the minimum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
RtoMax	Displays the maximum time (in milliseconds) permitted by a TCP implementation for the retransmission timeout.
MaxConn	Displays the maximum connections for the device.

Viewing TCP connections information

View information about TCP connections.

About this task

Among other things, the fields on the TCP connections tab provide important information about the health of connections that traverse your switch.

In particular, the state column lets you know the state of each TCP connection. Of these, synSent, synReceived, and established indicate whether or not a channel is established and listen indicates when an end system is waiting for a returning handshake (SYN).

Procedure

- 1. In the navigation pane, expand the : **Configuration > IP** folders.
- 2. Click TCP/UDP.
- 3. Click the TCP Connections tab.

TCP Connections field descriptions

Use the data in the following table to use the **TCP Connections** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
RemAddressType	Displays the type (IPv6 or IPv4) for the remote address of the TCP connection.
RemAddress	Displays the IPv6 address for the remote TCP connection.

Name	Description	
RemPort	Displays the remote port number for the TCP connection.	
State	Displays an integer that represents the state for the connection:	
	• closed	
	listen	
	• synSent	
	synReceived	
	established	
	• finWait1	
	• finWait2	
	closeWait	
	lastAck(9)	
	• closing	
	timeWait	
	deleteTCB	
Process	Displays the process ID for the system process associated with the TCP connection.	

Viewing TCP listeners information

View TCP listener information.

About this task

The TCP listeners table provides a detailed list of systems that are in the listening state.

When a connection is in the listen state an end point system is waiting for a returning handshake (SYN). The normal listening state should be very transient, changing all of the time.

Two or more systems going to a common system in an extended listening state indicates the need for further investigation.

End systems in an extended listening state can indicate a broken TCP connection or a DOS attack on a resource.

This type of DOS attack, known as a SYN attack, results from the transmission of SYNs with no response to return replies.

While many systems can detect a SYN attack, the TCP listener statistics can provide additional forensic information.

Procedure

1. In the navigation pane, expand the **Configuration** > **IP** folders.

- 2. ClickTCP/UDP.
- 3. Click the **TCP Listeners** tab.

TCP Listeners field descriptions

Use the data in the following table to use the **TCP Listeners** tab.

Name	Description
LocalAddressType	Displays the type (IPv6 or IPv4) for the address in the LocalAddress field.
LocalAddress	Displays the IPv6 address for the TCP connection.
LocalPort	Displays the local port number for the TCP connection.
Process	Displays the process ID for the system process associated with the TCP connection.

Chapter 4: Distributed Virtual Routing

Table 22: Distributed Virtual Routing product support

Feature	Product	Release introduced	
For configuration details, see Configuration	For configuration details, see Configuring IPv4 Routing for VOSS.		
Distributed Virtual Routing (DvR) Controller	VSP 4450 Series	Not Supported	
	VSP 4900 Series	VOSS 8.1.5	
Important: Because of a change in VOSS 6.0.1.2, the best		VSP4900-12MXU-12XE, VSP4900-24S, and VSP4900-24XE only	
practice is to use a minimum	VSP 7200 Series	VOSS 6.0.1	
software version of 6.0.1.2 in DvR deployments.	VSP 7400 Series	VOSS 8.0	
BVIX deployments.	VSP 8200 Series	VOSS 6.0.1	
	VSP 8400 Series	VOSS 6.0.1	
	VSP 8600 Series	VSP 8600 8.0	
	XA1400 Series	Not Supported	
Distributed Virtual Routing (DvR)	VSP 4450 Series	VOSS 6.1	
Leaf	VSP 4900 Series	VOSS 8.1	
Important:	VSP 7200 Series	VOSS 6.0.1	
Because of a change in	VSP 7400 Series	VOSS 8.0	
VOSS 6.0.1.2, the best practice is to use a <i>minimum</i>	VSP 8200 Series	VOSS 6.0.1	
software version of 6.0.1.2 in	VSP 8400 Series	VOSS 6.0.1	
DvR deployments.	VSP 8600 Series	Not supported	
	XA1400 Series	Not Supported	
Distributed Virtual Routing (DvR) dvr-one-ip	VSP 4450 Series	Not Supported	
	VSP 4900 Series	VOSS 8.2 demonstration feature	
	VSP 7200 Series	VOSS 8.2 demonstration feature	
	VSP 7400 Series	VOSS 8.2 demonstration feature	
	VSP 8200 Series	VOSS 8.2 demonstration feature	
	VSP 8400 Series	VOSS 8.2 demonstration feature	
	VSP 8600 Series	VSP 8600 8.0	

Feature	Product	Release introduced
	XA1400 Series	Not Supported
Distributed Virtual Routing (DvR) In-band Management	VSP 4450 Series	VOSS 6.1
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 6.0.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 6.0.1
	VSP 8400 Series	VOSS 6.0.1
	VSP 8600 Series	Not supported
	XA1400 Series	Not Supported

Distributed Virtual Routing (DvR) Fundamentals

Distributed Virtual Routing (DvR) is a technology for router redundancy in a fabric deployment where IP subnets are stretched across multiple switches. DvR provides Default Gateway Redundancy and optimizes traffic flows to avoid traffic tromboning due to inefficient routing, thereby increasing the total routing throughput.

DvR can be deployed in Campus environments for stretching IP subnets between multiple aggregation layer switches and also simplifies data center deployments by introducing a *Controller-Leaf* architecture. In this architecture, Layer 3 configuration is required only on the Controller nodes, whereas the Leaf nodes require only Layer 2 configuration. All Layer 3 configuration is automatically distributed to the Leaf nodes by the Controller nodes.

For typical Campus DvR deployments, configure aggregation layer switches as DvR Controllers. Wiring closet access switches are then typically dual-homed to a pair of DvR Controllers.

IP subnets, which stretch between aggregation layer switches and multiple wiring closets, enable seamless IP roaming for wireless users while at the same time ensure optimal traffic forwarding. To optimize automation, Fabric Attach is typically deployed between wiring closet and aggregation switches. In this construct, there would likely be no DvR Leaf configured.

In fabric deployments, DvR replaces VRRP (with VRRP-BackupMaster or RSMLT). The operator can chose for each I-SID/IP subnet what router redundancy method to use.

To migrate to a DvR-enabled I-SID/IP subnet, all member fabric switches of this I-SID must be either DvR Controllers or DvR Leafs. You can connect non fabric switches to DvR Leafs and DvR Controllers with manual configuration or Fabric Attach configuration. Until all fabric switches that are members of the I-SID/IP subnet are DvR-enabled, use VRRP or RSMLT as the router redundancy protocol.

DvR Domain

To enable multi-site DvR deployments, a DvR domain concept has been introduced. Within a DvR domain, a set of up to eight DvR Controllers control the DvR domain Leaf switches. A domain can also include just DvR controllers without DvR Leafs. Typically, a DvR domain is restricted to one physical location. Traffic leaving this physical location always passes through DVR Controllers.

A DvR domain is a logical group of switches or nodes that are DvR enabled. These nodes are not physically connected but are connected over the SPB Fabric such that each node is aware of the BMAC addresses of all other nodes within the domain. A DvR domain does not contain nodes that are not DvR enabled. However, those nodes can co-exist with other DvR enabled nodes within the same SPB Fabric network.

You configure a common DvR domain ID for all nodes belonging to a DvR domain. This domain ID translates internally to a Domain Data Distribution (DDD) I-SID. All switch nodes that share the same DvR domain ID or DDD ISID receive the Layer 3 information that is distributed from all other nodes belonging to that DvR domain.

A DvR domain can contain multiple Layer 3 VSNs and Layer 2 VSNs. Layer 2 and Layer 3 VSNs can span multiple DvR domains.

A DvR domain typically has the following members:

- 1. DvR Controller(s)
- 2. DvR Leaf nodes

For scaling information on the number of Controllers and Leaf nodes to configure in a DvR domain, see Release Notes for VOSS.

DvR Controller

In a DvR domain, the Controller nodes are the central nodes on which Layer 3 is configured. They own all the Layer 3 configuration and push the configuration information to the Leaf nodes within the SPB network.

A DvR domain can have one or more controllers for redundancy and you must configure every Layer 2 VSN (VLAN) and Layer 3 VSN within the domain, on the Controller(s). A node that you configure as a DvR Controller is considered the controller for all the Layer 2 and Layer 3 VSNs configured on that node. A Controller is configured with its own subnet IP address for every DvR enabled Layer 2 VSN within the domain.

All Layer 2 VSNs on a DvR Controller need not be DvR enabled. A controller can be configured with individual Layer 2 VSNs that are DvR disabled.

The Layer 3 configuration data that is pushed to the Leaf nodes include the Layer 3 IP subnet information for all Layer 2 VSNs within the DvR domain. It also includes the IP routes learned or redistributed by the Controllers from networks outside the SPB network, into the DvR Domain. Controllers also send information on whether Multicast is enabled on a specific DvR enabled Layer 2 VSN, and the version of IGMP. DvR Controllers inject a default route into the DvR domain for

external route reachability. Use route policies to inject specific routes into a DvR domain or inject host routes into OSPF or BGP.

A Controller can only belong to one DvR domain, based on the domain ID that you configure on the node.

DvR Controllers include all DvR Leaf functions, thus a Leaf node free deployment is a valid network deployment. Especially if you use DvR in Campus deployments to replace VRRP or RSMLT, a Controller-only deployment, as Fabric Attach server nodes, is a valid deployment option.

DvR Leaf Node

DvR Leaf nodes are typically data center top of the rack (TOR) Fabric switches that aggregate physical and virtual servers or storage devices. DvR Leaf nodes operate in a reduced configuration mode, where Layer 3 is not configured locally, but pushed to them from the DvR Controller(s) within the domain. You need to configure only the IS-IS infrastructure and the Layer 2 VSNs on the Leaf nodes.

A DvR Leaf node also monitors local host attachments and communicates updates about the current state of those host attachments to the DvR domain. All DvR nodes exchange host attachment information using the DvR host distribution protocol, which leverages a DvR domain I-SID.

DvR leaf nodes are managed in-band through a local loopback address, which is exchanged using the IP Shortcut protocol.

Eligibility Criteria for a Leaf Node

A Leaf node must support the following criteria:

- configuration of basic parameters of IP Multicast over Fabric Connect, such as the system ID, nickname, B-VLANs, SPBM instance, area, peer system ID and virtual BMAC
- configuration of a physical port as either an SPB NNI interface, a FLEX-UNI interface or an FA interface
- configuration of an MLT as either an SPB NNI interface, a FLEX-UNI interface or an FA interface
- configuration of an SMLT as a Flex-UNI Interface or an FA Interface
- configuration of Layer 2 VSN I-SID instances of type ELAN
- configuration of FLEX-UNI end-points as part of a Layer 2 VSN
- FA Server functionality on FA enabled interfaces
- SMLT and vIST
- configuration of a in-band management interface for in-band management of the node

Configuration Limitations

The DvR Controller and the DvR Leaf node have the following configuration limitations:

- On a Controller, for any VLAN, you can either configure DvR or VRRP but not both. Similarly, you can either configure DvR or RSMLT but not both.
- On a leaf node, you cannot configure the following:
 - Layer 3 (for example, IP interfaces, IP routing and VRFs). You can only perform Layer 2 configuration.
 - platform VLANs. You cannot configure a platform VLAN directly on a DvR leaf node. However, you can configure a VLAN Management Instance on a DvR leaf node. After you configure the management VLAN, you can configure a platform VLAN.
- On a Controller, you cannot configure the following:
 - dynamic routing protocols like OSPF, RIP and BGP on a DVR enabled VLAN.
 - static routes such that the next-hop exists on any DVR enabled VLAN.

Summary of Controller and Leaf Node Functions

A DvR Controller performs the following functions:

- pushes Layer 3 configuration data (IPv4 Unicast and Multicast) to the Leaf nodes for all the Layer 2 VSNs or subnets within the DvR domain.
- pushes the Layer 3 learned host routes (host routes learned on its own UNI ports) and route data learned through route redistribution or route policies, to the Leaf nodes.
- configures learned remote host routes from other Controllers and Leaf nodes, on its own device.

A DvR enabled Leaf node performs the following functions:

- configures the gateway MAC when the gateway IPv4 address is learned.
- pushes the Layer 3 learned remote host routes to other Controllers and Leaf nodes in the domain.
- configures learned remote host routes from other Controllers and Leaf nodes on its own device.
- configures ECMP routes (in the datapath only) for the Layer 2 VSN subnets, with each next hop as the Controller in the DvR domain.
- configures learned routes from the Controllers that are redistributed using DvR.
- handles host route response packet interception based on the Controller VLAN MAC or the gateway MAC.

DvR backbone

The DvR backbone is automatically established among the DvR Controllers from all DvR domains. Every Controller node has an edge gateway to its DvR domain, to the DvR backbone and all other non-DvR domains within the network.

Controllers exchange host route information such that any host can be reached in a shortcut switched manner, irrespective of its location. For these host route information exchanges, controllers use an automatically assigned backbone I-SID. Local subnets to the Controllers are automatically injected into the DvR host route exchanges.

To redistribute DvR host routes into OSPF or BGP, you can configure route policies. These host routes are not injected into IS-IS.

DvR Backbone Members

You can configure a non-DvR backbone edge bridge (BEB) to join the DvR backbone. This enables the node to receive redistributed DvR host routes from all DvR Controllers in the SPB network, just like a DvR Controller. However, unlike the Controller, you can neither configure a DvR interface on this node nor can the node inject its host routes into the DvR domain.

DvR operation

In a DvR domain, DvR enabled Controller(s) handle the learning and distribution of Layer 3 configuration and route data to the DvR enabled Leaf nodes. The Leaf nodes in turn, use this data to automatically create distributed Layer 3 datapaths on themselves. In this way, Layer 3 configuration and learning remains only with the Controller(s) and there are distributed Layer 3 datapaths at the edges of the fabric network. This allows for destination lookups at the edge to happen quickly, and traffic is sent directly to their destinations without multiple lookups.

An important benefit of DvR is that only minimal configuration is required on the Leaf node. Based on the Layer 2 VSN that the Leaf node is a part of, all Layer 3 configuration information (IPv4 Unicast and Multicast configuration) is pushed from the Controllers in the domain. Thus the leaf nodes, although basically Layer 2 configured switches, become fully layer 3 capable devices.

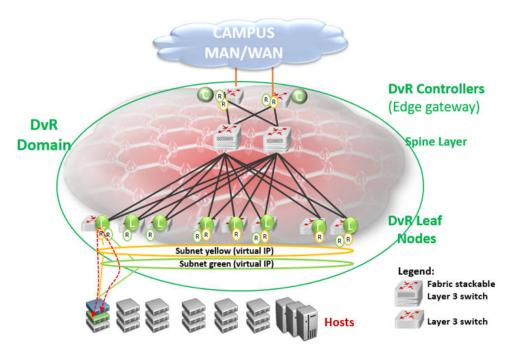


Figure 5: SPB Fabric network with central Layer 3 Controller and distributed Layer 3 datapath at the edges

ARP Learning

When DvR is enabled on a Controller, it initiates ARP requests for traffic to be routed to unknown destination hosts.

DvR enabled Controllers learn ARP requests from:

- DvR enabled Leaf nodes (here the Leaf node owns the ARPs)
- its own local UNI ports. Here, the controller owns the ARPs
- other DvR enabled Controllers

DvR enabled Leaf nodes learn ARP requests from:

- its own local UNI ports (here the Leaf node owns the ARPs)
- other DvR enabled Leaf nodes (that own the ARPs) and respond to ARP requests on their UNI ports
- DvR enabled Controllers (that own the local UNI ARPs)

Controllers only distribute ARP entries that are locally learned on its own UNI ports, to other DvR enabled nodes in the domain.

dvr-leaf-mode boot flag

To configure a node to operate as a DvR Leaf node, you must first enable the dvr-leaf-mode boot flag.

- The dvr-leaf-mode boot flag is disabled by default. You must explicitly enable this flag before you configure a switch node to operate as a Leaf node.
- After you enable or disable the boot flag, you must save the configuration and reboot with the saved configuration, for the changes to take effect.

Important:

A node on which the dvr-leaf-mode boot flag is enabled cannot be configured as a DvR Controller.

In-band management

Use in-band management to manage a DvR enabled Leaf node that does not have an out-of-band management port or a console port.

For in-band management of the node within the management subnet (for example, from a Controller node), you must configure a unique IPv4 address to be used as the in-band management IP address, on that node. This IPv4 address functions like a CLIP address.

DvR deployment scenarios

The following sections describe typical deployments of the DvR infrastructure.

DvR deployment in a single data center

The following topology shows DvR deployment in a single data center. This deployment consists of a single DvR domain comprising a Controller layer and a Leaf node layer. The Controller layer has two controllers (for redundancy), which are deployed closer to the boundary of the DvR domain and the rest of the SPB Fabric network. The DvR Leaf nodes or Top of Rack (TOR) switches are typically access or edge switches.

All switches that belong to the DvR domain are configured with the same DvR domain ID and communicate with each other over a predefined I-SID.

The Controller nodes control the Leaf nodes and also build the gateway between the DvR domain and the rest of the Fabric infrastructure. So traffic is either routed between the Leaf nodes, or through the Controllers, to the rest of the fabric infrastructure.

Two IP subnets (Layer 2 VSNs), yellow and green, span the Leaf nodes. Each subnet is configured with a virtual IP address that is a shared among all Controller and Leaf nodes that belong to the

DvR Controllers (Edge gateway)

Spine Layer

DvR Leaf
Nodes

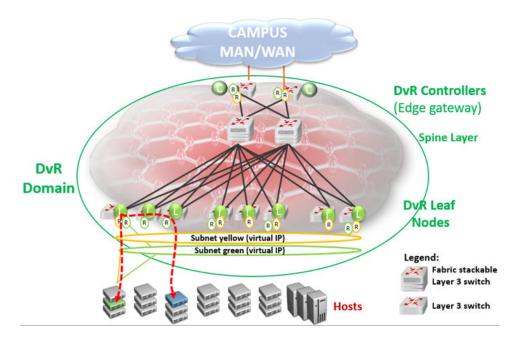
Subnet yellow (virtual IP)

Legend:
Fabric stackable
Layer 3 switch

subnet. The Controller and Leaf nodes are configured with routing interfaces to the subnets, as shown in the figure.

DvR works by enabling each Leaf node or Top of Rack (TOR) switch to bi-directionally route traffic for each IP subnet of which it is a member. This is done by distributing the Layer 3 configuration information (IP Unicast, IP Multicast and virtual IP configuration) needed to handle Layer 3 routing, from the Controllers to the Leaf nodes. Configuration information is pushed over the DvR Domain I-SID, as indicated by the blue arrows in the above figure.

Routing between the two IP subnets is achieved directly at the Leaf nodes when the Layer 3 distributed datapath is programmed at the Leaf Nodes, based on the Layer 3 configuration data that is pushed. Thus traffic within and between IP subnets is shortcut switched without having to traverse the central routing nodes, as shown in the figure below, if there are direct physical connections between them.



Thus, in a DvR deployment, all virtual IP and Layer 3 configuration is performed on the Controller nodes and pushed to the Leaf nodes, so that the Leaf nodes though basically Layer 2 configured switches, become fully layer 3 capable devices.

DvR deployment in a dual data center

The following example deployment shows two data centers each having its own DvR domain, connected through a backbone.

All nodes in data center Campus 1 belong to DvR domain shown in green, and the nodes in the data center Campus 3 belong to the DvR domain shown in orange. The two DvR domains are individually managed, so in this scenario, the controllers colored orange manage the orange Leaf nodes and the controllers colored green manage the green Leaf nodes. However, subnets can still be stretched across the DvR domains (and possibly between buildings), as shown in the figure.

Each DvR domain learns its own Layer 3 data and distributes this information to its own Leaf nodes. Layer 3 host information that is redistributed from other DvR Domains is learned by the Controllers only (through inter-DvR domain redistribution) and is programmed on the Leaf nodes in the same domain, but not in the other Domain. For example, Layer 3 information redistributed from domain 2 is learned by all controllers including the domain 1 controllers, but this information is not distributed to the Leaf nodes in domain 1.Hosts in one DvR domain can reach the hosts in the other DvR domain only through the Controllers.

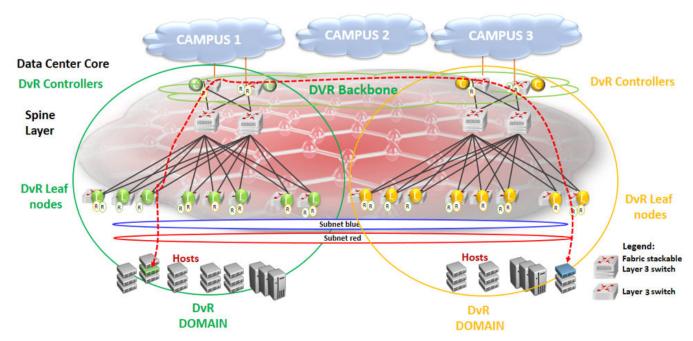


Figure 6: Shortest path routing between servers in different data centers

All controllers in all domains are always part of the DvR backbone by default, as they are connected by the SPB Fabric. The DvR backbone connects many DvR domains.

Thus DvR can scale to multiple campuses, allowing a simplified way to deploy a large scale fully-routed infrastructure.

DvR Route Redistribution

The following sections describe redistribution of IPv4 local and static routes from DvR Controllers into the DvR domain. It also describes redistribution of host routes that are learned on DvR enabled VLANs, to BGP and OSPF. You can configure route policies to control the selection of routes to be distributed. You can also configure IS-IS accept policies on DvR Controllers and non-DvR BEBs, to determine which DvR host routes to accept into the routing-table from the DvR backbone.

Redistribution of IPv4 Local and Static Routes

The DvR feature supports redistribution of IPv4 local and static routes into the DvR domain.



For every VRF instance and the Global Router, the Controller automatically injects a default route to the Leaf node, with a next hop as the advertising Controller. However, if you require only local or static routes to be advertised to the Leaf nodes, you can manually disable the injection of default routes on the Controller.

On a DvR Controller, you can configure (enable or disable) the redistribution of direct or static routes. Direct routes are redistributed with the route type as *internal*. Static routes are redistributed with the route type as *external*. You can apply route policies on the Controller to selectively permit

the redistribution of these routes and also configure a metric value for the route that is redistributed. The default metric for imported local routes is 1. For static routes, the configured route metric or cost is honored.

You can configure redistribution of static and direct routes from the Global Router, or within a VRF instance. For redistributed routes, the Controller configures the Layer 3 VSN as that of the VRF redistributing the route, and the next hop BEB as the system ID of the Controller injecting the route into the DvR domain.

The following example demonstrates how a DvR Leaf node benefits from the redistribution of local and static routes.

By default, if the injection of default routes is enabled on a DvR Controller, the DvR Leaf node can only route traffic to other nodes within the DvR enabled subnet. For the Leaf node to reach networks outside of the DvR enabled subnet, the Controllers must redistribute local and static routes from non-DvR subnets into the DvR domain. In the following figure, the DvR Leaf L1 can route traffic only to nodes in the DvR enabled subnet 10.10.10.0/24. To be able to reach hosts in VLAN 20 (20.20.20.0/24) or VLAN 30 (30.30.30.0/24), redistribution of local routes into DvR is required at each of the Controllers C1 and C2. For the Leaf node to reach hosts in remote networks 40.40.40.0/24 or 50.50.50.0/24, redistribution of static routes to the DvR domain is required.

You can apply route policies to control which local or static routes are to be redistributed into the DvR domain.

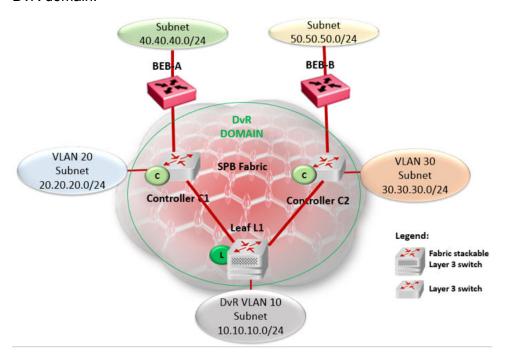


Figure 7: Redistribution of IPv4 local and static routes

Redistribution of Routes to OSPF or BGP

For non-SPB routers to benefit from the host accessibility information learned within a DvR domain, DvR supports the redistribution of host routes into OSPF or BGP. Redistribution of these host routes

is only by the DvR Controllers and only for the intra-domain host routes within the DvR enabled subnets.

A DvR Controller can redistribute host routes for all hosts from a DvR domain into OSPF or BGP. You can also apply route policies on the Controller to select the routes to be redistributed. The Controller supports redistribution of routes from the Global Router or within a VRF instance. You can also configure the metric of the route before redistribution.

The following example demonstrates the benefit of redistribution of routes to BGP.

Consider a 10.1.0.0/16 network with a stretched Layer 2 VSN spanning two data centers. On the campus side of the network, BGP peering is configured between a non-Extreme router and one or more routers in the data center. BGP advertises the network route 10.1.0.0/16 to the campus BGP routers. Depending on which edge router the traffic is delivered to, it is possible that traffic from a host on the campus traverses the WAN a second time to reach the server that is physically connected to one segment of the data center, as shown in the following figure.

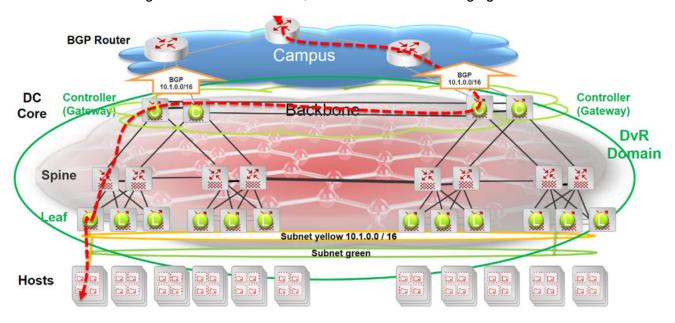


Figure 8: Inefficient traffic flow

Redistribution of the host routes from the DvR Controller to BGP solves this problem.

The following figure shows two DvR domains (show in green and orange) configured at each data center. Each campus edge router establishes a BGP peering session with one or more Controllers in each data center (DvR domain). This enables BGP to advertise more *specific* routes to the campus BGP router so that the optimal routing path is always taken. So, there is no need for traffic to traverse the WAN multiple times. Also, in the case of server movement within or between data centers, the updated DvR host routes are propagated to BGP, thus ensuring that traffic flowing into the data center continues along the most optimal path.

For example, in the following figure, only the Controller attached to the Leaf node where the 10.1.0.111 server exists, advertises its accessibility over the 10.1.0.111/32 route. Similarly, the DvR Controller associated with the Leaf node connected to the 10.1.0.222 server advertises the 10.1.0.222/32 host route.

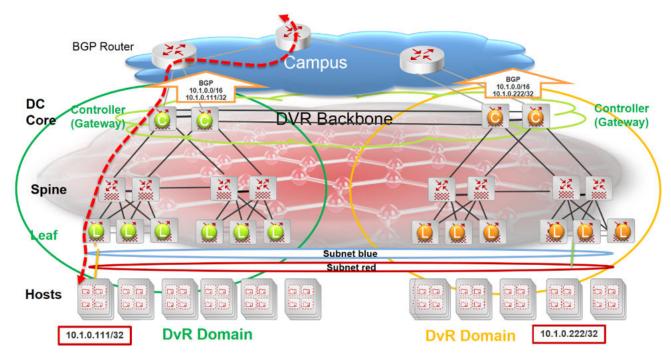


Figure 9: Traffic flow optimized with route redistribution

Controllers in each data center learn all host routes through the DvR backbone, but since those routes belong to different DvR domains, they are not all eligible for redistribution to OSPF or BGP.

Route Redistribution and IS-IS Accept Policies

DvR route redistribution leverages IS-IS accept policies to control (accept or reject) DvR routes learned from the DvR backbone. You can configure accept policies on both Controllers and non-DvR BEBs in the SPB network.

For more information about accept policies, see <u>Configuring Fabric Basics and Layer 2 Services for VOSS</u>.

DvR Deployment for Wireless Roaming in Campus Deployments

In fabric deployments where IP subnets/I-SIDs stretch between multiple buildings, you can use DvR instead of VRRP or RSMLT to avoid traffic tromboning issues. This deployment is supported with non-fabric switches that have Fabric Attach enabled, or switches that do not support Fabric Attach. You can use VSP 4000 Series as DvR Leafs in this context as well. IP subnets can stretch across one or multiple DvR domains as shown in the following figure.

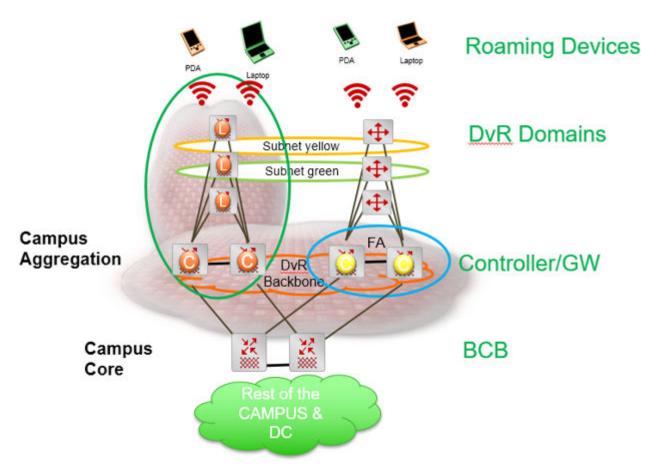


Figure 10: Wireless roaming in Campus

DvR Limitations

Review the following limitations and behavioral characteristics associated with DvR.

- The DvR feature does not affect out-of-band management on a switch chassis, if the chassis supports it.
- The DvR feature does not support a non-DvR BEB in a DvR enabled Layer 2 VSN.
- The number of host route records that can be stored in the datapath of a Leaf node is limited to the scaling capacity of the switch node. Different switch platforms have different scaling capacities.

For information on the scaling capacities of different platforms, see Release Notes for VOSS.

- You must first disable DvR on a Controller or Leaf node, before you attempt to change the domain ID of the node.
- You cannot configure IGMP snooping on DvR enabled nodes.

DvR is only supported for IPv4.

Configuration Limitations on a DvR Controller

• If you are using two different IP addresses for the DvR VLAN and the DvR GW IP, you must first configure a gateway IPv4 address and then configure an IP interface for the VLAN before you enable DvR on a Layer 2 VSN (VLAN). Both the VLAN IP address and the gateway IPv4 address must be in the same subnet.

For more information, see Configuring VLANs, Spanning Tree, and NLB for VOSS.

You cannot configure VRRP on a DvR VLAN.

Note:

A DvR VLAN is a VLAN configured on a DvR Controller with a VLAN IP address, a VLAN/I-SID, the DvR gateway IP address, and DvR enabled. This Layer 3 configuration for the DvR VLAN (the DvR gateway IP address and this DvR subnet) is pushed to the DvR Leaf nodes. The DvR gateway IP address must be the same address across all DvR Controllers for that DvR VLAN.

- You cannot configure RSMLT on a DvR VLAN.
- You cannot configure SPB-PIM Gateway (SPB-PIM GW) on a DvR VLAN.
- DvR-enabled VLAN/ISIDs are for host connectivity only; you cannot connect a router to a DvRenabled VLAN/ISID and use dynamic or static routing. Use a non-DvR VLAN/ISID instead to connect an external router.

Configuration Limitations on a DvR Leaf

- Enabling the DvR-leaf-mode boot flag before you configure a node as a DvR Leaf, automatically removes all existing non-DvR configuration on the node such as platform VLANs and their IP address configuration, CLIP configuration, routing protocol configuration and VRF configuration. The gateway IPv4 address, if configured, is also removed.
- You cannot configure SPB-PIM GW on a Leaf node. The configuration is supported only on a DvR Controller.
- You cannot configure Microsoft NLB on a Leaf node.
- You cannot configure Fabric Extend on a Leaf node.
- You cannot configure the VXLAN Gateway on a Leaf node.
- You cannot configure a T-UNI on a Leaf node.
- You cannot configure IPv4 multicast on a Leaf node. The configuration is supported only on a DvR Controller.
- You can configure only one instance of vIST on a Leaf node pair. Also, you cannot configure vIST on Leaf nodes from different domains.
- You cannot configure platform VLANs on a Leaf node. Only configuration of SPBM B-VLANs is supported.

- You cannot configure IP Shortcuts and IP Multicast over Fabric Connect on Leaf nodes. This configuration is pushed from the DvR Controllers in the domain.
- You must manually configure an I-SID on a Layer 2 VSN, on the Leaf node. This configuration is not pushed from a DvR Controller.
- DvR-enabled VLAN/ISIDs are for host connectivity only; you cannot connect a router to a DvR-enabled VLAN/ISID and use dynamic or static routing. Use a non-DvR VLAN/ISID instead to connect an external router.

Migrate from VRRP to DvR

About this task

If you have a VRRP network with a mix of existing routers that do not support DvR and VOSS devices that do support DvR, you can migrate your VRRP network to DvR using this high-level process. This migration process assumes the following design:

- Existing routers are the VRRP masters.
- Existing routers are the default gateways for all subnets.
- Fabric Connect network with DvR-capable nodes where DvR is configured globally, but not on I-SIDs, on the VOSS devices; and VOSS devices operate in Layer 2 mode for the VRRP VLANs that need to be migrated.

! Important:

When you configure DvR on Controllers with existing VRRP VLANs, ensure there is no VRRP VLAN with VRID 37 or VRID 38. VRID 37 conflicts with the DvR gateway MAC used by all DvR nodes. The DvR gateway MAC is a constant value 00:00:5e:00:01:25; VRRP VRID 37 translates to the same MAC. Similarly, VRRP VRID 38 translates to 00:00:5e:00:01:26, and is used within DvR. If you have a VRRP VLAN with either of these VRIDs, change the VRID to a different value.

Procedure

- Enable VRRP interfaces on the DvR Controllers but keep VRRP mastership on the existing routers.
- 2. Change VRRP mastership on the VLAN or IP Subnet in question on the DvR Controller by applying a higher priority than the current master.

Note:

You can easily fall back to the original VRRP master to change VRRP priorities back.

- 3. Disable VRRP on the existing routers.
- 4. Subnet by subnet (VLAN/I-SID), delete VRRP interfaces on all DvR Controllers first (this includes removing VRRP and removing the VLAN IP address), and then configure DvR interfaces (this includes adding the DVR-GW-IP, enabling DvR, and then adding the VLAN IP address) on the VLAN/I-SID instead. This might lead to a short traffic interruption.



Note:

For each VLAN/I-SID, ensure that VRRP is disabled on all nodes before you configure DvR interfaces on the Controllers for the VLAN/I-SID.

Keep in mind that you can only enable DvR on VLAN or I-SIDs where all participating BEBs are DvR-capable.

Anytime when falling back, you can delete the DvR interface on the I-SID (this includes disabling DvR, removing the DVR-GW-IP and removing the VLAN IP address) and configure the VRRP interface again (this includes adding the VLAN IP address and adding VRRP again), however, ensure you delete the DvR interfaces on all Controllers first before you enable VRRP again.

Example

Start with VRRP VLAN:

```
vlan create 250 name vlan_test250 type port-mstprstp 0
vlan i-sid 250 111250
interface vlan 250
ip address 192.0.2.3 255.255.255.0
interface vlan 250
ip vrrp version 2
ip vrrp address 10 192.0.2.1
ip vrrp 10 priority 180
ip vrrp 10 backup-master enable
ip vrrp 10 enable
```

Change to DvR VLAN:

```
interface vlan 250
no ip vrrp address 10 192.0.2.1
no ip address 192.0.2.3
exit.
interface vlan 250
dvr gw-ipv4 192.0.2.1
dvr enable
ip address 192.0.2.3 255.255.255.0
```

Change back to VRRP VLAN:

```
interface vlan 250
no dvr enable
no dvr gw-ipv4
no ip address 192.0.2.3
exit
interface vlan 250
ip address 192.0.2.3 255.255.255.0
ip vrrp version 2
ip vrrp address 10 192.0.2.1
ip vrrp 10 priority 180
ip vrrp 10 backup-master enable
ip vrrp 10 enable
exit
```

DvR configuration using the CLI

The following sections describe configuration of Distributed Virtual Routing (DvR) using the Command Line Interface (CLI).

Configuration Limitations on a DvR Controller

If you are using two different IP addresses for the DvR VLAN and the DvR GW IP, you must
first configure a gateway IPv4 address and then configure an IP interface for the VLAN before
you enable DvR on a Layer 2 VSN (VLAN). Both the VLAN IP address and the gateway IPv4
address must be in the same subnet.

For more information, see Configuring VLANs, Spanning Tree, and NLB for VOSS.

You cannot configure VRRP on a DvR VLAN.



A DvR VLAN is a VLAN configured on a DvR Controller with a VLAN IP address, a VLAN/I-SID, the DvR gateway IP address, and DvR enabled. This Layer 3 configuration for the DvR VLAN (the DvR gateway IP address and this DvR subnet) is pushed to the DvR Leaf nodes. The DvR gateway IP address must be the same address across all DvR Controllers for that DvR VLAN.

- You cannot configure RSMLT on a DvR VLAN.
- You cannot configure SPB-PIM Gateway (SPB-PIM GW) on a DvR VLAN.
- DvR-enabled VLAN/ISIDs are for host connectivity only; you cannot connect a router to a DvR-enabled VLAN/ISID and use dynamic or static routing. Use a non-DvR VLAN/ISID instead to connect an external router.

Configuring a DvR Controller

About this task

Configuring a node as a DvR Controller enables DvR globally on that node.

Perform this procedure to create a DvR domain with the domain ID that you specify, and configure the role of the node as the Controller of that domain. A Controller can belong to only one DvR domain.



For a node to perform the role of both a Controller and a Leaf within a DvR domain, you must configure it as a Controller.

Before you begin

• Ensure that you configure IP Shortcuts on the node. This is necessary for proper functioning of the node as a DvR Controller.

• Ensure that the dvr-leaf-mode boot flag is disabled on the node.

To verify the setting, enter show boot config flags in Privileged EXEC mode.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure a DvR Controller.

```
dvr controller <1-255>
```

3. **(Optional)** Disable DvR on a DvR Controller.

```
no dvr controller
```



Caution:

Disabling DvR on a DvR Controller destroys the domain ID and all dynamic content learned within the DvR domain.

However the switch retains the VLAN specific configuration and you can view the information using the command show running-config.

4. View a summary of the Controller configuration. Enter:

```
show dvr
```

Example

Configure a node as a DvR Controller:

```
Switch:1>enable
Switch: 1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #dvr controller 5
Switch:1(config) #show dvr
               DVR Summary Info
______
Domain ID
Domain ISID
: 16678219
: 16678216
: Controller
My SYS ID : 00:bb:00:00:81:21
Operational State : Up
GW MAC
InjectDefaultRouteDisabl
                             : 5
```

Variable definitions

Use the data in the following table to use the dvr controller command.

Variable	Value
<1-255>	Specifies the domain ID of the DvR domain that the controller belongs.

Disabling injection of default routes on a Controller

About this task

By default, a DvR Controller injects default routes into the DvR domain and all the Leaf nodes in that domain learn these routes with the next hop as the Controller that advertised them.

You can however disable default route injection for the GRT or a specific VRF on a Controller, to override this behavior.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Disable default route injection for the GRT or a specific VRF, on the Controller.

On the GRT:

```
dvr controller inject-default-route-disable
```

The default or the no operator enables injection of default routes for the GRT into the domain.

On a VRF instance:

```
dvr inject-default-route-disable
```

The default or the no operator enables injection of default routes for a specific VRF into the domain.

3. Verify the configuration.

On the GRT:

show dvr

On a VRF instance:

show dvr 13vsn

Example

Disable injection of default routes for the GRT on a Controller.

```
Domain ID : 5

Domain ISID : 16678219

Backbone ISID : 16678216

Role : Controller

My SYS ID : 00:bb:00:00:81:21

Operational State : Up

GW MAC : 00:00:5e:00:01:25

InjectDefaultRouteDisable(GRT) : Enabled
```

Disable injection of default routes for a specific VRF on a Controller.

Configuring DvR route redistribution

About this task

Configure redistribution of direct or static routes into the DvR domain, on the Global Router or for a specific VRF instance.

Procedure

 Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

- 2. Configure route redistribution of direct routes:
 - a. Configure route redistribution of direct routes on a VRF. The route type is *internal*.

```
dvr redistribute direct [metric <0-65535>]|[route-map WORD<1-64>]
```

b. Enable route redistribution.

```
dvr redistribute direct enable
```

c. Apply the configuration:

```
dvr apply redistribute direct
```

d. (Optional) Disable route redistribution of direct routes.

```
no dvr redistribute direct
```

- 3. Configure route redistribution of static routes:
 - a. Configure route redistribution of static routes on a VRF. The route type is external.

```
dvr redistribute static [metric <0-65535>]|[route-map WORD < 1-64>]
```

b. Enable route redistribution.

```
dvr redistribute static enable
```

c. Apply the configuration.

```
dvr apply redistribute static
```

d. (Optional) Disable route redistribution of static routes.

```
no dvr redistribute static
```

4. Verify the route redistribution configuration. You can also verify it on a specific VRF instance.

```
show dvr redistribute [vrf WORD<1-16>]
```

Example

Configure route redistribution of direct and static routes on the Global Router. Ensure that you apply the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config) #dvr redistribute static
Switch:1(config) #dvr redistribute static metric 200
Switch:1(config) #dvr redistribute static enable
Switch:1(config) #dvr apply redistribute static

Switch:1(config) #dvr redistribute direct
Switch:1(config) #dvr redistribute direct metric 100
Switch:1(config) #dvr redistribute direct enable
Switch:1(config) #dvr redistribute direct enable
Switch:1(config) #dvr apply redistribute direct
```

Verify configuration on the Global Router:

```
Switch:1(config) #show dvr redistribute

DVR Redistribute List - GlobalRouter

SOURCE MET MTYPE ENABLE RPOLICY

STAT 200 External TRUE -

LOC 100 Internal TRUE -
```

Configure redistribution of direct and static routes on the specific VRF instance vrf1. Ensure that you apply the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```
Switch:1(config) #router vrf vrf1
Switch:1(router-vrf) #dvr redistribute static
Switch:1(router-vrf) #dvr redistribute static metric 20000
Switch:1(router-vrf) #dvr redistribute static enable
Switch:1(router-vrf) #exit
Switch:1(config) #dvr apply redistribute static

Switch:1(router-vrf) #dvr redistribute direct
Switch:1(router-vrf) #dvr redistribute direct metric 10000
Switch:1(router-vrf) #dvr redistribute direct enable
Switch:1(router-vrf) #exit
Switch:1(config) #dvr apply redistribute static
```

Verify configuration on vrf1:

Variable definitions

Use the data in the following table to use the dvr redistribute direct or the dvr redistribute static commands.

Variable	Value
enable	Enables DvR route redistribution on the VRF instance.
	Route redistribution is enabled by default.
metric <0-65535>	Specifies the DvR route redistribution metric.
route-map WORD<1-64>	Specifies the route policy for DvR route redistribution.

Use the data in the following table to use the show dvr redistribute command.

Variable	Value
vrf WORD<1-16>	Specifies the VRF name.

Clearing DvR host entries

About this task

Clear DvR host entries (IPv4 remote host routes) on a Controller. The host entries are learned on the switch, either locally on its UNI port or dynamically from other nodes in the DvR domain.

Note:

You can clear DvR host entries only on a DvR Controller.

An error message displays if you attempt clearing of host entries on a DvR Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Clear the DvR host entries.

```
clear dvr host-entries [ipv4 {A.B.C.D}} | [12isid <0-16777215>] |
[13isid <0-16777215>]
```

Example

In this example, you clear host entries for IP address 50.0.1.0 to clear host entries for IP addresses 50.0.1.2 and 50.0.1.3.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clear dvr host-entries 50.0.1.0
```

Variable definitions

Use the data in the following table to use the clear dvr host-entries command.

Variable	Value	
ipv4	Specifies the IP address (IPv4) of the DvR host entries to clear.	
12isid	Specifies the Layer 2 VSN I-SID of the DvR host entries to clear	
	The range is 1 to 16777215.	
13isid	Specifies the Layer 3 VSN I-SID of the DvR host entries to clear.	
	The range is 0 to 16777215.	

Configuring a DvR Leaf

About this task

Perform this procedure to create a DvR domain with the domain ID that you specify, and configure the role of the node as a Leaf node. Configuring a node as a DvR Leaf automatically enables DvR globally on the node.

A Leaf node can belong to only one DvR domain.

Note:

For a node to perform the role of both a Controller and a Leaf within the domain, you must configure it as a Controller.

Note:

You must enable the VRF-scaling boot configuration flag on a DvR Leaf node, if more than 24 VRFs are required in the DvR domain.

For additional scaling information, see Release Notes for VOSS.

Before you begin

• You must enable the dvr-leaf-mode boot flag before you configure a node as a DvR Leaf node.

To verify the setting, enter show boot config flags in Privileged EXEC mode.



Caution:

Ensure that you save the current configuration on the switch, before you enable the flag.

Enabling the flag removes all existing non-DvR configuration on the switch, such as platform VLANs and their IP address configuration, CLIP configuration, routing protocol configuration and VRF configuration. The gateway IPv4 address, if configured, is also removed.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure a node as a DvR Leaf.

```
dvr leaf < 1-255 >
```

3. (Optional) Disable DvR on a DvR Leaf.

```
no dvr Leaf
```



Caution:

Disabling DvR on a Leaf node removes its membership with the DvR domain and all the dynamic content learned from the Controllers of that domain.

4. View a summary of the Leaf configuration.

```
show dvr
```

Example

Configure a node as a DvR Leaf:

```
Switch2:1>enable
Switch2:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

Variable definitions

Use the data in the following table to use the dvr leaf command.

Variable	Value
<1-255>	Specifies the domain ID of the DvR domain to which the Leaf node belongs.

Configuring vIST on a DvR Leaf node pair

Before you begin

Ensure that the nodes are configured as DvR Leaf nodes, before you configure vIST.

About this task

When you configure vIST on a DvR Leaf node pair, the switch generates an I-SID from the configured cluster ID. This I-SID is unique across the SPB network as long as the cluster ID is unique across the SPB network, for the vIST pair. You can configure only one instance of vIST on the Leaf node pair.

To configure vIST, both nodes must be Leaf nodes. You cannot configure vIST, for example, on a Controller-Leaf node pair.

Also both the nodes must belong to the same DvR domain. vIST configuration over Leaf nodes in different domains is not supported.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure vIST on the Leaf nodes:

```
dvr leaf virtual-ist \{\langle A.B.C.D/X | \langle A.B.C.D \rangle \langle A.B.C.D \rangle\} peer-ip
{A.B.C.D} cluster-id <1-1000>
```

3. (Optional) Disable vIST on the DvR Leaf node pair.

```
no dvr leaf virtual-ist
```



Caution:

Disabling DvR on a Leaf node in a vIST pair removes all vIST configuration on that node, but not on the pair. The node on which DvR is disabled also loses its membership with the DvR domain and all the dynamic content learned from the Controllers in that domain.

If DvR is re-enabled on the node, you must manually configure vIST on that node again.

4. View a summary of vIST configuration on the Leaf nodes.

show dvr

Example

Configure vIST on DvR Leaf nodes, with IP addresses 51.51.51.1 and 51.51.51.2 respectively:

```
Switch2:1>enable
Switch2:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config) #dvr leaf virtual-ist 51.51.51.1 peer-ip 51.51.51.2 cluster-id 255
Switch2:1#show dvr
______
                    DVR Summary Info
_____
Domain ID
Domain ISID
                               : 16678219
                               : Leaf
Role
My SYS ID
Role

      My SYS ID
      : 00:bb:00:00:71:23

      Operational State
      : Up

      GW MAC
      : 00:00:5e:00:01:25

      Inband Mgmt Clip TP

Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID : 16677226
```

Variable definitions

Use the data in the following table to use the dvr leaf virtual-ist command.

Variable	Value
{ <a.b.c.d x <a.b.c.d=""> <a.b.c.d>}</a.b.c.d></a.b.c.d>	Specifies the local IP (IPv4) address and subnet mask of the node.
{ <a.b.c.d>}</a.b.c.d>	Specifies the IP address (IPv4) of the vIST peer.

Table continues...

Variable	Value
<1–1000>	Specifies the cluster ID of vIST.
	It is set to 0 if vIST is not configured.

Configure a Management VLAN on a DvR Leaf Node

About this task

You cannot configure a platform VLAN directly on a DvR leaf node. However, you can configure a VLAN Management Instance on a DvR leaf node. After you configure the management VLAN, you can configure a platform VLAN.

Procedure

1. Enter Global Configuration mode:

enable
configure terminal

2. Create a management VLAN and associate it with an existing port-based VLAN:

mgmt vlan <2-4059>



If you attempt to configure a VLAN ID that is already in use by DVR for created DVR interfaces, the following error message displays:

Error: This vlan id is already used by DVR. Please configure another vlan id for the mgmt vlan

3. Create a platform VLAN:

vlan create <2-4059> [name WORD<0-64>] type port-mstprstp <0-63> [color <0-32>]

4. Enter VLAN Management Instance mode:

mgmt vlan

5. Enable the management VLAN:

enable

Variable Definitions

The following table defines parameters for the mgmt vlan command.

Variable	Value
<2-4059>	Specifies the existing port-based VLAN ID to associate with the management VLAN.

The following	table defines	parameters for the vlan	create command
THE IDIOWING	table delilles	parameters for the vian	Create Communication.

Variable	Value	
<2-4059>	Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.	
name WORD<0-64>	Specifies the VLAN name. The name attribute is optional.	
type port-mstprstp <0-63> [color <0-32>]	Creates a VLAN by port:	
	<0-63> is the STP instance ID from 0 to 63.	
	• color <0-32> is the color of the VLAN in the range of 0 to 32.	
	Note:	
	MSTI instance 62 is reserved for SPBM if SPBM is enabled on the switch.	

Delete a Management VLAN on a DvR Leaf Node

You must perform deletion steps in the correct order.

If you attempt to delete the Management Instance VLAN before you delete the platform VLAN, the following error message displays:

Error: On a DVR leaf, the vlan must be deleted before the mgmt vlan instance

If you attempt to delete the platform VLAN before disabling the Management Instance VLAN, the following error message displays:

Error: Cannot delete vlan if mgmt vlan interface is enabled

About this task Procedure

1. Enter Global Configuration mode:

enable
configure terminal

2. Enter VLAN Management Instance mode:

mgmt vlan

3. Disable the management VLAN:

no enable

4. Delete the platform VLAN:

no vlan <2-4059>

5. Delete the management VLAN:

```
no mgmt vlan
```

Variable Definitions

The following table defines parameters for the **no vlan** command.

Table 23:

Variable	Value
<2-4059>	Specifies the ID of the platform VLAN to be deleted.

Moving a vIST Leaf node pair from one domain to another

About this task

Use this procedure to move a vIST Leaf node pair from one DvR domain to another.

For vIST to work properly, both Leaf nodes must be in the *same* domain.

Procedure

1. Disable IS-IS on each vIST peer Leaf node, to remove the node from the SPB network.

```
no router isis enable
```

2. Disable DvR on each Leaf node.

no dvr leaf



Caution:

Disabling DvR on a Leaf node in a vIST pair automatically removes all vIST configuration on that node, but not on the pair. The node on which DvR is disabled also loses its membership with the DvR domain and all the dynamic content learned from the Controllers in that domain.

When you re-enable DvR on the node, you must manually configure vIST on that node again.

3. Configure each node as a DvR Leaf node, with the new domain ID.

Ensure that you configure both nodes as Leaf nodes and with the same domain ID.

```
dvr leaf < 1-255 >
```

4. Configure vIST on the DvR Leaf nodes.

```
dvr leaf virtual-ist {<A.B.C.D/X|<A.B.C.D> <A.B.C.D>} peer-ip
{A.B.C.D} cluster-id <1-1000>
```

5. Enable IS-IS on each vIST peer Leaf node, to add back the node to the SPB network.

```
router isis enable
```

Example

Consider two vIST peer Leaf nodes Switch1 (IP address 51.51.51.1) and Switch2 (51.51.51.2) that belong to a DvR domain (with domain ID 4), that you need to move to another domain (with domain ID 5).

View a summary of existing Leaf configuration on each node.

```
Switch1:1(config) #show dvr
______
                DVR Summary Info
______
Domain ID : 4
Domain ISID : 166
Domain ID

Domain ISID

Role

Role

State

My SYS ID

Operational State

GW MAC

Inband Mgmt Clip IP

Virtual Ist local address

Virtual Ist local subport mack

255, 255, 255, 0
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID : 16677226
Switch2:1(config) #show dvr
______
              DVR Summary Info
______
Domain ID
Domain ISID
Domain ISID : 16678220

Role : Leaf

My SYS ID : 00:00:72:55:45:00

Operational State : Up

GW MAC : 00:00:5e:00:01:25

Inband Mgmt Clip IP :

Virtual Ist local address : 51.51.51.1
Virtual Ist local subnet mask : 255.255.255.0
Virtual Ist peer address : 51.51.51.2
Virtual Ist cluster-id : 255
Virtual Ist ISID : 16677226
```

Disable IS-IS globally on each Leaf node.

```
Switch1:1>en
Switch1:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config) #router isis
Switch1:1(config-isis) #no router isis enable
Switch1:1(config-isis) #exit
Switch1:1(config) #

Switch2:1>en
Switch2:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config) #router isis
Switch2:1(config-isis) #no router isis enable
Switch1:1(config-isis) #exit
Switch1:1(config-isis) #exit
```

Disable DvR on each node. This automatically removes all vIST configuration on the node, but not on the vIST pair. The node also loses its membership with the DvR domain and all the dynamic content learned from the Controllers in that domain.

```
Switch1:1(config) #no dvr leaf
Switch2:1(config) #no dvr leaf
```

Configure each node as a DvR Leaf, with domain ID 5.

```
Switch1:1(config) #dvr leaf 5
Switch2:1(config) #dvr leaf 5
```

Configure vIST on each of the DvR Leaf nodes.

```
Switch1:1(config) #dvr leaf virtual-ist 51.51.51.1 peer-ip 51.51.51.2 cluster-id 255
Switch2:1(config) #dvr leaf virtual-ist 51.51.51.2 peer-ip 51.51.51.1 cluster-id 255
```

Enable IS-IS globally on each Leaf node.

```
Switch1:1(config) #router isis
Switch1:1(config-isis) #router isis enable
Switch1:1(config-isis) #exit
Switch1:1(config) #
Switch2:1(config) #router isis
Switch2:1(config-isis) #router isis enable
Switch2:1(config-isis) #exit
Switch2:1(config-isis) #exit
```

View a summary of Leaf configuration on each node.

```
Switch2:1(config)#show dvr
```

Virtual Ist cluster-id : 255 : 16677226 Virtual Ist ISID

Moving a vIST Controller pair from one domain to another

About this task

Use this procedure to move a vIST Controller node pair from one DvR domain to another.

For vIST to work properly, both Controller nodes must be in the same domain.

Procedure

1. Disable IS-IS on each vIST peer Controller node, to remove the node from the SPB network.

```
no router isis enable
```

2. Disable DvR on each Controller node:

no dvr controller



Caution:

Disabling DvR on a DvR Controller destroys the domain ID and all dynamic content learned within the DvR domain. However, the switch retains the VLAN specific configuration which you can view using the command show running-config.

3. Configure each node as a DvR Controller node, with the new domain ID. Ensure that you configure both nodes as Controller nodes and with the same domain ID.

```
dvr controller <1-255>
```

4. Enable IS-IS on each vIST peer Controller node, to add back the node to the SPB network.

```
router isis enable
```

Example

Consider two vIST peer Controller nodes Switch1 (IP address 51.51.51.3) and Switch2 (51.51.51.4) that belong to a DvR domain (with domain ID 4), that you need to move to another domain (with domain ID 5).

View a summary of Controller configuration on each node:

```
Switch1:1(config) #show dvr
_____
         DVR Summary Info
______
Domain ID
Domain ISID
                : 16678220
Backbone ISID
                : 16678216
Role
                 : Controller
My SYS ID
               : 00:bb:00:00:81:21
Operational State : Up
```

```
: 00:00:5e:00:01:25
GW MAC
InjectDefaultRouteDisable(GRT) : Disabled
Switch2:1(config) #show dvr
             DVR Summary Info
______
              : 4
                     : 16678220
Domain ISID
                     : 16678216
Backbone ISID
My SYS ID
                      : Controller
                     : 00:bb:00:00:82:22
My SYS ID
Operational State : Up
GW MAC
                      : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
```

Disable IS-IS globally on each Controller node:

```
Switch1:1>en
Switch1:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch1:1(config) #router isis
Switch1:1(config-isis) #no router isis enable
Switch1:1(config-isis) #exit
Switch1:1(config) #

Switch2:1>en
Switch2:1>conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch2:1(config) #router isis
Switch2:1(config-isis) #no router isis enable
Switch1:1(config-isis) #exit
Switch1:1(config) #
```

Disable DvR on each node:

```
Switch1:1(config) #no dvr controller
Switch2:1(config) #no dvr Controller
```

Configure each node as a DvR Controller, with domain ID 5.

```
Switch1:1(config) #dvr controller 5
Switch2:1(config) #dvr controller 5
```

Enable IS-IS globally on each Controller node.

```
Switch1:1(config) #router isis
Switch1:1(config-isis) #router isis enable
Switch1:1(config-isis) #exit
Switch1:1(config) #
Switch2:1(config) #router isis
Switch2:1(config-isis) #router isis enable
Switch2:1(config-isis) #exit
Switch2:1(config-isis) #exit
```

View a summary of Controller configuration on each node.

```
Switch1:1(config) #show dvr

-------

DVR Summary Info
------

Domain ID : 5
```

Distributed Virtual Routing

Domain ISID : 16678221
Backbone ISID : 16678216
Role : Controller
My SYS ID : 00:bb:00:00:81:21
Operational State : Up
GW MAC : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled

Switch2:1(config) #show dvr

DVR Summary Info

Domain ID : 5
Domain ISID : 16678221
Backbone ISID : 16678216
Role : Controller

My SYS ID : 00:bb:00:00:82:22

Operational State : Up

GW MAC : 00:00:5e:00:01:25

InjectDefaultRouteDisable(GRT) : Disabled

View the vIST configuration on each of the Controller nodes.

Switch1:1>show virtual-ist ______ IST Info VLAN ENABLE IST
ID IST STA STATUS -----51.51.51.2 4002 true up NEGOTIATED MASTER/ DIALECT IST STATE NONE up Master Switch2:1>show virtual-ist ______ IST Info ______ VLAN ENABLE IST ID IST STAT PEER-IP ADDRESS STATUS -----

NEGOTIATED MASTER/DIALECT IST STATE SLAVE

true

up

4002

NONE up Slave

51.51.51.1

Configuring a non-DvR BEB to join the DvR backbone

About this task

Configure a non-DvR backbone edge bridge (BEB) to join the DvR backbone so that it can receive redistributed DvR host routes from all DvR Controllers in the SPB network.



On a non-DVR BEB, the redistributed host routes from the DvR backbone are not automatically installed in the IP routing table. To utilize the backbone host routes to optimize traffic forwarding (forwarding in the data plane), you must explicitly configure an IS-IS accept policy with a backbone route policy using the command accept backbone-route-map <rul>
route-map
name>, and specifying a suitable route-map to select the list or range of DvR backbone host routes to be installed in the routing table.

For more information on configuring an IS-IS accept policy with a backbone route policy, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Procedure

1. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

2. Configure a non-DvR BEB to join the DvR backbone.

```
[no]|[default] backbone enable
```

- 3. Verify the configuration using the following commands.
 - show dvr backbone-members
 - show dvr backbone-members non-dvr-beb
 - show dvr backbone-entries
 - show isis

Example

Configure the non-DvR BEB to join the DvR backbone.

Switch3:1(config-isis) #backbone enable

Verify the configuration. View the DvR backbone members.

```
Switch3:1(config-isis) #show dvr backbone-members
                                           DVR BB Members
______

        System Name
        Nick-Name
        Nodal Factor

        DVR-8284-D2-C1-40
        0.82.40
        00:00:82:84:40:00
        NON-DVR-BEB DVR-8284-D2-C2-41

        DVR-8284-D2-C2-41
        0.82.41
        00:00:82:84:41:00
        Controller

System Name
                                     Nick-Name
                                                    Nodal MAC
2 out of 2 Total Num of DVR Backbone Members displayed
Switch3:1(config-isis) #show dvr backbone-members non-dvr-beb
                                           DVR BB Members
______
                                                    Nodal MAC Role
                                     Nick-Name
                                    0.82.40 00:00:82:84:40:00 NON-DVR-BEB
DVR-8284-D2-C1-40
1 out of 2 Total Num of DVR Backbone Members displayed
```

View the backbone DvR host routes that the non-DvR BEB receives from other Controllers in the SPB network.

	DVR Backbone-Entries					
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
39.1.1.4	10:cd:ae:70:5d:01	401	10390	200	DVR-8284-D2-C2-41	DVR-8284-D2-C2-41
39.2.1.4	10:cd:ae:70:5d:01	401	10391	200	DVR-8284-D2-C2-41	DVR-8284-D2-C2-41
39.3.1.4	10:cd:ae:70:5d:01	401	10392	200	DVR-8284-D2-C2-41	DVR-8284-D2-C2-41
39.4.1.4 39.5.1.4	10:cd:ae:70:5d:01 10:cd:ae:70:5d:01	401 401	10393 10394	200 200	DVR-8284-D2-C2-41 DVR-8284-D2-C2-41	DVR-8284-D2-C2-41 DVR-8284-D2-C2-41
39.6.1.4	10:cd:ae:70:5d:01	401	10395	200	DVR-8284-D2-C2-41	DVR-8284-D2-C2-41

View the IS-IS related information.

```
Switch3:1(config-isis) #show isis
______
                        ISIS General Info
                       AdminState : enabled
                       RouterType : Level 1
                       System ID : 00bb.0000.8121
               Max LSP Gen Interval : 900
                          Metric : wide
                Overload-on-startup: 20
                        Overload : false
                    Csnp Interval: 10
                    PSNP Interval : 2
                 Rxmt LSP Interval : 5
                      spf-delay : 100
Router Name : router_r1
                 ip source-address:
                ipv6 source-address:
           ip tunnel source-address :
                      Tunnel vrf :
                         ONA Port :
```

```
ip tunnel mtu:

Num of Interfaces: 3

Num of Area Addresses: 1

Inband Mgmt Clip IP: 72.54.44.1

backbone: enabled

Dynamically Learned Area: 00.0000.0000

FAN Member: No
```

DvR show commands

The following section explains the show commands for DvR.

Viewing DvR summary

Use this procedure to view a summary of the DvR configuration on a DvR Controller or a DvR Leaf.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View a summary of DvR configuration:

show dvr

Example

View the information on a DvR Controller:

View the information on a DvR Leaf:

Virtual Ist cluster-id : 255 Virtual Ist ISID : 16677226

Job aid

Use the data in the following table to use the **show dvr** command output.

On a Controller:

Field	Descriptions
Domain ID	Specifies the domain ID of the DvR domain to which the Controller belongs.
Domain I-SID	Specifies the DvR domain I-SID.
	The range is 16678217 to 16678727.
Backbone I-SID	Specifies the backbone I-SID.
	The value is 16678216.
Role	Specifies the role of the node in the DvR domain, namely Controller.
My SYS ID	Specifies the MAC address of the Controller.
Operational State	Specifies the operational state of the Controller.
GW MAC	Specifies the gateway MAC address.
InjectDefaultRouteDisable	Specifies whether injection of default routes is disabled on the Controller. The default is disabled.

On a Leaf node:

Field	Descriptions
Domain ID	Specifies the domain ID of the DvR domain to which the Leaf node belongs.
Domain I-SID	Specifies the DvR domain I-SID.
	The range is 16678217 to 16678727.
Role	Specifies the role of the node in the DvR domain, namely Leaf.
My SYS ID	Specifies the MAC address of the Leaf node.
Operational State	Specifies the operational state of the Leaf node.
GW MAC	Specifies the gateway MAC address.
Inband Mgmt Clip IP	Specifies the in-band management CLIP IP address.
Virtual Ist local address	Specifies the local IP address of the node, if vIST is configured.
Virtual Ist local subnet mask	Specifies the subnet mask of the local IP address of the node, if vIST is configured.
Virtual Ist peer address	Specifies the IP address of the peer node, in the vIST pair.

Table continues...

Field	Descriptions
Virtual Ist cluster-id	Specifies the cluster ID if vIST is configured.
Virtual Ist ISID	Specifies the I-SID if vIST is configured.

Viewing members of a DvR domain

About this task

View the members of all DvR domains, namely the Controllers and Leaf nodes.

You can view this information on either a Controller or a Leaf node. Both the Controller and the Leaf node displays those members of the DvR domain to which it belongs.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. show dvr members [controller|leaf]

Example

View all members of a DvR domain:

	DVR Members (Domain ID: 255)	
System Name	Nick-Name	Nodal MAC	Role
Leaf-4:110	0.41.10	00:bb:00:00:41:10	Leaf
Leaf-1:Q:123	0.71.23	00:bb:00:00:71:23	Leaf
Leaf-2:K:124	0.71.24	00:bb:00:00:71:24	Leaf
Leaf-3:K:125	0.71.25	00:bb:00:00:71:25	Leaf
Ctrl-1:Q:121	0.81.21	00:bb:00:00:81:21	Controller
Ctrl-2:0:122	0.81.22	00:bb:00:00:81:22	Controller

View member DvR Controllers:

	DVR Members (Domain ID: 255)	
System Name	Nick-Name	Nodal MAC	Role
Ctrl-1:Q:121 Ctrl-2:Q:122	0.81.21 0.81.22	00:bb:00:00:81:21 00:bb:00:00:81:22	Controller Controller

View member DvR Leaf nodes:

Switch:1#show dvr members	leaf
	DVR Members (Domain ID: 255)

System Name	Nick-Name	Nodal MAC	Role
Leaf-4:110 Leaf-1:Q:123 Leaf-2:K:124 Leaf-3:K:125	0.41.10 0.71.23 0.71.24 0.71.25	00:bb:00:00:41:10 00:bb:00:00:71:23 00:bb:00:00:71:24 00:bb:00:00:71:25	Leaf Leaf Leaf Leaf
4 out of 6 Total Num of DVR Memb	ers displayed		

Job aid

Use the data in the following table to use the show dvr members command output.

Field	Descriptions
System Name	Specifies the system name of the DvR member (Controller or Leaf node).
Nick-Name	Specifies the nick name of the DvR member.
Nodal MAC	Specifies the nodal MAC address of the DvR member.
Role	Specifies the role of the DvR member within the DvR domain, that is Controller or Leaf.

Viewing DvR interfaces

View the DvR interfaces on either a Controller or a Leaf node.

On Controllers, DvR interfaces are created when you configure IP on a DvR enabled Layer 2 VSN (VLAN, I-SID). Only Controllers display the administrative state of the interfaces because this is where you enable or disable the interfaces. The Leaf nodes display DvR interface information that is pushed from the Controllers, for example, subnet routes or gateway IP addresses for the Layer 2 VSNs.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the DvR interface information.

On a Controller:

show dvr interfaces [13isid <0-16777215>] [vrf WORD<1-16>] [vrfids WORD<0-512>]

On a Leaf node:

```
show dvr interfaces [13isid <0-16777215>]
```

Viewing the DvR interface information for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View DvR interfaces on a Controller node:

You can view DvR interface information on all interfaces or for a specific Layer 3 I-SID, VRF, or VRF ID.

				DVR	Interface	es		
Interface	Mask	L3ISID	VRFID	L2ISID	VLAN	GW IPv4	Admin State	SPBMC IGMP State Versic
50.0.1.2	255.255.0.0	55500	1	50500	500	50.0.1.1	enable	disable 2

View DvR interfaces on a Leaf node:

You can view DvR interface information on all interfaces or for a specific Layer 3 I-SID. Viewing the interface information for a specific VRF or VRF ID is not supported on a DvR Leaf node.

		DVR Inte	rfaces ======		======	
Interface	Mask	L3ISID	VRFID	L2ISID	VLAN	GW IPv4
40.1.0.0	255.255.0.0	401	2	10401	77	40.1.1.11
40.2.0.0	255.255.0.0	401	2	10402	78	40.2.1.11
40.3.0.0	255.255.0.0	401	2	10403	79	40.3.1.11
40.4.0.0	255.255.0.0	401	2	10404	80	40.4.1.11

Variable definitions

Use the data in the following table to use the show dvr interfaces command.

Variable	Value
13isid	Specifies the Layer 3 I-SID of the DvR interface.
	The range is 0 to 16777215.
vrf	Specifies the VRF name.
vrfids	Specifies the VRF ID.
	The range is 0 to 512.

Job aid

Use the data in the following table to use the **show dvr interfaces** command output.

Field	Descriptions
Interface	Specifies the VLAN IP address (IPv4) of the DvR interface.
Mask	Specifies the subnet mask of the VLAN IP address.

Table continues...

Field	Descriptions
L3ISID	Specifies the Layer 3 I-SID of the DvR interface.
	The range is 0 to 16777215.
VRFID	Specifies the VRF ID of the DvR interface
L2ISID	Specifies the Layer 2 I-SID of the DvR interface.
	The range is 1 to 16777215.
VLAN	Specifies the VLAN ID of the DvR interface.
GW IPv4	Specifies the DvR gateway IP address (IPv4).
Admin State	Specifies the administrative state of the DvR interface.
	Note:
	This field displays only on a Controller node.
SPBMC State	Specifies the SBPMC state of the DvR interface.
IGMP version	Specifies the version of IGMP running on the DvR interface.

Viewing DvR host entries

About this task

View DvR host entries (IPv4 remote host routes) on either a Controller or a Leaf node. The node displays the host entries learned either locally on its Switched UNI port or dynamically from other nodes within the DvR domain.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the DvR host entries.

On a Controller:

```
show dvr host-entries [domain-id <1-255>]|[ipv4 {A.B.C.D}]|[l2isid <1-16777215>]|[l3isid <0-16777215>]|[nh-as-mac]|[type <1-2>]|[vrf WORD<1-16>] [vrfids WORD<0-512>]
```

On a Leaf node:

```
show dvr host-entries [domain-id <1-255>] | [ipv4 {A.B.C.D}] | [l2isid <1-16777215>] | [l3isid <0-16777215>] | [nh-as-mac] | [type <1-2>]
```

Viewing the DvR host entries for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View DvR host entries on either a Controller or a Leaf node.

Viewing the DvR host entries for a specific VRF or VRF ID is not supported on a DvR Leaf node.

DVR Host-Entries								
P-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	VRFID	PORT	DOMAIN ID	TYPE	NEXT HOP
50.0.1.2 50.0.1.3	b0:ad:aa:42:ed:04 b0:ad:aa:4c:3d:01	55500 55500	50500 50500	0	2/23 cpp	255 255	DYNAMIC LOCAL	Cont-1:121 Cont-2:122

View DvR host entries for a specific IP address.

In this example, you enter IP address 50.0.1.0 to display host entries for IP addresses 50.0.1.2 and 50.0.1.3.

				DVR Host-E	Intries			
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	VRFID	PORT	DOMAIN ID	TYPE	NEXT HOP
50.0.1.2 50.0.1.3	b0:ad:aa:42:ed:04 b0:ad:aa:4c:3d:01	55500 55500	50500 50005	0	2/23 cpp	2 2	DYNAMIC LOCAL	Cont-1:121 Cont-2:122

View DvR host entries where the next hop displays the MAC address instead of the system name.

DVR Host-Entries								
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	VRFID	PORT	DOMAIN ID	TYPE	NEXT HOP
50.0.1.2 50.0.1.3	b0:ad:aa:42:ed:04 b0:ad:aa:4c:3d:01	55500 55500	50500 50500	0	2/23 cpp	2 2	DYNAMIC LOCAL	00:bb:00:00:01:03

View DvR host entries based on the host type. Type 1 indicates local hosts and type 2 dynamic hosts.

DVR Host-Entries								
P-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	VRFID	PORT	DOMAIN ID	TYPE	NEXT HOP
0.0.1.2	b0:ad:aa:42:ed:04	55500	50500	0	2/23	2	DYNAMIC	00:bb:00:00:01:0

Variable definitions

Use the data in the following table to use the show dvr host-entries command.

Variable	Value
domain-id	Specifies the domain ID of the DvR host entry.
	The range is 1 to 255.

Table continues...

Variable	Value
ipv4	Specifies the IP address (IPv4) of the DvR host entry.
12isid	Specifies the Layer 2 VSN I-SID of the DvR host entry.
	The range is 1 to 16777215.
l3isid	Specifies the Layer 3 VSN I-SID of the DvR host entry.
	The range is 0 to 16777215.
nh-as-mac	Specifies the MAC address of the next hop node instead of the system name.
type	Specifies the host type of the DvR host entry.
	A value of 1 indicates local hosts and a value of 2 indicates dynamic hosts.
vrf	Specifies the VRF name of the DvR host entry.
vrfids	Specifies the VRF ID of the DvR host entry.
	The range is 0 to 512.

Job aid

Use the data in the following table to use the show dvr host-entries command output.

Field	Descriptions
IP-ADDRESS	Specifies the IP address of the DvR host entry (IPv4 remote ARP).
HOST MAC-ADDRESS	Specifies the MAC address of the DvR host entry (IPv4 remote ARP).
L3VSN ISID	Specifies the Layer 3 VSN I-SID of the DvR host entry.
VRFID	Specifies the VRF ID of the DvR host entry.
L2VSN ISID	Specifies the Layer 2 VSN I-SID of the DvR host entry.
PORT	Specifies the port of the DvR host entry.
DOMAIN ID	Specifies the DvR domain ID of the DvR host entry.
TYPE	Specifies the host type of the DvR host entry.
NEXT HOP	Specifies the next hop system MAC address of the DvR host entry.

Viewing DvR routes

About this task

View the DvR routes (IPv4 network routes) on a DvR Controller or a Leaf node.

Controllers display all the IP subnet routes configured for that DvR domain. The Leaf nodes display the IP subnet routes that are learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. Leaf nodes also display routes that are redistributed by Controllers (direct routes, static routes and the default route), into the DvR domain.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the DvR routes.

On a Controller:

```
show dvr routes [ipv4 \{A.B.C.D\}] | [13isid <0-16777215>] | [nh-as-mac] | [vrf WORD<1-16>] | [vrfids WORD<0-512>]
```

On a Leaf node:

```
show dvr routes [ipv4 \{A.B.C.D\}] | [13isid < 0-16777215>] | [nh-as-mac]
```

Viewing the DvR routes for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View DvR routes on either a Controller or a Leaf node.

Viewing the DvR routes for a specific VRF or VRF ID is not supported on a DvR Leaf node.

			OVR Routes				
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	TYPE	COST
0.0.0.0	255.255.0.0	Ctrl-1:8400:121	0	55500	50500		1

View DvR routes where the next hop MAC address is displayed instead of the system name:

DVR Routes							
DEST	MASK	NEXT HOP	VRFID	L3VSN TSTD	L2VSN TSTD	TYPE	COST
50.0.0.0	255.255.0.0	00:bb:00:01:02		55500	50500		

Variable definitions

Use the data in the following table to use the show dvr routes command.

Variable	Value
ipv4 {A.B.C.D}	Specifies the IP address (IPv4) of the DvR route.
I3isid <0-16777215>	Specifies the Layer 3 I-SID of the DvR route.

Table continues...

Variable	Value
	The range is 0 to 16777215.
nh-as-mac	Specifies the MAC address of the next hop node instead of the system name.
vrf	Specifies the VRF name of the DvR route.
vrfids	Specifies the VRF ID of the DvR route.
	The range is 0 to 512.

Job aid

Use the data in the following table to use the show dvr routes command output.

Field	Descriptions
DEST	Specifies the IPv4 destination address of the DvR route.
MASK	Specifies the subnet mask of the IPv4 destination address of the DvR route.
NEXT HOP	Specifies the host name of the next hop BEB, in the DvR route.
VRFID	Specifies the VRF ID of the DvR route.
L3VSN ISID	Specifies the Layer 3 VSN I-SID of the DvR route.
L2VSN ISID	Specifies the Layer 2 VSN I-SID of the DvR route.
TYPE	Specifies the route type of the DvR route.
COST	Specifies the SPB cost of the DvR route.

Viewing DvR database information

About this task

View all DvR routes on a Controller or a Leaf node.

The Controller node displays all the IP subnet routes configured for that DvR domain. A Leaf node displays all IP subnet routes learned from the Controller(s) for the L2 VSNs in the DvR Domain. It also displays the Host Routes (ARPs) learned from other DvR enabled nodes.

Before you begin

Ensure that DvR is enabled globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the DvR database.

On a Controller:

```
show dvr database [ipv4 {A.B.C.D}] | [13isid<0-16777215>] | [nh-as-mac] | [vrf WORD<1-16>] | [vrfids WORD<0-512>]
```

On a Leaf node:

show dvr database [ipv4 $\{A.B.C.D\}$] | [13isid<0-16777215>] | [nh-as-mac]

Viewing the DvR database for a specific VRF or VRF ID is not supported on a DvR Leaf node.

Example

View the DvR database on either a Controller or a Leaf node.

Viewing the DvR database for a specific VRF or VRF ID is not supported on a DvR Leaf node.

				DVR DATA	BASE				
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE		PREFIX COST	AGE
40.0.0.0 40.0.1.2 40.0.1.3	255.255.0.0 255.255.255.255 255.255.255.255		0 0 101	0 0 0	40400 40400 40400	cpp cpp Ctrl1-Ctrl2	10 10 10	1	0 day(s), 05:44:55 0 day(s), 05:44:55 0 day(s), 05:44:30

View the DvR database for a specific IPv4 address:

Switch:1#sho	w dvr database ipv4	40.3.1.2							
				DVR DATA	ABASE				
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPE COST	PREFI	IX AGE
40.3.1.2	255.255.255.255	Ctrl-1:K:121	0	0	40403	cpp	10	1	0 day(s), 05:50:0
1 out of 122	5 Total Num of DVR	Database entries di	isplayed	ı					

View DvR database entries for a specific L3 I-SID.

SWITCH: I#Show C	vr database 13is	10 U		DVR DATA	====== BASE		=====			
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB COST	PREFI COST	X AGE	======
40.0.0.0 40.0.1.2 40.0.1.3	255.255.0.0 255.255.255.255 255.255.255.255		0 0 0	0 0 0	40400 40400 40400	cpp cpp Ctrl1-Ctrl2		1 1 1	0 day(s), 0 day(s), 0 day(s),	05:44:55
3 out of 3 Tota	l Num of DVR Data	abase entries displ	ayed							

View DvR database entries with next hop MAC address displayed instead of the system name:

			Ι	OVR DATABAS	E					
DEST	MASK	NEXT HOP	VRFID	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB COST	PREFI COST	X AGE	
40.0.0.0 40.0.1.2 40.0.1.3	255.255.255.255	00:bb:00:00:81:21 00:bb:00:00:81:21 00:bb:00:00:81:22	0 0 0	0 0 0	40400 40400 40400	cpp cpp Ctrl1-Ctrl2	10 10 10	1 1 1	0 day(s), 0 0 day(s), 0 0 day(s), 0	5:44:55

Variable definitions

Use the data in the following table to use the show dvr database command.

Variable	Value
ipv4 {A.B.C.D}	Specifies the IP address (IPv4) of the DvR database entry.
I3isid <0-16777215>	Specifies the Layer 3 I-SID of the DvR database entry.
	The range is 0 to 16777215.
nh-as-mac	Specifies the MAC address of the next hop node instead of the system name.
vrf	Specifies the VRF name of the DvR database entry.
vrfids	Specifies the VRF ID of the DvR database entry.
	The range is 0 to 512.

Job aid

Use the data in the following table to use the show dvr database command output.

Field	Descriptions
DEST	Specifies the address type of the IPv4 destination address of the DvR database entry.
MASK	Specifies the destination mask of the DvR database entry.
NEXT HOP	Specifies the MAC address of the next hop BEB, in the DvR database entry.
VRFID	Specifies the VRF ID for the database entry.
L3VSN ISID	Specifies the Layer 3 VSN I-SID of the DvR database entry.
L2VSN ISID	Specifies the Layer 2 VSN I-SID of the DvR database entry.
OUTGOING INTERFACE	Specifies the outgoing interface (port or MLT) of the DvR database entry.
SPB COST	Specifies the SPB cost of the DvR database entry.
PREFIX COST	Specifies the prefix cost of the DvR database entry.
AGE	Specifies the uptime since creation of the DvR database table entry.

Viewing DvR backbone entries

About this task

View the DvR backbone entries (redistributed host routes) learned from all Controllers in all DvR domains.



Note:

DvR backbone entries can be viewed only on a Controller. Viewing backbone entries is not applicable on a Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View DvR backbone entries:

show dvr backbone-entries [adv-controller WORD < 1-255 >] | [domain-id < 1-255 >] | [host-mac-address $0 \times 00: 0 \times 00]$ | [ipv4 {A.B.C.D}] | [12isid < 1-16777215 >] | [13isid < 0-16777215 >] | [next-hop WORD < 1-255 >] | [nh-as-mac]

Example

View all DvR backbone entries:

				DVR Backb	one-Entries	
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
40.0.1.2 40.0.1.2 40.0.1.3 40.0.1.3	b0:ad:aa:4c:55:00 b0:ad:aa:4c:55:00 b0:ad:aa:43:31:00 b0:ad:aa:43:31:00	0 0 0 0	40400 40400 40400 40400	255 255 255 255 255	Ctrl-2:8200:122 Ctrl-1:8400:121 Ctrl-1:8400:121 Ctrl-2:8200:122	Ctrl-1:8400:121 Ctrl-1:8400:121 Ctrl-2:8200:122 Ctrl-2:8200:122

View DvR backbone entries on a specific DvR Controller:

				DVR Backbo	one-Entries	
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
0.0.1.2	b0:ad:aa:4c:55:00 b0:ad:aa:43:31:00	0 0	40400 40401	255 255	Ctrl-2:8200:122 Ctrl-2:8200:122	Ctrl-1:8400:121 Ctrl-2:8200:122

View DvR backbone entries for a specific host MAC address:

DVR Backbone-Entries									
P-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP			
0.0.1.2 0.0.1.2	b0:ad:aa:4c:55:00 b0:ad:aa:4c:55:00	0 0	40400 40400	255 255	Ctrl-2:8200:122 Ctrl-1:8400:121	Ctrl-1:8400:121 Ctrl-1:8400:121			

View DvR backbone entries for a specific IP address:

In this example, you enter IP address 40.0.1.0 to display backbone entries for IP addresses 40.0.1.2 and 40.0.1.3.

```
Switch:1#show dvr backbone-entries ipv4 40.0.1.0

DVR Backbone-Entries
```

IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
40.0.1.2 40.0.1.2 40.1.1.3 40.1.1.3 4 out of 4 Total	b0:ad:aa:4c:55:00 b0:ad:aa:4c:55:00 b0:ad:aa:43:31:00 b0:ad:aa:43:31:00 al Num of DVR Backbor	0 0 0 0 ne Routes d	40400 40400 40401 40401 isplayed	255 255 255 255	Ctrl-2:8200:122 Ctrl-1:8400:121 Ctrl-2:8200:122 Ctrl-2:8200:121	Ctrl-1:8400:121 Ctrl-1:8400:121 Ctrl-2:8200:122 Ctrl-2:8200:122

View DvR backbone entries for a specific L3 VSN I-SID:

				DVR Backbo	one-Entries	
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
40.0.1.2 40.0.1.2 40.0.1.3 40.0.1.3	b0:ad:aa:4c:55:00 b0:ad:aa:4c:55:00 b0:ad:aa:43:31:00 b0:ad:aa:43:31:00	0 0 0 0	40400 40400 40400 40400	255 255 255 255 255	Ctrl-2:8200:122 Ctrl-1:8400:121 Ctrl-1:8400:121 Ctrl-2:8200:122	Ctrl-1:8400:121 Ctrl-1:8400:121 Ctrl-2:8200:122 Ctrl-2:8200:122

View DvR backbone entries for a specific next hop node:

				DVR Backbo	one-Entries	
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
40.0.1.2 40.0.1.2	b0:ad:aa:4c:55:00 b0:ad:aa:4c:55:00	0	40400 40400	255 255	Ctrl-2:8200:122 Ctrl-1:8400:121	Ctrl-1:8400:121 Ctrl-1:8400:121

View DvR backbone entries where the next hop nodes are displayed as MAC addresses:

DVR Backbone-Entries						
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	DOMAIN ID	ADV-CONTROLLER	NEXT HOP
10.0.1.2 10.0.1.2	20.444.44.10.00.00	0 0	40400 40400	255 255	Ctrl-2:8200:122 Ctrl-1:8400:121	00:bb:00:00:81:21 00:bb:00:00:81:21

Variable definitions

Use the data in the following table to use the show dvr backbone entries command.

Variable	Value
adv-controller WORD<1-255>	Specifies the system name of the advertising Controller.
domain-id <1-255>	Specifies the domain ID of the DvR backbone entry.
	The range is 1 to 255.
host-mac-address 0x00:0x00:0x00:0x00:0x00:0x00	Specifies the host MAC address of the DvR backbone entry.

Table continues...

Variable	Value
ipv4 {A.B.C.D}	Specifies the IP address (IPv4) of the DvR backbone entry.
I2isid <1-16777215>	Specifies the Layer 2 I-SID of the DvR backbone entry.
	The range is 1 to 16777215.
I3isid <0-16777215>	Specifies the Layer 3 I-SID of the DvR backbone entry.
	The range is 0 to 16777215.
next-hop WORD<1-255>	Specifies the system name of the next hop node.
nh-as-mac	Specifies the MAC address of the next hop node instead of the system name.

Job aid

Use the data in the following table to use the show dvr backbone-entries command output.

Field	Descriptions
IP-ADDRESS	Specifies the IPv4 address of the DvR backbone host.
HOST MAC-ADDRESS	Specifies the MAC address of DvR backbone host.
L3VSN ISID	Specifies the Layer 3 VSN I-SID of the DvR backbone host.
L2VSN ISID	Specifies the Layer 2 VSN I-SID of the DvR backbone host.
DOMAIN ID	Specifies the domain ID of the DvR backbone host.
ADV-CONTROLLER	Specifies the host name of the advertising Controller.
NEXT HOP	Specifies the MAC address of the next hop backbone host in the DvR route.

Viewing DvR backbone members

About this task

DvR backbone members are either DvR Controllers or non-DvR BEBs that receive redistributed host routes from all other DvR Controllers in the SPB network.

Before you begin

Ensure that DvR is enabled globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View DvR backbone member information:

show dvr backbone-members [controller|non-dvr-beb]

Example

View all DvR backbone members:

Switch:1#show dvr backbone-members

DVR BB Members				
System Name	Nick-Name	Nodal MAC	Role	Domain Id
DVR-D2-C1-40 Ctrl-2:8200:122	0.82.40 0.81.22	00:00:82:84:40:00 00:bb:00:00:81:22	NON-DVR-BEB Controller	2 2
2 out of 2 Total Num of DVR Backbone Members displayed				

View backbone members that are DvR controllers:

Switch:1#show dvr backbone-members controller					
DVR BB Members (Domain ID: 255)					
System Name	Nick-Name	Nodal MAC	Role	Domain Id	
Ctrl-2:8200:122	0.81.22	00:bb:00:00:81:22	Controller	2	
1 out of 2 Total Num of DVR Backbone	out of 2 Total Num of DVR Backbone Members displayed				

View backbone members that are non-DvR BEBs:

Role Domain Id
ROIE DOMAIN IQ
40:00 NON-DVR-BEB 2
Į

Variable definitions

Use the data in the following table to use the show dvr backbone-members command.

Variable	Value
controller	Specifies backbone members that are DvR Controllers.
non-dvr-beb	Specifies backbone members that are non-DvR BEBs.

Job aid

Use the data in the following table to use the show dvr backbone-members command output.

Field	Description
System Name	Specifies the system name of the DvR backbone member.
Nick-Name	Specifies the nick name of the DvR backbone member.
Nodal MAC	Specifies the nodal MAC address of the DvR backbone member.
Role	Specifies the role of the DvR backbone member.
Domain Id	Specifies the domain ID of the backbone member.

Viewing Layer 3 VSN information

About this task

View VRFs corresponding to Layer 3 (routed) VSN I-SIDs on either a Controller or a Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View the Layer 3 VSN information:

```
show dvr 13vsn [13isid <0-16777215>] | [vrf WORD<1-16>] | [vrfids WORD<0-512>]
```

Example

View Layer 3 VSN information on a DvR Controller:

View Layer 3 VSN information on a DvR Leaf node:

Variable definitions

Use the data in the following table to use the show dvr 13vsn command.

Variable	Value
I3isid <0-16777215>	Specifies the Layer 3 VSN I-SID.
	The range is 0 to 16777215.
vrf WORD<1-16>	Specifies the VRF name of the VRF corresponding to the Layer 3 VSN I-SID.
vrfids WORD<0-512>	Specifies the VRF ID of the VRF.

Job aid

Use the data in the following table to use the **show dvr** 13vsn command output on a DvR Controller.

Field	Description
VRF ID	Specifies the VRF ID of the VRF corresponding to the Layer 3 VSN I-SID.
L3VSN ISID	Specifies the Layer 3 VSN I-SID.
VRF NAME	Specifies the VRF name of the VRF corresponding to the Layer 3 VSN I-SID.
INJECT-DEFAULT-ROUTE-DISABLE	Specifies whether injection of default routes is disabled.

Use the data in the following table to use the **show dvr 13vsn** command output on a DvR Leaf node.

Field	Description
VRF ID	Specifies the VRF ID of the VRF corresponding to the Layer 3 VSN I-SID.
L3VSN ISID	Specifies the Layer 3 VSN I-SID.
VRF NAME	Specifies the VRF name of the VRF corresponding to the Layer 3 VSN I-SID.

Viewing DvR domain redistribution information

About this task

View DvR domain redistribution information on a Controller or a Leaf node.



You can view DvR domain redistribution information only on a DvR Controller.

An error message displays if you attempt to view this information on a DvR Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. Enter Privileged EXEC mode:

enable

2. View DvR domain redistribution information:

```
show dvr redistribute [vrf WORD<1-16>] | [vrfids WORD<0-512>]
```

Example

View DvR domain redistribution information on a Controller:

Switch:1#show dvr redistribute

```
DVR Redistribute List - GlobalRouter

SOURCE MET MTYPE ENABLE RPOLICY

STAT 1 External TRUE -
```

View DvR domain redistribution information for a particular VRF.

```
Switch:1#show dvr redistribute vrf vrf1

DVR Redistribute List - VRF vrf1

SOURCE MET MTYPE ENABLE RPOLICY

STAT 20000 External TRUE -
LOC 10000 Internal TRUE -
```

Variable definitions

Use the data in the following table to use the **show dvr redistribute** command.

Variable	Definitions
vrf WORD<1-16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID of the VRF.

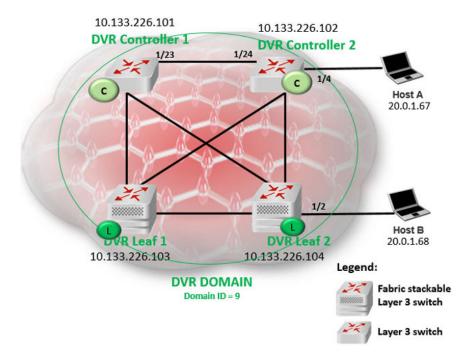
Job aid

Use the data in the following table to use the show dvr redistribute command output.

Field	Description
SOURCE	Specifies the source of the DvR route redistribution.
MET	Specifies the DvR route redistribution metric. The range is 0 to 65535.
MTYPE	Specifies the DvR route redistribution metric type.
ENABLE	Specifies whether DvR route redistribution is enabled on the VRF instance.
RPOLICY	Specifies the route policy for DvR route redistribution.

Configuring a DvR solution

The following section describes a simple configuration example to configure Distributed Virtual Routing (DvR) over a Fabric Connect (SPB) network.



About this task

In this example, you configure two DvR Controllers (with IP addresses 10.133.226.101 and 10.133.226.102) and two DvR Leaf nodes (with IP addresses 10.133.226.103 and 10.133.226.104), in a single DvR domain with domain ID 9. Hosts connect to the DvR nodes as shown in the figure.

Before you begin

On the switches to be configured as DvR Controllers:

- Ensure that you configure Fabric Connect.
- Ensure that you configure IP Shortcuts on the node. This is necessary for proper functioning of the node as a DvR Controller.
- Verify that the dvr-leaf-mode boot flag is disabled on the node. To verify the setting, enter show boot config flags in Privileged EXEC mode.

On the switches to be configured as DvR Leaf nodes:

• Ensure that you configure Fabric Connect.

Procedure

DvR Controller configuration — Controller 1 and Controller 2:

 Verify configuration of Fabric Connect on each of the switches to be configured as the DvR Controllers. The following examples show verification on one of the switches. Perform this verification on both switches.

a. Verify the SPB configuration:

```
Switch1:1>en
Switch1:1#show spbm
                      spbm : enable
                   ethertype : 0x8100
              nick-name server : enable
            nick-name allocation : static
           nick-name server range : B.00.00-B.FF.FF
Switch1:1#show isis spbm
_____
                  ISIS SPBM Info
SPBM B-VID PRIMARY NICK LSDB IP IPV6 MULTICAST SPB-PIM-GW INSTANCE VLAN NAME TRAP
1 4051-4052 4051 0.10.01 disable enable disable enable disable
______
                ISIS SPBM SMLT Info
______
     SMLT-SPLIT-BEB SMLT-VIRTUAL-BMAC
SPBM
                              SMLT-PEER-SYSTEM-ID
INSTANCE
  primary
                  00:00:00:00:00:00
```

b. Verify the global IS-IS configuration:

Total Num of SPBM instances: 1

```
Switch1:1#show isis
______
                       ISIS General Info
______
                     AdminState : enabled
                     RouterType : Level 1
                      System ID : 00bb.0000.0101
              Max LSP Gen Interval : 900
                        Metric : wide
              Overload-on-startup: 20
                      Overload : false
                   Csnp Interval: 10
                   PSNP Interval : 2
                Rxmt LSP Interval : 5
                      spf-delay: 100
                    Router Name : Cont-1
                ip source-address: 10.0.0.101
              ipv6 source-address :
          ip tunnel source-address :
                     Tunnel vrf :
                   ip tunnel mtu :
                Num of Interfaces: 4
             Num of Area Addresses : 1
              Inband Mgmt Clip IP :
                      backbone : disabled
          Dynamically Learned Area: 00.0000.0000
                    FAN Member : No
```

- 2. Configure the DvR Controllers.
 - a. Configure Controller 1 (IP address 10.133.226.101) with DvR domain ID 9.

```
Switch:1>en
Switch: 1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #dvr controller 9
Switch:1(config) #show dvr
______
           DVR Summary Info
______
Domain ID
Domain ISID
                        : 16678219
Backbone ISID
Role : Cor
My SYS ID : 00:
Operational State : Up
                        : Controller
                       : 00:bb:00:00:81:21
                        : 00:00:5e:00:01:25
InjectDefaultRouteDisable(GRT) : Disabled
```

b. Configure Controller 2 (IP address 10.133.226.102), also with DvR domain ID 9.

```
Switch:1>en
Switch: 1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config) #dvr controller 9
Switch: 1 (config) #show dvr
______
        DVR Summary Info
______
Domain ID : 9
Domain ISID : 16
                       : 16678219
Backbone ISID
                      : 16678216
Role
My SYS ID
Operational State
                      : Controller
                       : 00:bb:00:00:81:21
                       : 00:00:5e:00:01:25
GW MAC
InjectDefaultRouteDisable(GRT) : Disabled
```

c. Verify the configuration. View the members of the DvR domain.

Layer 2 VSN (VLAN) configuration on the DvR Controllers:

- 3. Configure Layer 2 VSN on the DvR Controllers, Controller 1 and Controller 2.
 - a. Configure platform VLANs on Controller 1 (VLAN ID=200 and VLAN ID=202).

 Associate the VLANs with the I-SIDs 20200 and 20202 respectively. Configure

gateway IPv4 addresses 20.0.1.1 and 20.2.1.1 respectively, and enable DvR on those interfaces.

```
Switch1:1(config) #vlan create 200 type port-mstprstp 0
Switch1:1(config) #vlan i-sid 200 20200
Switch1:1(config) #interface vlan 200
Switch1:1(config) #dvr gw-ipv4 20.0.1.1
Switch1:1(config) #dvr enable
Switch1:1(config) #ip address 20.0.1.2 255.255.0.0

Switch1:1(config) #vlan create 202 type port-mstprstp 0
Switch1:1(config) #vlan i-sid 202 20202
Switch1:1(config) #interface vlan 202
Switch1:1(config) #dvr gw-ipv4 20.2.1.1
Switch1:1(config) #dvr enable
Switch1:1(config) #ip address 20.2.1.2 255.255.0.0
Switch1:1(config) #exit
Switch1:1#
```

b. Configure the platform VLANs on Controller 2. Ensure that you configure the same gateway IPv4 addresses on the corresponding VLANs, as on Controller 1.

```
Switch2:1(config) #vlan create 200 type port-mstprstp 0
Switch2:1(config) #vlan i-sid 200 20200
Switch2:1(config) #interface vlan 200
Switch2:1(config) #dvr gw-ipv4 20.0.1.1
Switch2:1(config) #dvr enable
Switch2:1(config) #ip address 20.0.1.3 255.255.0.0

Switch2:1(config) #vlan create 202 type port-mstprstp 0
Switch2:1(config) #vlan i-sid 202 20202
Switch2:1(config) #interface vlan 202
Switch2:1(config) #dvr gw-ipv4 20.2.1.1
Switch2:1(config) #dvr enable
Switch2:1(config) #ip address 20.2.1.3 255.255.0.0
Switch2:1(config) #exit
Switch2:1#
```

c. Verify Layer 2 VSN (VLAN) configuration on the Controllers. The following example shows the verification on Controller 1. Perform this verification on both Controllers.

View the DvR interfaces.

On Controllers, DvR interfaces are created when you configure IP on a DvR enabled Layer 2 VSN (VLAN, I-SID). You can also view the administrative state of these interfaces on the Controller.

				DVR Int	terfaces				
Interface	Mask	L3ISID	VRFID	L2ISID	VLAN	GW IPv4	Admin State	SPBMC State	IGMP Version
20.0.1.2 20.2.1.2	255.255.0.0 255.255.0.0	0	0	20200 20202	200 202	20.0.1.1 20.2.1.1	enable enable	disable disable	2 2

View the DvR host entries learned locally on the S-UNI port.

Switch1:1#shov	w dvr host-entries						
				DVR Host-E	ntries		
IP-ADDRESS	HOST MAC-ADDRESS	L3VSN ISID	L2VSN ISID	PORT	DOMAIN ID	TYPE	NEXT HOP
20.0.1.2	b0:ad:aa:42:ed:04	0	20200	срр	9	LOCAL	Cont-1

20.2.1.2	b0:ad:aa:42:ed:04	0	20202	срр	9	LOCAL	Cont-1
2 out of 2 Total	Num of DVR Host Ent	ries di:	splayed				

View the DvR database. All IP subnet routes configured on the Controller, for the DvR domain, are displayed.

			DVR DATA	BASE				
DEST	MASK	NEXT HOP	L3VSN ISID	L2VSN ISID	OUTGOING INTERFACE	SPB COST	PREFIX COST	AGE
20.0.1.2 06:41:40	255.255.255.255	Cont-1	0	20200	срр	10	1	1 day(s),
20.2.1.2	255.255.255.255	Cont-1	0	20202	срр	10	1	1 day(s),

View the DvR routes for the subnets 20.0.0.0 and 20.2.0.0.

```
| DVR Routes | DVR
```

Layer 3 configuration on the DvR Controllers

- 4. Configure Layer 3 (VRF) on the DvR Controllers, Controller 1 and Controller 2.
 - a. Configure Layer 3 on Controller 1. As part of this configuration, you configure a VRF vrf501 and associate it with a DvR VLAN.

```
Switch1:1(config) #ip vrf vrf501 vrfid 501
Switch1:1(config) #vlan create 501 type port-mstprstp 0
Switch1:1(config) #vlan i-sid 501 50501
Switch1:1(config) #interface Vlan 501
Switch1:1(config) #vrf vrf501
Switch1:1(config) #dvr gw-ipv4 50.1.1.1
Switch1:1(config) #dvr enable
Switch1:1(config) #ip address 50.1.1.2 255.255.0.0

Switch1:1(config) #router vrf vrf501
Switch1:1(router-vrf) #i-sid 55501
Switch1:1(router-vrf) #ipvpn enable
Switch1:1(router-vrf) #exit
Switch1:1(config) #
```

b. Configure Layer 3 on Controller 2.

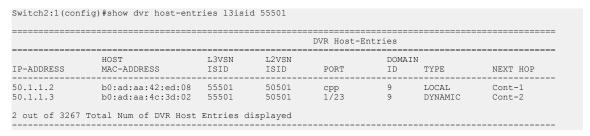
```
Switch2:1(config) #ip vrf vrf501 vrfid 501
Switch2:1(config) #vlan create 501 type port-mstprstp 0
Switch2:1(config) #vlan i-sid 501 50501
Switch2:1(config) #interface Vlan 501
Switch2:1(config) #vrf vrf501
Switch2:1(config) #dvr gw-ipv4 50.1.1.1
Switch2:1(config) #dvr enable
Switch2:1(config) #ip address 50.1.1.3 255.255.0.0

Switch2:1(config) #router vrf vrf501
Switch2:1(router-vrf) #i-sid 55501
Switch2:1(router-vrf) #ipvpn enable
```

```
Switch2:1(router-vrf)#exit
Switch2:1(config)#
```

c. Verify Layer 3 configuration. The following example shows verification on Controller 1. Perform this verification on both Controllers.

View the DvR host entries.



View the DvR interfaces.

				DVR I	nterface	3				
Interface	Mask	L3ISID	VRFID	L2ISID	VLAN	GW IPv4	Adm Sta		SPBMC State	IGMP Version
50.1.1.2	255.255.0.0	55501	501	50501	501	50.1.1.1	ena	ble	disable	2
1 out of 291 T	otal Num of DVR	Interfaces	displayed							
Switch2:1(conf	ig) #show dvr data	abase 13isi	d 55501							
Switch2:1(conf	ig)#show dvr data	abase 13isi	.d 55501	======	DVR DAT	ABASE			======	======
DEST	MASK	NEXT HOP	L3VS ISID	IS	JSN (DUTGOING INTERFACE	COST	PREFIX	AGE	======
DEST		NEXT HOP Cont-1	L3VS ISID	1 50 1 50	VSN (1D 501 501 501	DUTGOING INTERFACE	10 10	COST 1		01:26:49 01:26:49

DvR Leaf configuration — Leaf 1 and Leaf 2

5. Configure the boot flag dvr-leaf-mode on the switches to be configured as DvR Leaf nodes.



Caution:

Ensure that you save the current configuration on the switch, before you enable the flag.

Enabling the flag removes all existing non-DvR configuration on the switch, such as platform VLANs and their IP address configuration, CLIP configuration, routing protocol configuration and VRF configuration. The gateway IPv4 address, if configured, is also removed.

On switch with IP address 10.133.226.104, configure the boot flag and reboot the switch.

```
Switch3:1>en
Switch3:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch3:1(config) #boot config flags dvr-leaf-mode
Switch3:1(config) #save config
Switch3:1(config) #reset
```

On switch with IP address 10.133.226.105, configure the boot flag and reboot the switch.

```
Switch4:1>en
Switch4:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch4:1(config) #boot config flags dvr-leaf-mode
Switch4:1(config) #save config
Switch4:1(config) #reset
```

6. After the switches come back up, configure the nodes as DvR Leaf nodes.

Configure switch with IP address 10.133.226.104 as DvR Leaf 1; verify the configuration.

```
Switch3:1(config) #dvr Leaf 9
Switch3:1(config) #show dvr
_____
                             DVR Summary Info
Domain ID : 9
Domain ISID : 16
Role
                         : 16678219
Role
My SYS ID
Role
                        : Leaf
                        : 00:bb:00:00:80:05
Inband Mgmt Clip IP
Virtual Ist local
                        : Up
                         : 00:00:5e:00:01:25
Virtual Ist local subnet mask :
Virtual Ist peer address :
Virtual Ist cluster-id
Virtual Ist ISID
```

Configure switch with IP address 10.133.226.105 as DvR Leaf 2; verify the configuration.

```
Switch4:1(config) #dvr Leaf 9
Switch4:1(config) #show dvr
______
                       DVR Summary Info
______
Domain ID
Domain ISID
                   : 16678219
My SYS ID
                   : Leaf
                   : 00:bb:00:00:80:05
Inband Mgmt Clip IP
                    : Up
                   : 00:00:5e:00:01:25
Virtual Ist local subnet mask :
Virtual Ist peer address
Virtual Ist cluster-id
Virtual Ist ISID
```

7. Associate the I-SIDs on the DvR Leaf nodes to the DvR VLANs configured on the Controller.

On Leaf node 1 (IP address 10.133.226.105):

```
Switch3:1(config)#i-sid 20200 elan
Switch3:1(elan:20200)#c-vid 200 port 1/2
Switch3:1(config)#exit
```

View the host connections.

```
Switch3:1#show dvr host-entries nh-as-mac
```

				DVR Host-E	ntries		
	HOST	L3VSN	L2VSN		DOMAI	N	
IP-ADDRESS	MAC-ADDRESS	ISID	ISID	PORT	ID	TYPE	NEXT HOP
20.0.1.67	00:00:00:00:00:67	0	20200	1/4	9	DYNAMIC	00:bb:00:00:81:21
20.0.1.68	00:00:00:00:00:68	0	20200	1/2	9	DYNAMIC	00:bb:00:00:81:21
2 out of 2 Total	Num of DVR Host Ent	ries displ	ayed				

On Leaf node 2 (IP address 10.133.226.105):

```
Switch4:1(config) #i-sid 20200 elan
Switch4:1(elan:20200) #c-vid 200 port 1/2
Switch4:1(config) #exit
```

View the host connections.

				DVR Host-E	ntries		
	HOST	L3VSN	L2VSN		DOMA	I N	
P-ADDRESS	MAC-ADDRESS	ISID ISID		PORT	ID	TYPE	NEXT HOP
0.0.1.67	00:00:00:00:00:67	0	20200	1/4	9	DYNAMIC	00:bb:00:00:81:21
20.0.1.68	00:00:00:00:00:68	0	20200	1/2	9	DYNAMIC	00:bb:00:00:81:21

8. View all members of the DvR domain. You can view this information on either a Leaf node or a Controller node.

	DVR Members	(Domain ID: 2)	
ystem Name	Nick-Name	Nodal MAC	Role
ont-1	0.10.01	00:bb:00:00:01:01	Controller
ont-2	0.10.02	00:bb:00:00:01:02	Controller
eaf1	0.10.04	00:bb:00:00:80:04	Leaf
eaf2	0.10.05	00:bb:00:00:80:05	Leaf

DvR Configuration Using the EDM

The following sections describe configuration of Distributed Virtual Routing (DvR) using the Enterprise Device Manager (EDM).

Configure a DvR Controller or a DvR Leaf Globally

About this task

Configure a node to perform the role of either a Controller or a Leaf, within the DvR domain.

Before you begin



For DvR Leaf Configuration only:

You must enable the dvr-leaf-mode boot flag before you configure a node as a DvR Leaf node. Navigate to Configuration > Edit > Chassis. On the Boot Config tab, select EnableDvrLeafMode.

Ensure that you save the current configuration on the switch, before you enable the flag. Enabling the flag removes all non-DvR configuration on the switch.

Procedure

- 1. In the navigation pane, expand the **Configuration > Edit** folders.
- 2. Click DVR.
- 3. Click the Globals tab.
- 4. Enter the domain ID in the **DomainId** field.
 - Note:

A Controller or a Leaf node can belong to only one DvR domain.

- 5. Select the role of the node in the **Role** field.
- 6. **(Optional)** On a Controller node, disable injection of default routes into the DvR domain. Select **InjectDefaultRouteDisable**.
 - Note:

This field applies only to Controllers. Attempting to select this field on a Leaf node displays an error message.

7. Update the fields as necessary, and then click **Apply** to save your configuration.

Globals Field Descriptions

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

Field	Descriptions
DomainId	Uniquely identifies the domain that the node belongs to.
	The range for a Controller or a Leaf is 1 to 255. Set to 0 if is not configured.
Role	Specifies the role of the node in the domain, that is, either a Controller or a Leaf.
Enable	Specifies whether DvR is enabled on the node.
	Configuring a Controller or Leaf sets this parameter to true.

Field	Descriptions
DomainIsid	Uniquely identifies the domain I-SID that the node belongs to.
	The range is 16678217 to 16678727. 0 indicates that is not configured.
Backbonelsid	Uniquely identifies the backbone ISID that the node belongs to.
	The valid backbone I-SID is 16678216. It is set to 0 if is not configured.
GatewayMac	Specifies the Gateway MAC address used by all Domains.
InbandMgmtlp Note:	Specifies the In-band Management IP address configured under IS-IS.
Exception: not supported on VSP 8600 Series or XA1400 Series.	You can use this IP address to manage the node, irrespective of whether DvR is enabled on it.
InjectDefaultRouteDisable	Specifies whether injection of default routes is disabled on the Controller in the domain.
	By default, Controllers inject default routes into the domain so that all Leaf nodes in the domain learn these routes with the next hop as the Controller that advertised it. Selecting this field disables this behavior.
VirtuallstLocalAddr	Specifies the local IP address of vIST, if vIST is configured on a Leaf.
	vIST cannot be configured on a Controller.
VirtuallstLocalMask	Specifies the local subnet mask of vIST, if vIST is configured on a Leaf.
	vIST cannot be configured on a Controller.
VirtuallstPeerAddr	Specifies the peer IP address of vIST, if vIST is configured on a Leaf.
	vIST cannot be configured on a Controller.
VirtuallstClusterId	Specifies the cluster ID of vIST, if vIST is configured on a Leaf.
	vIST cannot be configured on a Controller.
	Set to 0 if vIST is not configured.
Virtuallstlsid	Specifies the I-SID if vIST is configured.
OperState	Specifies the operational state of the node.

View DvR Routes

About this task

View the DvR routes (host routes and the IPv4 network routes) that are learned on a DvR Controller or a Leaf node.

Controllers display all the IP subnet routes configured for that DvR domain. Leaf nodes display the IP subnet routes learned from the Controller(s) for the L2 VSNs in the DvR Domain. Leaf nodes also display any redistributed routes into the DvR Domain that are learned from the Controllers (direct routes, static routes and the default route).

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand the **Configuration > Edit** folders.
- 2. Click DVR.
- 3. Click the Routes tab.
- 4. To filter the rows based on the specific criteria, click **Filter**.

Routes field descriptions

Use the data in the following table to use the Routes tab.

Name	Description
DestlpAddrType	Specifies the IPv4 destination address type of the DvR route.
DestlpAddr	Specifies the IPv4 destination address of the DvR route.
DestMask	Specifies the destination mask of the DvR route.
L3Isid	Specifies the L3 I-SID of the DvR route.
EcmpIndex	Specifies the ECMP index for the ECMP routes of the DvR route.
NextHopMac	Specifies the MAC address of the next hop BEB in the DvR route.
L2Isid	Specifies the L2 I-SID of the DvR route.
Vrfld	Specifies the VRF ID.
Cost	Specifies the SPB cost of the DvR route.
NextHopName	Specifies the host name of the next hop BEB, in the DvR route.
Туре	Specifies the route type of the DvR route.

View Members of a DvR Domain

About this task

View the members of all DvR domains namely the Controllers and Leaf nodes.

You can view this information on either a Controller or a Leaf node. Both the Controller and the Leaf node displays the members of the DvR domain to which it belongs.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand **Configuration > Edit**.
- 2. Select **DVR**.
- 3. Select the **Members** tab.
- 4. (Optional) To filter the rows based on specific criteria, click Filter.

Members field descriptions

Use the data in the following table to use the Members tab.

Name	Description
MacAddress	Specifies the system ID or the nodal MAC address of this DvR member.
SysName	Specifies the system name of this DvR member.
NickName	Specifies the nick name of this DvR member.
Role	Specifies the DvR role (Controller or Leaf) of this DvR member.
DomainId	Specifies the domain ID of the DvR domain that this member belongs to.

View DvR Backbone Members

About this task

DvR backbone members are either DvR Controllers or non-DvR BEBs that receive redistributed host routes from all other DvR Controllers in the SPB network.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** folders.

- 2. Click DVR.
- 3. Click the **Backbone Members** tab.
- 4. (Optional) To filter the rows based on specific criteria, click Filter.

Backbone Members field descriptions

Use the data in the following table to use the Backbone Members tab.

Name	Description
MacAddress	Specifies the system ID or the nodal MAC address of this DvR backbone member.
SysName	Specifies the system name of this DvR backbone member.
NickName	Specifies the nick name of this DvR backbone member.
Role	Specifies the role of this DvR backbone member.
	It is either a DvR Controller or a non-DvR BEB.
DomainId	Specifies the domain ID of the DvR domain that this backbone member belongs to.
	The domain ID is 0 for a non-DvR BEB.

View DvR Interfaces

About this task

View the DvR interfaces on either a Controller or a Leaf node.

On Controllers, DvR interfaces are created when you configure IP on a DvR enabled Layer 2 VSN (VLAN, I-SID). Only Controllers display the administrative state of the interfaces because this is where you enable or disable the interfaces. On a Leaf node, the DvR interface information that the Controllers push, for example, subnet routes and the gateway IP addresses for the Layer 2 VSNs, are displayed.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand the **Configuration > Edit** folders.
- 2. Click DVR.
- 3. Click the Interfaces tab.

Click **Filter** to filter rows based on specific filter criteria.

Interfaces field descriptions

Use the data in the following table to use the Interfaces tab.

Name	Description
VlanlpAddrType	Specifies the VLAN IP address type of the DvR interface.
VlanlpAddr	Specifies the VLAN IP address (IPv4) of the DvR interface.
L3Isid	Specifies the Layer 3 I-SID of the DvR interface.
	The range is 1 to 16777215.
L2lsid	Specifies the Layer 2 I-SID of the DvR interface.
	The range is 1 to 16777215.
VlanlpMask	Specifies the VLAN IP address mask of the DvR interface.
Vrfld	Specifies the VRF ID of the DvR interface.
	The VRF ID is 0 for the GRT.
VlanId	Specifies the VLAN ID of the DvR interface.
GwlpAddrType	Specifies the address type of the DvR gateway IP address (IPv4).
GwlpAddr	Specifies the DvR gateway IP address (IPv4).
AdminState	Specifies the administrative state of the DvR interface.
SpbmcState	Specifies the state of IP Multicast over Fabric Connect, on the DvR interface.
IgmpVersion	Specifies the version of IGMP that runs on the DvR interface.

View DvR Host Entries

About this task

View DvR host entries (IPv4 remote ARPs) on either a Controller or a Leaf node. The node displays the host entries learned either locally on its UNI port or dynamically from other nodes in the DvR domain.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **Edit** folders.
- 2. Click DVR.

- 3. Click the Host Entries tab.
- 4. (Optional) To filter the rows based on the specific criteria, click Filter.

Host Entries field descriptions

Use the data in the following table to use the Host Entries tab.

Name	Description
IpAddrType	Specifies the address type of the DvR host entry (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR host entry.
Mask	Specifies the subnet mask of the DvR host entry.
L3Isid	Specifies the Layer 3 I-SID of the DvR host entry.
MacAddr	Specifies the MAC address of the DvR host entry.
L2lsid	Specifies the Layer 2 I-SID of the DvR host entry.
Vrfld	Specifies the VRF ID associated with the DvR host entry.
Port	Specifies the port of the DvR host entry.
DomainId	Specifies the DvR domain ID of the DvR host entry.
Туре	Specifies the host type of the DvR host entry.
NextHopName	Specifies the next hop system name of the DvR host entry.
NextHopMac	Specifies the next hop system MAC address of the DvR host entry.
ClearEntry	Clears the entry if the configured value is true.

Clear DvR Host Entries

About this task

Clear DvR host entries (IPv4 remote host routes) on a Controller. The host entries are learned on the switch either locally on its UNI port or dynamically from other nodes in the DvR domain.

Note:

You can clear DvR host entries only on a DvR Controller.

An error message displays if you attempt clearing of host entries on a DvR Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **Edit** folders.
- 2. Click DVR.

- 3. Click the Clear Host Entries tab.
- 4. Update the fields as necessary, and then click **Apply** to save your configuration.

Clear Host Entries field descriptions

Use the data in the following table to use the Clear Host Entries tab.

Name	Description
ClearAll	Select to clear all DvR host entries.
Clearlpv4	Specifies the IPv4 address of the DvR host entries to clear.
	The IPv4 address must not be the VLAN IP address on any Controller within the DvR domain.
ClearL2Isid	Specifies the Layer 2 VSN I-SID of the DvR host entries to clear.
	The range is 0 to 16777215.
ClearL3Isid	Specifies the Layer 3 VSN I-SID of the DvR host entries to clear.
	The range is 0 to 16777215.

View Layer 3 VSN Information

About this task

View VRFs corresponding to Layer 3 (routed) VSN I-SIDs on either a Controller or a Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand the **Configuration > Edit** folders.
- 2. Click DVR.
- 3. Click the L3-VSN tab.

Click **Filter** to filter rows based on specific filter criteria.

L3-VSN field descriptions

Use the data in the following table to use the L3-VSN tab.

Name	Description
Vrfld	Specifies the VRF ID of the VRF corresponding to the Layer 3 VSN I-SID.

Name	Description
Isid	Specifies the Layer 3 VSN I-SID.
VrfName	Specifies the VRF name of the VRF corresponding to the Layer 3 VSN I-SID.
InjectDefaultRouteDisable	Specifies whether injection of default routes is disabled.

View the DvR Database

About this task

View all DvR routes on a Controller or a Leaf node.

The Controller node displays all the IP subnet routes configured for that DvR domain. A Leaf node displays all IP subnet routes learned from the Controller(s) for the Layer 2 VSNs in the DvR Domain. It also displays the Host Routes (ARPs) learned from other DvR enabled nodes.

Before you begin

Ensure that you enable DvR on the node.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **Edit** folders.
- 2. Click DVR.
- 3. Click the **Database** tab.
- 4. (Optional) To filter the rows based on the specific criteria, click Filter.

Database field descriptions

Use the data in the following table to use the Database tab.

Name	Description
DestlpAddrType	Specifies the address type of the IPv4 destination address of the DvR database entry.
DestlpAddr	Specifies the IPv4 destination address of the DvR database entry.
DestMask	Specifies the destination mask of the DvR database entry.
L3Isid	Specifies the Layer 3 I-SID of the DvR database entry.
EcmpIndex	Specifies the ECMP index for the DvR database entry.
NextHop	Specifies the MAC address of the next hop BEB, in the DvR database entry.

Name	Description
L2Isid	Specifies the Layer 2 I-SID of the DvR database entry.
Vrfld	Specifies the VRF ID for the DvR database entry.
OutgoingInterface	Specifies the outgoing interface (port or MLT) of the DvR database entry.
SpbCost	Specifies the SPB cost of the DvR database entry.
PrefixCost	Specifies the prefix cost of the DvR database entry.
NextHopName	Specifies the host name of the next hop BEB, in the DvR database table entry.
Age	Specifies the uptime since creation of the DvR database table entry.

View DvR Backbone Entries on a Controller

About this task

View the DvR backbone entries (redistributed host routes) learned from all Controllers in all DvR domains.



You can view DvR backbone entries only on a Controller. Viewing backbone entries does not apply to a Leaf node.

Before you begin

Ensure that you enable DvR globally on the node.

Procedure

- 1. In the navigation pane, expand **Configuration > Edit** folders.
- 2. Select DVR.
- 3. Select the **Backbone Entries** tab.
- 4. (Optional) To filter the rows based on the specific criteria, click Filter.

Backbone Entries field descriptions

Use the data in the following table to use the Backbone Entries tab.

Name	Description
IpAddrType	Specifies the address type of the DvR backbone host (IPv4 remote ARP).
IpAddr	Specifies the IPv4 address of the DvR backbone host.

Name	Description
L3Isid	Specifies the L3 I-SID of the DvR backbone host.
DomainId	Specifies the domain ID of the DvR backbone host.
Ecmpindex	Specifies the ECMP index of the DvR backbone host.
HostMacAddr	Specifies the MAC address of DvR backbone host.
L2Isid	Specifies the L2 I-SID of the DvR backbone host.
AdvControllerName	Specifies the host name of the advertising Controller.
AdvController	Specifies the host MAC address of the advertising Controller.
NextHopName	Specifies the host name of the next hop Backbone host in the DvR route.
NextHopMac	Specifies the MAC address of the next hop Backbone host in the DvR route.

Chapter 5: Address Resolution Protocol

Table 24: Address Resolution Protocol product support

Feature	Product	Release introduced	
For configuration details, see Confi	For configuration details, see Configuring IPv4 Routing for VOSS.		
Address Resolution Protocol (ARP) including Proxy ARP and Static ARP	VSP 4450 Series	VSP 4000 4.0	
	VSP 4900 Series	VOSS 8.1	
Static ARF	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VSP 8200 4.0	
	VSP 8400 Series Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 4.5	
	XA1400 Series	VOSS 8.0.50	
Gratuitous ARP filtering	VSP 4450 Series	VOSS 4.2	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VOSS 4.2	
	VSP 8400 Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 4.5	
	XA1400 Series	VOSS 8.0.50	

Address Resolution Protocol

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host IP address, the network station uses Address Resolution Protocol (ARP) to determine the physical address for a network host by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts.

The network station uses ARP to determine the host physical address as follows:

• The network station broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.

- All network hosts receive the broadcast request.
- Only the specified host responds with its hardware address.
- The network station then maps the host IP address to its physical address and saves the results in an address-resolution cache for future use.
- The network station ARP table displays the associations of the known MAC address to IP address.

You can create ARP entries, and you can delete individual ARP entries.

Enable ARP traffic

The switch accepts and processes ARP traffic, spanning tree bridge packet data units (BPDU), and Topology Discovery Protocol packets on port-based VLANs with the default port action of drop. If a filter port action is drop for a packet, ARP packets are also dropped. As a result, ARP entries on that port are cleared and are not relearned when the ARP aging timer expires.

To prevent dropped ARP packets, configure the following options:

- A user-defined protocol-based VLAN for ARP EtherType (byprotocol usrDefined 0x0806).
- Ports as static members to this VLAN with the default port action of drop.
- The port default VLAN ID to the correct port-based VLAN where the ARPs are processed.

You do not need to make configuration changes for the BPDU and Topology Discovery Protocol packets.

Only one user-defined protocol-based VLAN for ARP is allowed for each Spanning Tree Group (STG). If the ports with the default port action of drop are in different STGs, you must create additional user-defined protocol-based VLANs.

Proxy ARP

A network station uses proxy ARP to respond to an ARP request from a locally attached host or end station for a remote destination. The network station sends an ARP response back to the local host with its own MAC address of the network station interface for the subnet on which the ARP request was received. The reply is generated only if the device has an active route to the destination network.

The following figure shows an example of proxy ARP operation. In this example, host C with mask 24 appears to be locally attached to host B with mask 16, so host B sends an ARP request for host C. However, the switch is between the two hosts. To enable communication between the two hosts, the switch responds to the ARP request with the IP address of host C but with its own MAC address.

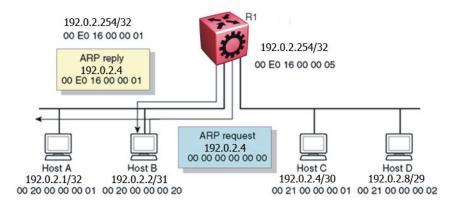


Figure 11: Proxy ARP operation

Loop detection

To prevent cases of ARP looping, configure the ARP loop detection flag to detect this situation. When a loop is detected, the port is shut down.

Flushing router tables

For administrative or troubleshooting purposes, sometimes you must flush the routing tables. Flush routing tables either by VLAN or by port. In a VLAN context, all entries associated with the VLAN are flushed. In a port context, all entries associated with the port are flushed.

Reverse Address Resolution Protocol

Certain devices use the Reverse Address Resolution Protocol (RARP) to obtain an IP address from a RARP server. MAC address information for the port is broadcast on all ports associated with an IP protocol-based or port-based VLAN. To enable a device to request an IP address from a RARP server outside its IP VLAN, you must create a RARP protocol-based VLAN.

RARP has the format of an ARP frame but its own Ethernet type (8035). You can remove RARP from the IP protocol-based VLAN definition and treat it as a separate protocol, thus creating a RARP protocol-based VLAN.

A typical network topology provides desktop switches in wiring closets with one or more trunk ports that extend to one or more data center switches where attached servers provide file, print, and other services. Use RARP functionality to define all ports in a network that require access to a RARP server as potential members of a RARP protocol-based VLAN. You must define all tagged ports and data center RARP servers as static or permanent members of the RARP VLAN. Therefore, a desktop host broadcasts an RARP request to all other members of the RARP VLAN. In normal operation, these members include only the requesting port, tagged ports, and data center RARP server ports. Because all other ports are potential members of this VLAN and RARP is only transmitted at startup, all other port VLAN memberships expire. With this feature, one or more centrally located RARP servers extend RARP services across traditional VLAN boundaries to reach desktops globally.

ARP configuration using the CLI

Network stations that use IP protocol require both a physical address and an IP address to transmit packets. In situations where the station knows only the network host IP address, the Address Resolution Protocol (ARP) lets you use the network station to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address.

A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

ARP response is enabled by default.

Enabling ARP on a port or a VLAN

Enable ARP on the device so that it answers local ARP requests.

About this task

You can enable or disable ARP responses on the device. You can also enable ARP proxy, which lets a router answer a local ARP request for a remote destination.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```

Note:

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

2. Enable ARP on the device:

```
ip arp-response
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #interface vlan 200
Switch:1(config-if) #ip arp-response
```

Enabling ARP proxy

About this task

Configure an ARP proxy to allow the platform to answer a local ARP request for a remote destination. ARP proxy is disabled by default.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```

Note:

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

2. Enable ARP proxy on the device:

```
ip arp-proxy enable
```

Use the no operator to disable ARP proxy: no ip arp-proxy [enable]

Example

Enable ARP proxy on VLAN 200:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #interface vlan 200
Switch:1(config-if) #ip arp-proxy enable
```

View ARP Information

The **show** ip **arp** command displays all of the configured and dynamically learned ARP entries in the ARP table.

Procedure

- 1. To enter User EXEC mode, log on to the switch.
- 2. Display ARP information for a specified port or for all ports:

```
show ip arp interface gigabitethernet [slot/port[/sub-port][-slot/
port[/sub-port]][,...]
```

3. Display ARP information for a VLAN:

```
show ip arp interface vlan <1-4059>
```

Example

Switch:1>show ip arp interface			
		Pc	ort Arp
PORT_NUM	DOPROXY	DORESP	
1/2 1/3 1/4 1/5 1/6 1/7 1/8 1/9 1/10 1/11 1/12 1/13 1/14 1/15 1/16	false	true true true true true true true true	
	<pre>false (q = quit)</pre>	true	

Variable definitions

Use the data in the following table to use the **show ip arp** command.

Variable	Value
A.B.C.D	Specifies the IP address of a network.
{slot/port[/sub-port][-slot/port[/sub-port]][,]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
interface	Displays ARP interface configuration information.
spbm-tunnel-as-mac	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
-S	Specifies a subnet.
	You must indicate the IP address followed by the subnet mask expressed as <a.b.c.d> <a.b.c.d>.</a.b.c.d></a.b.c.d>
vlan <1-4059>	Displays ARP entries for a particular VLAN ID.
	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and

Variable	Value
	spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
vrf WORD<1–16>	Specifies a VRF name expressed as text from 1 to 16 characters in length.
	The total number of ARPs listed in the summary line of the show ip arp output represents the total number of ARPs on the chassis, including all VRFs (which includes the Management Router VRF).
vrfids WORD<0-512>	Specifies a range of VRFIDs as text from 0 to 512 characters in length.
	The total number of ARPs listed in the summary line of the show ip arp output represents the total number of ARPs on the chassis, including all VRFs (which includes the Management Router VRF).

Use the data in the following table to help you understand the **show ip arp interface** command output.

Variable	Value
PORT_NUM	Indicates the port number.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.

Use the data in the following table to help you understand the **show ip arp interface vlan** command output.

Variable	Value
VLAN_ID	Indicates the VLAN ID.
DOPROXY	Indicates if ARP proxy responses are enabled or disabled on the specified interface.
DORESP	Indicates if the sending of ARP responses is enabled or disabled on the specified interface.

Configuring IP ARP static entries

About this task

Configure ARP static entries to modify the ARP parameters on the device. The only way to change a static ARP is to delete the static ARP entry and create a new entry with new information.



Note:

Static multicast ARP entries are not supported for NLB Unicast or NLB Multicast operations.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Configure ARP static entries on the device:

```
ip arp <A.B.C.D> 0x00:0x00:0x00:0x00:0x00:0x00 {slot/port[-slot/
port][,...]}
```

Example

Configure ARP static entries:

```
Switch:1>enable
Switch: 1#configure terminal
Switch:1(config) #ip arp 192.0.2.10 00-16-76-7D-80-C2 2/1
```

Variable definitions

Use the data in the following table to use the ip arp command.

Table 25: Variable definitions

Variable	Value
request-threshold <50-1000>	Configures the maximum number of outstanding ARP requests that a device can generate. The range is 50–1000. The default value is 500.
	To configure this option to the default value, use the default operator with this command.
timeout <1-32767>	Configures the length of time in seconds an entry remains in the ARP table before timeout. The range is 1–32767.
	To configure this option to the default value, use the default operator with this command.
	Note:
	The aging of ARP records is tied to the aging of MAC records. The ARP record for a given IP address is not removed unless the associated MAC record ages out and the router stops receiving a response to ARP requests for that IP address. In cases where the ARP aging time is set to less than the MAC aging time, the switch waits until the MAC ages out before deleting the ARP for an inactive host.
<a.b.c.d></a.b.c.d>	Adds ARP entries.

Clearing ARP entries

Use this procedure to clear dynamic ARP table entries associated with the interface or VLAN.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Clear ARP entries:

```
clear ip arp interface <gigabitethernet|vlan> <slot/port[/sub-port]
[-slot/port[/sub-port]][,...]| <1-4059>>
```

Example

Clear ARP entries:

```
Switch:1> enable
Switch:1# clear ip arp interface gigabitethernet 1/16
```

Variable definitions

Use the data in the following table to use the clear ip arp interface command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
gigabitethernet vlan	Specifies the interface type.
{slot/port[/sub-port][-slot/port[/sub-port]][,]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Showing ARP table information

Show ARP information to view the configuration information in the ARP table.

About this task

When you use the interface parameter with the **show ip arp** command you can display ARP configuration information only for a specific switch.

The **show** ip **arp** command displays all of the configured and dynamically learned ARP entries in the ARP table.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display the ARP table:

show ip arp [$\langle A.B.C.D \rangle$] [$-s \langle A.B.C.D \rangle$] [gigabitEthernet $\langle slot/port[/sub-port] \rangle$] [interface $\langle gigabitethernet|vlan \rangle$] [nlb] [spbm-tunnel-as-mac][$vlan \langle 1-4059 \rangle$] [$vrf WORD \langle 1-16 \rangle$] [$vrfids WORD \langle 0-512 \rangle$]

Example

```
Switch: 1#show ip arp
                       IP Arp - GlobalRouter
______
                                  TYPE TTL(10 Sec) TUNNEL
IP_ADDRESS MAC_ADDRESS VLAN PORT
192.0.2.1 00:09:0f:09:00:08 20 1/3 DYNAMIC
192.0.2.12 b4:a9:5a:ff:f8:40 20 1/3
                                        DYNAMIC
192.0.2.25 e4:5d:52:3c:65:00 20
                                        LOCAL
2160
192.0.2.154 d4:ea:0e:c2:08:00 20 1/3
                                        DYNAMIC
2131
192.0.2.157 00:1c:17:b1:ec:80 20 1/3
                                        DYNAMIC
2131
192.0.2.161 fc:a8:41:fb:40:00 20 1/3
                                        DYNAMIC
2131
192.0.2.253 e0:db:55:d4:e5:7c 20 1/3
                                        DYNAMIC
192.0.2.255 ff:ff:ff:ff:ff: 20
                                         LOCAL
2160
______
                  IP Arp Extn - GlobalRouter
MULTICAST-MAC-FLOODING AGING (Minutes)
                               ARP-THRESHOLD
                                500
c: customer vid u: untagged-traffic
8 out of 8 ARP entries displayed
ARPs on TX-NNI: Current = 0, re-ARP count = 0
```

Variable definitions

Use the data in the following table to help you use the show ip arp command.

Variable	Value
-s	Specifies the subnet for the table.
gigabitEthernet	Displays the entries for a particular brouter port.
interface	Displays ARP interface configuration information.
	Use the following parameters to display ARP table information specifically for:
	gigabitethernet {slot/port[-slot/port][,]} displays IP ARP gigabitethernet interface information
	VLAN <1-4059> displays IP ARP VLAN interface information
	Example: show ip arp interface vlan 1
nlb	Displays the Network Load Balancing (NLB) ARP entries on the switch.
spbm-tunnel-as-mac	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
vlan	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
	Use these parameters to display ARP table information specifically for:
	vrf WORD<1–16>—the VLAN VRF name in a range from 1 to 16 characters
	vrfids WORD<0-512>—the VLAN VRF ID in a range from 0 to 512
	Example: show ip arp vlan 1 vrf 1
vrf WORD <1-16>	Specifies the name of the VRF.
	The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis including all VRFs.
vrfids WORD <0-512>	Specifies the VRF ID.
	The total number of ARPs listed in the summary line of the "show ip arp" display represents the total number of ARPs on the chassis including all VRFs.
<a.b.c.d></a.b.c.d>	Specifies the network IP address for the table.

Use the data in the following table to help you understand the output of the ${\tt show}$ ip ${\tt arp}$ command.

Parameter	Description
IP_ADDRESS	Indicates the IP address where ARP is configured.
MAC_ADDRESS	Indicates the MAC address where ARP is configured.
VLAN	Indicates the VLAN address where ARP is configured.
PORT	Indicates the port where ARP is configured.
TYPE	Indicates the type of learning (dynamic or local) where ARP is configured.
TTL<10 secs>	Indicates the time to live as tenths of a second where ARP is configured.
TUNNEL	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
MULTICAST-MAC- FLOODING	Displays whether IP ARP multicast MAC flooding is enabled or disabled. When enabled, the ARP entries for multicast MAC addresses are associated with the VLAN or port interface on which they were learned.
AGING (Minutes)	Displays when the ARP aging timer expires.
ARP-THRESHOLD	Displays the maximum number of outstanding ARP requests that a device can generate.

Configuring Gratuitous ARP

Use the following procedure to configure Gratuitous Address Resolution Protocol (ARP). When Gratuitous ARP is enabled the switch allows all Gratuitous ARP request packets. The default is enabled.

If you disable Gratuitous ARP, the switch only allows Gratuitous ARP packets associated with Routed Split Multi-Link Trunking (RSMLT) or Virtual Router Redundancy Protocol (VRRP), and the switch discards all other Gratuitous ARP request packets.

About this task

ARP translates network layer (layer 3) IP addresses into link layer (layer 2) MAC addresses. A host sends a Gratuitous ARP request packet to inform other hosts of the existence of an interface on the network, so other local hosts can update their ARP tables. If the IP or MAC address changes, or in the event of a failover, a host sends a Gratuitous ARP request packet to inform other hosts to update their ARP tables.

VRRP and RSMLT use gratuitous ARP to update the MAC address tables on switches.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable Gratuitous ARP:

```
ip gratuitous-arp
```

3. **(Optional)** Disable Gratuitous ARP:

no ip gratuitous-arp

4. (Optional) Configure Gratuitous ARP to the default value:

default ip gratuitous-arp

5. Save the changed configuration.

save config [backup WORD<1-99>][file WORD<1-99>][verbose]

ARP configuration using Enterprise Device Manager

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host IP address, the network station can use Address Resolution Protocol (ARP) to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

Enabling or disabling ARP on the brouter port or a VRF instance

About this task

After you assign the IP address, you can configure ARP. By default, ARP Response is enabled and Proxy ARP is disabled.

Procedure

- 1. In the Device Physical View tab, select a port.
- 2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.
- 3. Click IP.
- 4. Click the ARP tab.
- 5. In the **DoProxy** check box, select **enable** to enable the Proxy ARP function.
- 6. In the **DoResp** check box, select **enable** to configure the system to respond to an ARP. The default is enable.
- 7. Click Apply.

The ARP function is available only when the port or VLAN is routed; that is, it is assigned an IP address.

ARP field descriptions

Use the data in the following table to use the **ARP** tab fields.

Name	Description
DoProxy	Configures the system to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.

Enabling or disabling ARP on a VLAN or a VRF instance

About this task

Use the following procedure to enable ARP on VLAN level.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANs > Basic.
- 3. Select a VLAN.
- 4. Click IP.
- 5. Click the **ARP** tab.
- 6. In the **DoProxy** field, click **enable** to enable the Proxy ARP function.
- 7. In the **DoResp** field, click **enable** to configure the system to respond to an ARP. The default is enable.
- 8. Click Apply.

The ARP dialog box is available only if the port or VLAN is routed; that is, it is assigned an IP address.

ARP field descriptions

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

Name	Description
DoProxy	Configures the system to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Configures the system to send ARP responses for this IP interface address. The default value is enable.

Viewing and managing ARP

About this task

You can view and manage known MAC address to IP address associations. In addition, you can create or delete individual ARP entries.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the ARP tab.

ARP field descriptions

Use the data in the following table to use the ARP tab.

Name	Description
NetAddress	Specifies the IP address corresponding to the media-dependent physical address.
IfIndex	Identifies the router interface for this ARP entry:
	Brouter interfaces are identified by the slot/port number of the brouter port.
	VLAN interfaces are identified by the vlan name.
PhysAddress	Specifies the media-dependent physical address (that is, the Ethernet address).
Туре	Specifies the type of ARP entry:
	local—a locally configured ARP entry
	static—a statically configured ARP entry
	dynamic—a learned ARP entry
TimeToLive	Indicates the time to live where the ARP is configured.
Destifindex	Indicates the slot/port on which the ARP entry was learned. For brouter interfaces this is the same value as IfIndex, but for VLAN interfaces, it designates the particular port in the VLAN on which the ARP was learned.
DestVlanId	VLAN ID where the ARP is configured.
ВМас	Identifies the backbone MAC address if the entry is learned from an SPBM network.
DestCvid	Identifies the customer VLAN ID for a Switched UNI port.

Creating static ARP entries

About this task

Use the following procedure to create a static ARP entry.



Note:

Static multicast ARP entries are not supported for NLB Unicast or NLB Multicast operations.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the ARP tab.

- 4. Click Insert.
- 5. In the **NetAddress** field, type the IP address.
- 6. Click Port.

OR

Click Port in VLAN

- 7. In the dialog box, select the interface.
- 8. Click OK.
- 9. In the **PhysAddress** field, type the MAC address.
- 10. Click Insert.

Configuring ARP proxy

About this task

With an ARP proxy, the switch can respond to an ARP request from a locally attached host or end station for a remote destination. Proxy ARP does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet on which the ARP request was received. The reply is generated only if the system has an active route to the destination network.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANs > Basic.
- 3. Choose a VLAN.
- 4. Click IP.
- 5. Click ARP tab.
- 6. Select **DoProxy enable**.
- 7. Click Apply.

Chapter 6: Dynamic Host Configuration Protocol and User Datagram Protocol Configuration

Table 26: Dynamic Host Configuration Protocol product support

Feature	Product	Release introduced
For configuration details, see Configuring IPv4 Routing for VOSS.		
Dynamic Host Configuration	VSP 4450 Series	VSP 4000 4.0
Protocol (DHCP) Relay	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50
DHCP Option 82	VSP 4450 Series	VSP 4000 4.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 4.2.1
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VSP 8200 4.0
	VSP 8400 Series	VOSS 4.2
	VSP 8600 Series	VSP 8600 4.5
	XA1400 Series	VOSS 8.0.50

DHCP option 82

The DHCP option 82 is the DHCP Relay Agent Information option. The DHCP relay agent inserts option 82 when it forwards the client-originated DHCP packets to a DHCP server. The Relay Agent Information option is organized as a single DHCP option that contains one or more sub-options that convey information known by the relay agent. The DHCP server echoes the option back to the relay

agent in server-to-client replies, and the relay agent removes the option before forwarding the reply to the client.

The DHCP option 82 is added at the DHCP relay level as shown in the following image.

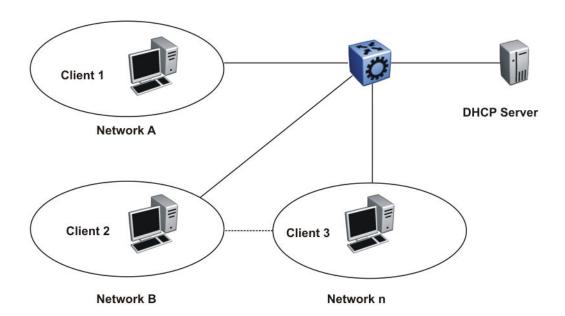


Figure 12: DHCP Client-Relay-Server Architecture

The Relay Agent Information option (code 82) is a container for specific agent-supplied suboptions; Agent Circuit ID (code 1) and Agent Remote ID (code 2). The suboptions can represent different information relevant for the relay. The fields are encoded in the following manner, where N or n is the total number of octets in the Agent Information Field (all bytes of the suboptions):

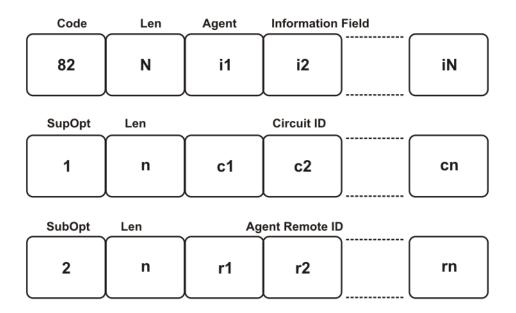


Figure 13: Format of the Relay Agent Information

Because at least one of the sub-options must be defined, the minimum Relay Agent Information length is two (2), and the length n of the suboption can be zero (0). The sub-options do not have to appear in any particular order. No pad suboption is defined and the Information field is not terminated with 255 suboption.

DHCP Suboptions

The suboptions are Agent Circuit ID and Agent Remote ID.

The DHCP relay agents can add the Agent Circuit ID to terminate switched or permanent circuits. The Agent Circuit ID encodes an agent-local identifier of the circuit from which a DHCP client-to-server packet was received. Agents can use the Circuit ID to relay DHCP responses back to the proper circuit. In the switch, the Agent Circuit ID field contains the ifIndex of the interface on which the packet is received.

DHCP relay agents can add the Agent Remote ID to terminate switched or permanent circuits, and can identify the remote host end of the circuit. The switch uses the Agent Remote ID field to encode the MAC address of the interface on which the packet is received. The Agent Remote ID must be globally unique.

Agent Operations

A DHCP relay agent adds a Relay Agent Information field as the last option in the DHCP options field of any recognized BOOTP or DHCP packet forwarded from a client to a server. However, if the End Option 255 is present, then the DHCP relay agent adds a Relay Agent information field before the End Option 255 field.

Relay agents can receive a DHCP packet from an untrusted circuit with the gateway IP address (GIADDR) set to zero to indicate that the relay agent is the first-hop router from the gateway. If a Relay Agent Information option is present in the packet, the relay agent discards the packet and increments an error counter. A trusted circuit can contain a trusted downstream network element, for example, a bridge, between the relay agent and the client. The bridge can add a relay agent option but does not set the GIADDR field. In this case, the relay agent forwards the DHCP packet per normal DHCP relay agent operations, and sets the GIADDR field to the relay address. The relay agent does not add a second relay agent option.

You can distinguish between a trusted circuit and an untrusted circuit based on the type of circuit termination equipment you use. To make a circuit trusted, set the trusted flag under DHCP for each interface.

After packets append the Relay Agent Information option, the packets that exceed the MTU or the vendor size buffer of 64 bits, are forwarded without adding the Agent Information option, and an error counter is incremented.

The relay agent or the trusted downstream network element removes the Relay Agent Information option echoed by a server that is added when forwarding a server-to-client response back to the client.

The following list outlines the operations that the relay agent does not perform:

- The relay agent does not add an Option Overload option to the packet or use the file or sname fields to add the Relay Agent Information option. The agent does not parse or remove Relay Agent Information options that can appear in the sname or file fields of a server-to-client packet forwarded through the agent.
- The relay agent does not monitor or modify client-originated DHCP packets addressed to a server unicast address; this includes the DHCP-REQUEST sent when entering the RENEWING state.
- The relay agent does not modify DHCP packets that use the IPSEC Authentication Header or IPSEC Encapsulating Security Payload.

A DHCP relay agent can receive a client DHCP packet forwarded from a BOOTP/DHCP relay agent closer to the client. This packet has a GIADDR as non-zero, and may or may not already have a DHCP Relay Agent option in it.

Relay agents configured to add a Relay Agent option which receive a client DHCP packet with a nonzero GIADDR, discards the packet if the GIADDR spoofs a GIADDR address implemented by the local agent itself. Otherwise, the relay agent forwards any received DHCP packet with a valid non-zero GIADDR without adding any relay agent options. The GIADDR value does not change.

UDP broadcast forwarding

Some network applications, such as the NetBIOS name service, rely on a User Datagram Protocol (UDP) broadcast to request a service or locate a server for an application. If a host is on a network, subnet segment, or VLAN that does not include a server for the service, UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. You can resolve this problem by forwarding the broadcasts to the server through physical or virtual router interfaces.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface out to other router IP interfaces as a rebroadcast or to a configured IP address. If the address is that of a server, the packet is sent as a unicast packet to this address. If the address is that of an interface on the router, the frame is rebroadcast.

After a UDP broadcast is received on a router interface, it must meet the following criteria to be eligible for forwarding:

- It must be a MAC-level broadcast.
- · It must be an IP limited broadcast.
- It must be for the specified UDP protocol.
- It must have a time-to-live (TTL) value of at least 2.

For each ingress interface and protocol, the policy specifies how the UDP broadcast is retransmitted: to a unicast host address or to a broadcast address.

DHCP and UDP configuration using the CLI

Use Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), to provide host configuration information to the workstations dynamically. Use the DHCP relay commands to configure DHCP relay behavior on a port or on a VLAN.

This section describes CLI commands for DHCP and User Datagram Protocol (UDP) configuration.

Configure DHCP Parameters Globally

Before you begin

Configure an IP address on the interface to be used as the DHCP relay interface.

About this task

Configure DHCP relay parameters for the port or the VLAN.

Procedure

1. Enter Global Configuration mode:

enable

```
configure terminal
```

2. Create the forwarding path from the client to the server:

```
ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D>
```

3. Enable the forwarding path from the client to the server:

```
ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> enable
```



If the agent IP address (the first <A.B.C.D> variable) is a VLAN or port IP address, you must enable DHCP Relay on that VLAN or port by running ip dhcp-relay within the VLAN context. However, if the first <A.B.C.D> variable is a VRRP address, you do not need to enable DHCP Relay on the VLAN or port in which the VRRP address resides.

4. Modify DHCP mode to forward BootP messages only, DHCP messages only, or both:

```
ip dhcp-relay fwd-path <A.B.C.D> <A.B.C.D> mode <bootp|bootp_dhcp|
dhcp>
```

Example

Create the forwarding path from the client to the server. Enable the forwarding path from the client the server. Modify DHCP mode to forward both BootP and DHCP messages.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config) #ip dhcp-relay fwd-path 192.0.2.120 192.0.2.50
Switch:1(config) #ip dhcp-relay fwd-path 192.0.2.128 192.0.2.50 enable
Switch:1(config) #ip dhcp-relay fwd-path 192.0.2.128 192.0.2.50 mode bootp_dhcp
```

Variable Definitions

The following table defines parameters for the ip dhcp-relay fwd-path command.

Table 27: Variable definitions

Variable	Value
fwd-path <a.b.c.d> <a.b.c.d></a.b.c.d></a.b.c.d>	Configures the forwarding path from the client to the server.
	The first <a.b.c.d> variable is the agent IP address configured on an interface (a locally configured IP address).</a.b.c.d>
	The second < <i>A.B.C.D</i> > variable is the IP address of the DHCP server in the network. If this IP address corresponds to the locally configured IP network, the DHCP packet is broadcast out from the interface.
fwd-path <a.b.c.d> <a.b.c.d> disable</a.b.c.d></a.b.c.d>	Disables DHCP relaying on the path from the IP address to the server. This feature is disabled by default.
	The first <a.b.c.d> variable is the agent IP address configured on an interface (a locally configured IP address).</a.b.c.d>

Variable	Value
	The second < <i>A.B.C.D</i> > variable is the IP address of the DHCP server in the network.
fwd-path <a.b.c.d> <a.b.c.d> enable</a.b.c.d></a.b.c.d>	Enables DHCP relaying on the path from the IP address to the server.
	The first < <i>A.B.C.D</i> > variable is the agent IP address configured on an interface (a locally configured IP address).
	Note:
	If the agent IP address (the first <a.b.c.d> variable) is a VLAN or port IP address, you must enable DHCP Relay on that VLAN or port by running ip dhcp-relay within the VLAN context. However, if the first <a.b.c.d> variable is a VRRP address, you do not need to enable DHCP Relay on the VLAN or port in which the VRRP address resides.</a.b.c.d></a.b.c.d>
	The second < <i>A.B.C.D</i> > variable is the IP address of the DHCP server in the network. If this IP address corresponds to the locally configured IP network, the DHCP packet is broadcast out from the interface.
fwd-path <a.b.c.d> <a.b.c.d> mode <bootp bootp_dhcp dhcp></bootp bootp_dhcp dhcp></a.b.c.d></a.b.c.d>	Modifies DHCP mode to forward BootP messages only, DHCP messages only, or both. The default is both.
	The mode is {bootp bootp_dhcp dhcp}.

Showing DHCP relay information

Display relay information to show relay information about DHCP routes and counters.

For scaling information on DHCP Relay forwarding (IPv4 or IPv6), see Release Notes for VOSS.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display information about DHCP relay forward paths:

```
show ip dhcp-relay fwd-path [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

3. Display information about DHCP relay counters:

```
show ip dhcp-relay counters [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

4. Display the options for each listed interface:

```
show ip dhcp-relay interface [gigabitethernet \{\text{slot/port}[/\text{sub-port}] = (-\text{slot/port}[/\text{sub-port}]] = (-\text{slot/port}[/\text{sub-port}]] = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) = (-16) =
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1#show ip dhcp-relay interface

Port Dhcp

Port VRF
MAX MIN
NAME
ENABLE HOP SEC MODE

Vlan Dhcp

Vlan Dhcp

Vlan VRF
MAX MIN
ALWAYS CIRCUIT REMOTE TRUST
BCAST ID

Vlan VRF
DNAME
ENABLE HOP SEC MODE

Vlan Dhcp

Vlan VRF
DNAME
ENABLE HOP SEC MODE

ALWAYS CIRCUIT REMOTE TRUST
BCAST ID

O CIRC

Always CIRCUIT REMOTE TRUST
BCAST ID

O CIRC

Always CIRCUIT REMOTE TRUST
BCAST ID

O CIRC

All 0 out of 0 of Vlan Dhcp Entries displayed
```

Variable definitions

Use the data in the following table to use the **show** ip **dhcp-relay** command.

Variable	Value
vrf WORD<1-16>	The name of the VRF.
vrfids WORD<0-512>	The ID of the VRF. The value is an integer in the range of 0–512.

Use the data in the following table to use the show ip dhcp-relay interface command.

Variable	Value
{slot/port[/sub-port][-slot/port[/sub-port]][,]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
[vlan <1-4059>]	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
[vrf WORD<1-16>]	Specifies the name of the VRF.

Variable	Value
[vrfids WORD<0-512>]	Specifies the ID of the VRF. The value is an integer from 0– 512.

Configuring DHCP option 82

Configure the DHCP option 82 to enable the circuit ID to encode an agent-local identifier of the circuit from which a DHCP client-to-server packet is received. Configure the DHCP option 82 to enable the remote ID to encode the MAC address of the interface on which the packet is received. By default, the DHCP option 82 is disabled.

Before you begin

• You must enable ip and dhcp-relay on the VLAN.

About this task

To configure the DHCP option 82 on a VLAN, you must enter the VLAN Interface Configuration mode.

To configure the DHCP option 82 on a brouter port, you must enter the GigabitEthernet Interface Configuration mode.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```

Note:

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

2. Enable the circuit ID:

```
ip dhcp-relay circuitID
```

3. Enable the remote ID:

```
ip dhcp-relay remoteID
```

4. Configure the circuit as trusted:

```
ip dhcp-relay trusted
```

5. Show statistics for option 82, which is the relay agent information option:

```
show ip dhcp-relay counters option82 [vrf WORD <0-16>] [vrfids WORD <0-512>]
```

Example

Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/10

Enable the circuit ID:

Switch:1(config-if) # ip dhcp-relay circuitID

Enable the remote ID:

Switch:1(config-if) # ip dhcp-relay remoteID

Configure the circuit as trusted:

Switch:1(config-if) # ip dhcp-relay trusted

Show statistics for option 82, which is the relay agent information option:

Switch:1(config-if) # show ip dhcp-relay counters option82

Variable definitions

Use the data in the following table to configure the DHCP option 82.

Table 28: Variable definitions

Variable	Value
circuitID	Enables the Circuit ID.
remoteID	Enables the Remote ID.
trusted	Sets the circuit as trusted.

Use the data in the following table to use the show ip dhcp-relay counters option82 [vrf WORD <0-16>] [vrfids WORD <0-512>] command.

Variable	Value
vrf WORD <0–16>	Displays DHCP counters for a particular VRF. WORD <0–16> specifies the VRF name.
vrfids WORD <0–512>	Displays a DHCP forward path for a particular VRF. WORD <0–512> specifies the VRF ID.

Configuring DHCP relay on a port or VLAN

You can view and configure the DHCP parameters on specific ports or on a VLAN.

Before you begin

You must configure IP on the interface.

Procedure

1. Enter Interface Configuration mode:

enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port][-slot/port[/subport]][,...]} **or** interface vlan $\langle 1-4059 \rangle$



Note:

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

2. Enable DHCP parameters on a specified port or VLAN:

```
ip dhcp-relay
```

Example

```
Switch:1> enable
Switch: 1# configure terminal
Switch:1(config) # interface gigabitethernet 1/10
```

Enable DHCP parameters on a specified port or VLAN:

```
Switch:1(config-if) # ip dhcp-relay
```

Variable definitions

Use the data in the following table to use the ip dhcp-relay command.

Use the no operator to disable DHCP parameters on specified ports: no ip dhcp-relay.



The no ip dhcp-relay command disables DHCP Relay, it does not delete the DHCP entry.

To configure this option to the default value, use the default operator with this command.

Variable	Value
broadcast	Enables the device to send the server reply as a broadcast to the end station. After you disable this variable, the device sends the server reply as a unicast to the end station. Use the no operator to disable broadcast: no ip dhcp-relay broadcast.
	To configure this option to the default value, use the default operator with this command.
circuitld	Enables Option 82 circuit ID on the interface.
clear-counters	Clears DHCP Relay counters for the interface.
fwd-path <a.b.c.d> [vrid <1-255>]</a.b.c.d>	Creates a forward path server with a virtual router ID (or VRRP ID), a mode, and a state.
	A.B.C.D is the IP address.
	vrid <1-255> is the ID of the virtual router and is an integer from 1 to 255.

Variable	Value
	Use the no operator to delete a forward path server with a specific value and virtual router ID: no ip dhcp-relay fwd-path <a.b.c.d> [vrid <1-255>]</a.b.c.d>
	To configure this option to the default value, use the default operator with this command.
fwd-path <a.b.c.d> disable [vrid <1-255>]</a.b.c.d>	Disables a forward path server with a specific value and virtual router ID.
	A.B.C.D is the IP address.
	vrid <1-255> is the ID of the virtual router (or VRRP ID) and is an integer from 1 to 255.
fwd-path <a.b.c.d> enable [vrid <1-255>]</a.b.c.d>	Enables a forward path server with a specific value and virtual router ID (or VRRP ID).
	A.B.C.D is the IP address in the form a.b.c.d.
	vrid <1-255> is the ID of the virtual router and is an integer from 1 to 255.
fwd-path <a.b.c.d> mode <bootp bootp_dhcp dhcp> [vrid <1-255>]</bootp </a.b.c.d>	Configures the forward path mode for a VLAN. This command string is available only in VLAN Interface Configuration mode.
	A.B.C.D is the IP address in the form a.b.c.d.
	mode is a choice of bootp, dhcp, or bootp_dhcp.
	vrid <1-255> is the ID of the virtual router (or VRRP ID) and is an integer from 1 to 255.
	To configure this option to the default value, use the default operator with this command.
max-hop <1-16>	Configures the maximum number of hops before a BootP/DHCP packet is discarded (1 to 16). The default is 4.
	To configure this option to the default value, use the default operator with this command.
min-sec <0-65535>	Configures the minimum seconds count for DHCP. If the secs field in the BootP/DHCP packet header is greater than this value, the device relays or forwards the packet; otherwise, the packet is dropped (0 to 65535). The default is 0 seconds.
	To configure this option to the default value, use the default operator with this command.
mode <bootp bootp_dhcp dhcp></bootp bootp_dhcp dhcp>	Configures DHCP mode to forward BootP messages only, DHCP messages only, or both. The default is both.
	To configure this option to the default value, use the default operator with this command.
remoteld	Enables Option82 remote ID on the interface.
trusted	Configures the DHCP circuit as trusted.

Configuring UDP broadcast forwarding

About this task

By default, routers do not forward broadcasts. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts. You must set up UDP broadcast forwarding on the system. Configure UDP broadcast forwarding to forward the UDP broadcasts of network applications to the required server through physical or virtual router interfaces.

Procedure

- 1. Enter protocols into a table.
- 2. Create policies (protocol/server pairs).
- 3. Assemble the policies into lists or profiles.
- 4. Apply the list to the appropriate interfaces.

Configuring UDP protocols

About this task

Configure UDP protocols to determine which UDP broadcasts are forwarded.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Configure a UDP protocol:

```
ip forward-protocol udp <1-65535> WORD<1-15>
```

3. Confirm your configuration:

```
show ip forward-protocol udp interface [vrf WORD<0-16>]|[vrfids WORD<0-512>] portfwd [vrf WORD<0-16>]| [vrfids WORD<0-512>] portfwdlist <1-1000>[vrf WORD<0-16>]|[vrfids WORD<0-512>] vrf WORD<0-16> vrfids WORD<0-512>
```

Example

Configure a UDP protocol and confirm your configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip forward-protocol udp 53 DNS
Switch:1(config)#show ip forward-protocol udp
```

```
Udp Protocol Tbl - GlobalRouter

UDP_PORT PROTOCOL_NAME

37    Time Service
49    TACACS Service
53    DNS
69    TFTP
137    NetBIOS NameSrv
138    NetBIOS DataSrv
```

Variable definitions

Use the data in the following table to use the ip forward-protocol udp command.

Variable	Value
<1-65535> WORD<1-15>	Creates a new UDP protocol.
	<1-65535> WORD<1-15> is the UDP protocol name as a string.
	Use the no operator to delete a UDP protocol no ip forward-protocol udp <1-65535>.
portfwd	Displays portfwd information.
portfwdlist	Displays port forward list information.
vrf WORD<0-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF.

Configuring a UDP port forward entry

Configure a UDP port forward entry to add or remove a port forward entry.

Procedure

1. Enter either Global Configuration mode or VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
Optional: router vrf WORD<1-16>
```

2. Configure a UDP port forward entry:

```
ip forward-protocol udp portfwd <1-65535> {A.B.C.D}
```

3. Confirm your configuration:

```
show ip forward-protocol udp portfwd [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

Configure a UDP port forward entry:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip forward-protocol udp portfwd 150 192.0.2.10
```

Variable definitions

Use the data in the following table to use the ip forward-protocol udp portfwd command.

Variable	Value
<1-65535> {A.B.C.D}	Adds a UDP protocol port to the specified port forwarding list.
	1-65535 is a UDP protocol port in the range of 1–65535.
	A.B.C.D is an IP address in a.b.c.d format.
	Use the no operator to remove a protocol port forwarding entry and IP address from the list: no ip forward-protocol udp portfwd <1-65535> <a.b.c.d>.</a.b.c.d>
	To configure this option to the default value, use the default operator with this command.
vrf WORD<0-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF.

Configuring the UDP port forwarding list

Configure the UDP port forwarding list to assign protocols and servers to the port forward list.

About this task

You can perform this procedure in Global Configuration mode, VLAN Interface Configuration mode, or VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the UDP port forwarding list:

```
ip forward-protocol udp portfwdlist <1-1000>
```

! Important:

The following two steps are not available in the Global Configuration or VRF Router Configuration mode. The following two commands are available in VLAN Interface Configuration mode only.

3. Enter VLAN Interface Configuration mode:

interface vlan <1-4059>

4. Configure the broadcast mask:

ip forward-protocol udp broadcastmask {A.B.C.D}

5. Configure the maximum time to live:

ip forward-protocol udp maxttl <1-16>

6. Confirm your configuration:

show ip forward-protocol udp portfwdlist <1-1000> [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

Switch:1> enable Switch:1# configure terminal

Configure the UDP port forwarding list:

Switch:1(config) # ip forward-protocol udp portfwdlist 1

Log on to the VLAN interface configuration mode:

Switch:1(config) # interface vlan 3

Configure the broadcast mask:

Switch:1(config-if) # ip forward-protocol udp broadcastmask 192.0.2.255

Configure the maximum time to live:

Switch:1(config-if)# ip forward-protocol udp maxttl 10

Confirm the configuration:

Switch:1(config-if) # show ip forward-protocol udp portfwdlist

Variable definitions

Use the data in the following table to use the ip forward-protocol udp portfwdlist command.

Variable	Value
<1-1000>	Creates a UDP port forwarding list in the range of 1–1000.
<1–65535> {A.B.C.D}	Adds a UDP protocol port to the specified port forwarding list.
	1-65535 is a UDP protocol port in the range of 1–65535.
	A.B.C.D is an IP address in a.b.c.d format.
	Use the no operator to remove or delete a port forwarding list ID,

Variable	Value
	no ip forward-protocol udp portfwdlist <1-1000> <1-65535> <a.b.c.d>.</a.b.c.d>
	To configure this option to use the default value, use the default operator with this command.
name WORD<0-15>	Changes the name of the port forwarding list.

Use the data in the following table to use the ip forward-protocol udp command.

Variable	Value	
broadcastmask {A.B.C.D}	Configures the interface broadcast mask (the interface broadcast mask can be different from the interface mask).	
	A.B.C.D is an IP address in a.b.c.d format.	
	Use the no operator to delete the broadcast mask:	
	no ip forward-protocol udp broadcastmask {A.B.C.D}	
	To configure this option to the default value, use the default operator with this command.	
maxttl <1-16>	Configures the maximum time-to-live value (TTL) for the UDP broadcast forwarded by the interface. The range is 1–16.	
portfwdlist <1–1000>	Assigns the list to the VLAN.	
vlan <1-4059> [portfwdlist <1– 1000>]	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.	
	If you use the portfwdlist variable with the VLAN variable, it assigns the list to the specified VLAN, regardless of which VLAN context you currently configure.	

Showing UDP forward information

Show UDP forward information to view information about the UDP forwarding characteristics of the device. UDP forwarding only supports 128 entries.

About this task

There are four show options:

- · Show the interface information
- Show the port forward information
- Show the port forward list information
- Show the protocol information

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display information about the UDP interface for all IP addresses or a specified IP address:

```
show ip forward-protocol udp interface [<A.B.C.D>] [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

3. Display the UDP port forwarding table:

```
show ip forward-protocol udp portfwd [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

4. Display the UDP port forwarding list table for the specified list or all lists on the device:

```
show ip forward-protocol udp portfwdlist [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

5. Display the UDP protocol table with the UDP port numbers for each supported or designated protocol:

```
show ip forward-protocol udp [vrf WORD<0-16>] [vrfids WORD<0-512>]
```

Example

Display the UDP protocol table with the UDP port numbers for each supported or designated protocol:

```
Switch:1>enable
Switch:1#show ip forward-protocol udp

Udp Protocol Tbl - GlobalRouter

UDP_PORT PROTOCOL_NAME

Time Service
49 TACACS Service
53 DNS
69 TFTP
137 NetBIOS NameSrv
138 NetBIOS DataSrv
```

Variable definitions

Use the data in the following table to use the show ip forward-protocol udp interface command.

Table 29: Variable definitions

Variable	Value
<a.b.c.d></a.b.c.d>	Specifies the IP address for the interface in a.b.c.d format.
vrf WORD<0-16>	Specifies the name of the VRF.

Variable	Value
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0 to 512.

DHCP and UDP configuration using Enterprise Device Manager

Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), dynamically provides host configuration information to workstations. To lower administrative overhead, network managers prefer to configure a small number of DHCP servers in a central location. Using few DHCP servers requires the routers connecting to the subnets or bridge (or VLAN) domains to support the BootP/DHCP relay function so that hosts can retrieve the configuration information from servers several router hops away.

User datagram protocol (UDP) is a connectionless protocol that adds reliability and multiplexing to IP. It describes how messages reach application programs within a destination computer. Some network applications, such as the NetBIOS name service, rely on a UDP broadcast to request a service or to locate a service. By default, broadcasts are not forwarded by a router. UDP broadcast forwarding is a generalized mechanism for the router to selectively forward UDP broadcasts.

Important:

BootP/DHCP relays are supported only on IP routed port-based VLANs and protocol-based VLANs.

Before you begin

You must enable DHCP relay on the path for port or VLAN configuration to take effect.

Configuring DHCP on a brouter port or a VRF instance

Before you begin

- You must first enable BootP/DHCP relay on a port (or VLAN).
- You must enable DHCP and forwarding path.
- · You must enable IP Routing on the interface.

About this task

Use the DHCP tab to configure the DHCP behavior on a brouter port or a VRF instance. The DHCP tab is available only if the port is routed (that is, assigned an IP address).

Procedure

- 1. In the Device Physical View tab, select a port.
- 2. In the navigation tree, expand the following folders: **Configuration > Edit > Port**.

- 3. Click IP.
- 4. Click the **DHCP Relay** tab.
- 5. Click **Enable** to select the DHCP option. The default is disable.
- 6. Configure the other parameters as needed.
- 7. Click Apply.

DHCP field descriptions

Use data from the following table in the DHCP Relay tab.

Name	Description
Enable	Lets you use BootP/DHCP on the port. The default is disable.
МахНор	Sets the maximum number of hops before a BootP/DHCP packet is discarded (1 to 16). The default is 4.
MinSec	The secs field in the BootP/DHCP packet header represents the elapsed time since the client first sent the message. If the secs field in the packet header is greater than this value, the system relays or forwards the packet; otherwise, the packet is dropped. The default is 0 seconds.
Mode	Sets the interface to process only BootP, only DHCP, or both types of packets. The default is both.
AlwaysBroadcast	When enabled, the server reply is sent as a broadcast back to the end station. The default is disable.
CircuitId	Indicates whether DHCP Relay inserted the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Remoteld	Indicates whether DHCP Relay inserted the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Trusted	Indicates if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.

Configuring BootP/DHCP on a VLAN or VRF instance

Before you begin

• You must enable IP Routing on the interface.

About this task

Use the DHCP Relay tab to configure the DHCP behavior on a VLAN. The DHCP Relay tab is available only if the VLAN is routed and is assigned an IP address.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.

- 2. Click VLANs > Basic.
- 3. Select a VLAN.
- 4. Click IP.
- 5. Click the **DHCP Relay** tab.
- 6. Select Enable.
- 7. Configure the parameters as required.
- 8. Click Apply.

DHCP Relay field descriptions

Use the data in the following table to use the **DHCP Relay** tab.

Variable	Value
Enable	Lets you use BootP/DHCP on the port. The default is disable.
МахНор	Sets the maximum number of hops a BootP/DHCP packet can take from the DHCP client to the DHCP server. The maximum number of hops is 16. The default is 4.
MinSec	Represents the minimum number of seconds to wait between receiving a DHCP packet and forwarding the DHCP packet to the DHCP server. A value of 0 indicates that forwarding is done immediately. The default value is 0.
Mode	Indicates the type of DHCP packet required. The options are:
	• bootp
	• dhcp
	• both
	The default is both.
AlwaysBroadcast	When enabled, the DHCP Reply packets are sent as a broadcast to the DHCP client. The default is disable.
Circuitld	Indicates whether DHCP Relay inserted the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.
Remoteld	Indicates whether DHCP Relay inserted the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.

Variable	Value
Trusted	Indicates if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.

Configuring DHCP relay

About this task

After you configure the BootP/DHCP relay on an IP interface, you can configure forwarding paths to indicate where packets are forwarded. The forwarding paths are based on the type of packet and where the packet is received.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click DHCP Relay.
- 3. Click the Globals tab.
- 4. Click Insert.
- 5. In the **AgentAddr** box, type the agent address.
- 6. In the **ServerAddr** list, type the server address.
- 7. Click **Enable** to enable BootP/DHCP relay. You can enable or disable each agent server forwarding policy. The default is enabled.
- In the **Mode** box, select the type of messages to relay.
 Both the mode setting for the DHCP interface and the mode setting for the agent interface
- determine which packets are forwarded.

Globals field descriptions

9. Click Insert.

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

Name	Description
AgentAddr	The IP address of the input interface (agent) on which the BootP/DHCP request packets are received for forwarding. This address is the IP address of either a brouter port or a VLAN for which forwarding is enabled.
ServerAddr	This parameter is either the IP address of the BootP/DHCP server or the address of another local interface.
	If it is the address of the BootP/DHCP server, the request is unicast to the server address.

Name	Description
	If the address is one of the IP addresses of an interface on the system, the BootP/ DHCP requests are broadcast out of that local interface.
Enable	Enables BootP/DHCP relay.
Mode	Specifies the type of messages relayed:
	Only BootP
	Only DHCP
	Both types of messages
	The default is to forward both BootP and DHCP messages.

Viewing DHCP relay configuration information

About this task

Use the DHCP Relay Interfaces tab to view configuration information about the DHCP relay. To change the configuration information, double-click the value in the field under the required interface, and enter a new value.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click **DHCP Relay**.
- 3. Click the Interfaces tab.

Interfaces field descriptions

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

Variable	Value
IfIndex	A read-only interface number that represents a physical interface, or the VLAN logical interface.
МахНор	Sets the maximum number of hops a DHCP packet can take from the DHCP client to the DHCP server. The maximum number of hops is 16. The default is 4.
MinSec	Represents the minimum number of seconds to wait between receiving a DHCP packet and forwarding the DHCP packet to the DHCP server. A value of 0 indicates that forwarding is done immediately. The default value is 0.
Mode	Indicates the type of DHCP packet required. The options are: • bootp

Variable	Value	
	• dhcp	
	• both	
	The default is both.	
AlwaysBroadcast	Indicates if DHCP Reply packets can be sent as a broadcast to the DHCP client. The default is false.	
Circuitld	Indicates whether DHCP Relay inserted the option 82 circuit ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.	
Remoteld	Indicates whether DHCP Relay inserted the option 82 remote ID information into the DHCP packets before sending the DHCP packets to the DHCP server. The default is disable.	
Trusted	Indicates if DHCP packets come through a trusted DHCP circuit. Only packets with GIADDR configured to 0 and containing option 82 are forwarded if the circuit is trusted. The default value is false.	

Managing UDP forwarding protocols

About this task

The switch configures the following protocols, by default:

- Time Service
- Terminal Access Controller Access Control System (TACACS) Service
- Domain Name System (DNS)
- Trivial file transfer protocol (TFTP)
- Network Basic Input/Output System (NetBIOS) NameSrv
- NetBIOS DataSrv

You can use these protocols to create forwarding entries and lists but you cannot delete them; you can add or remove other protocols to the list of protocols.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click **UDP Forwarding**.
- 3. Click Insert.
- 4. In the **PortNumber** field, type a UDP port number.

This number defines the UDP port used by the server process as its contact port. The range is from 1 to 65535 and cannot be one of the UDP port numbers or a number previously assigned.

- 5. In the **Name** field, type a name for the protocol.
- 6. Click Insert.

The protocol is added to the Protocol table. After you create a protocol, you cannot change its name or number.

Protocols field descriptions

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

Name	Description	
PortNumber	Defines the UDP port (1 to 65535).	
Name	Specifies an administratively assigned name for this list (0 to 15 characters).	

Managing UDP forwarding

About this task

You manage UDP forwarding by defining the destination addresses for the UDP protocol.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration** > **IP**.
- 2. Click **UDP Forwarding**.
- 3. Click the **Forwardings** tab.
- 4. Click Insert.
- 5. In the Insert Forwardings dialog box, select a destination UDP port from the defined protocols in the **DestPort** box.
- 6. Enter a destination IP address in the **DestAddr** box.

The destination address can be any IP server address for the protocol application or the IP address of an interface on the router.

7. Click **Insert**. The information is added to the Forwarding tab.

Forwardings field descriptions

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

Name	Description	
DestPort	Specifies the port number defined for UDP, depending upon the protocol type.	

Name	Description	
DestAddr	Specifies the destination address can be any IP server address for the protocol application or the IP address of an interface on the router:	
	If the address is that of a server, the packet is sent as a unicast packet to this address.	
	If the address is that of an interface on the router, the frame is rebroadcast.	
Id	Specifies an integer that identifies this entry internally.	
NumFwdPackets	Specifies the total number of UDP broadcast packets forwarded using this policy.	
NumDropPacketsTtlExpired	Specifies the total number of UDP broadcast packets dropped because the time-to-live value (TTL) expired.	
NumDropPacketsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the specified destination address was unreachable.	

Creating the forwarding profile

About this task

A forwarding profile is a collection of port and destination pairs. When you configure UDP forwarding list entries, be sure to first configure the UDP forwarding list. Then, configure your UDP forwarding list entries and assign them to a UDP forwarding list. If you do not assign a UDP forwarding list entry to at least one UDP forwarding list, the UDP forwarding list is lost after a restart.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click UDP Forwarding.
- 3. Click the **Forwarding Lists** tab.
- 4. Click Insert.
- 5. In the **Id** field, type the forwarding list ID.
- In the Name field, type the name of the forwarding list if required.The forwarding list appears in the FwdldList box.
- 7. Click Insert.

Forwarding Lists field descriptions

Use the data in the following table to use the **Forwarding Lists** tab and **Insert Forwarding Lists** dialog box.

Name	Description
ld	Specifies a value that uniquely identifies this list of entries (1 to 1000).
Name	Specifies an administratively assigned name for this list (0 to 15 characters).
FwdldList	Specifies the zero or more port forwarding entries associated with this list. Each list identifier is stored as 2 bytes in this array, starting from 0 bytes (size=64). Clicking on the ellipsis () button in this field displays the ID list.

Managing the broadcast interface

About this task

Manage the broadcast interface by specifying and displaying which router interfaces can receive UDP broadcasts to forward.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > IP**.
- 2. Click **UDP Forwarding**.
- Click the Broadcast Interfaces tab.
- 4. Click Insert.
- 5. In the **LocalifAddr** field, click the ellipsis (...) to select a local interface IP address from the list, and then click **OK**.
- 6. In the **UdpPortFwdListId** field, click the ellipsis (...) to select a forwarding list ID from the list, and then click **OK**.
- 7. In the **MaxTtl** field, type the maximum number of hops an IP broadcast can take from the source device to the destination device (the default is 4; the range is 1 to 16).
- 8. In the **BroadCastMask** field, enter the subnet mask of the local interface that broadcasts the UDP broadcast packets.
 - When you configure the UDP forwarding broadcast mask, the broadcast mask must be less specific (shorter in length) or equally specific (equal in length) to the subnet mask of the IP interface on which it is configured. If the UDP forwarding broadcast mask is more specific than the subnet mask of the corresponding IP interface, UDP forwarding does not function properly.
- 9. Click Insert.

Broadcast Interfaces field descriptions

Use the data in the following table to use the **Broadcast Interfaces** tab.

Name	Description	
LocalifAddr	Specifies the IP address of the local router interface that receives forwarded UDP broadcast packets.	
UdpPortFwdListId	Specifies the number of the UDP lists or profiles that this interface is configured to forward (0 to100). A value of 0 indicates that the interface cannot forward any UDP broadcast packets.	
MaxTtl	Specifies the maximum number of hops an IP broadcast packet can take from the source device to the destination device (the default is 4; the range is 1 to 16).	
NumRxPkts	Specifies the total number of UDP broadcast packets received by this local interface.	
NumFwdPkts	Specifies the total number of UDP broadcast packets forwarded by this local interface.	
NumDropPktsMaxTtlExpired	Specifies the total number of UDP broadcast packets dropped because the time-to-live (TTL) value expired.	
NumDropPktsDestUnreach	Specifies the total number of UDP broadcast packets dropped because the destination was unreachable.	
NumDropPktsUnknownPort	Specifies the total number of UDP broadcast packets dropped because the destination port or protocol specified has no matching forwarding policy.	
BroadCastMask	Specifies the subnet mask of the local interface that broadcasts the UDP broadcast packets.	

Viewing UDP endpoint information

View UDP Endpoints to confirm correct configuration.

About this task

You can use UDP endpoint information to display local and remote UDP activity.

Since UDP is a protocol used to establish connectionless network sessions, you need to monitor local and remote UDP activity and to know which applications are running over UDP.

You can determine which applications are active by checking the port number.

Processes are further identified with a UDP session to allow for the multiplexing of a port mapping for UDP.

Procedure

- 1. In the navigation pane, expand the **Configuration > IP** folders.
- 2. Click TCP/UDP.
- 3. Click the **UDP Endpoints** tab.

UDP Endpoints field descriptions

Use the data in the following table to use the **UDP Endpoints** tab.

Name	Description
LocalAddressType	Displays the local address type (IPv6 or IPv4).
LocalAddress	Displays the local IPv6 address.
LocalPort	Displays the local port number.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IPv6 address.
RemotePort	Displays the remote port number.
Instance	Distinguishes between multiple processes connected to the UDP endpoint.
Process	Displays the ID for the UDP process.

Chapter 7: Route Filtering and IP Policies

Table 30: IP Route Policies product support

Feature	Product	Release introduced	
For configuration details, see Configuring IPv4 Routing for VOSS.			
IP route policies	VSP 4450 Series	VSP 4000 4.0	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VSP 8200 4.0	
	VSP 8400 Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 4.5	
	XA1400 Series	VOSS 8.0.50	

Route Filtering and IP Policies

When the switch routes IP traffic, you can apply a number of filters to manage, accept, redistribute, and announce policies for unicast routing table information. Filters apply differently to different unicast routing protocols.



IPv6 ingress QoS ACL/Filters and IPv6 Egress Security and QoS ACL/Filters are not supported. For information on the maximum number of IPv6 ingress port/vlan security ACL/filters supported on the switch, see Release Notes for VOSS.

The following figure shows how filters apply to BGP, RIP, and OSPF protocols.

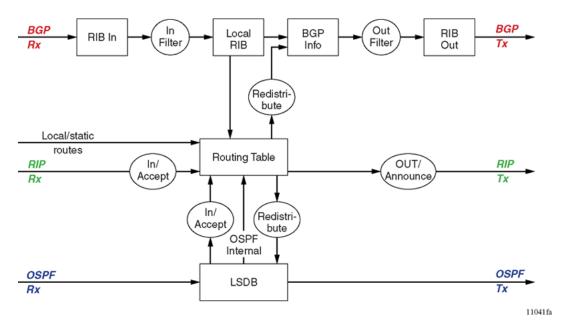


Figure 14: Route filtering for BGP, RIP, and OSPF routing protocols

The following figure shows how filters apply to the IS-IS protocol for Fabric Connect Layer 3 VSNs or IP Shortcuts.

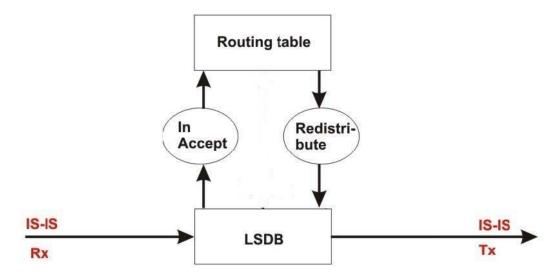


Figure 15: Route filtering for the IS-IS routing protocol

Accept Policies

Accept policies are applied to incoming traffic to determine whether to add the route to the routing table. Accept policies are applied differently to protocols, as follows:

- RIP and BGP—filters apply to all incoming route information.
- OSPF—filters apply only to external route information. Internal routing information is not filtered because otherwise, other routers in the OSPF domain can have inconsistent databases that can affect the router view of the network topology.
- IS–IS —filters apply to all incoming route information.

In a network with multiple routing protocols, you can prefer specific routes from RIP instead of from OSPF. The network prefix is a commonly used match criterion for accept policies.

Redistribution Filters

Redistribution filters notify changes in the route table to the routing protocol (within the device). With redistribution filters, providing you do not breach the protocol rules, you can choose not to advertise everything that is in the protocol database, or you can summarize or suppress route information. By default, no external routes are leaked to protocols that are not configured.

Announce Policies

Announce policies are applied to outgoing advertisements to neighbors or peers in the protocol domain to determine whether to announce specific route information. Out filtering applies to RIP updates and BGP NLRI updates.

In contrast, announce policies are not applied to IS-IS or OSPF information because routing information must always be consistent across the domain. To restrict the flow of external route information in the IS-IS or OSPF protocol database, apply redistribution filters instead of announce policies.

Route Filtering Stages

The following figure shows the three distinct filter stages that are applied to IP traffic.

These stages are:

- Filter stage 1 is the accept policy or in filter that applies to incoming traffic to detect changes in the dynamic (protocol-learned) routing information, which are then submitted to the routing table.
- Filter stage 2 is the redistribution filter that applies to the entries in the routing table to the protocol during the leaking process.

• Filter stage 3 is the announce policy or out filter that applies to outgoing traffic within a protocol domain.

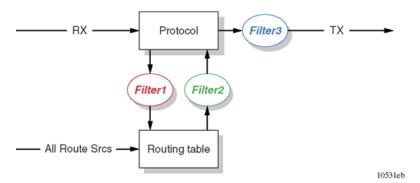


Figure 16: Route filtering stages

The following figure shows the logical process for route filtering on the switch.

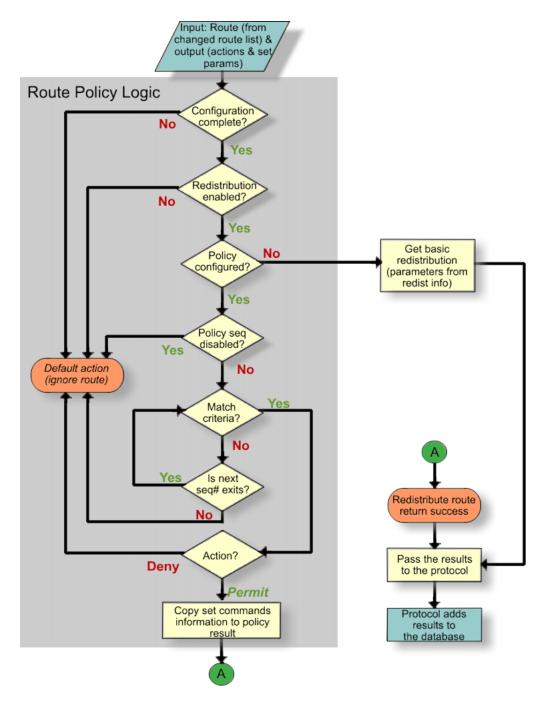


Figure 17: Route filtering logic

Prefix list

In the switch software, you can create one or more IP prefix lists and apply these lists to IP route policy.

Route Policy Definition

You can define an IP route policy and its attributes globally, and then apply them individually to interfaces and protocols. You can also form a unified database of route policies that the RIP or OSPF protocol can use for type of filtering purpose. A name or ID identifies a policy.

Under a policy you can have several sequence numbers. If you do not configure a field in a policy, the field appears as 0 in CLI show command output. This value indicates that the device ignores the field in the match criteria. Use the clear option to remove existing configurations for the field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference field, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce or redistribute purposes.

You can only apply one policy for each purpose (RIP Announce, for example) on a given RIP interface. In this case, all sequence numbers under the policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

The following tables display the accept, announce, and redistribute policies for RIP, OSPF, IS-IS and BGP. The tables also display which matching criteria apply for a certain routing policy. In these tables, 1 denotes advertise router, 2 denotes RIP gateway, and 3 denotes that external type 1 and external type 2 are the only options.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Note:

IPv4 and IPv6 route-maps cannot be configured on the same match statement.

Table 31: Protocol route policy table for RIP

		Announce				
	OSPF	Direct	RIP	BGP	RIP	
Match Protocol	Yes	Yes	Yes	Yes		
Match Network	Yes	Yes	Yes	Yes	Yes	
Match IpRoute Source	Yes ¹		Yes ²			

	Announce				Accept
	OSPF	Direct	RIP	BGP	RIP
Match NextHop	Yes	Yes	Yes	Yes	Yes
Match Interface			Yes		
Match Route Type	Yes				
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
NssaPbit					
SetRoute Preference					Yes
SetMetric TypeInternal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type					
SetNextHop					
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					Yes
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 32: Protocol route policy table for OSPF

		Redistribute				
	Direct	Static	RIP	BGP	IS-IS	OSPF
Match Protocol				Yes	Yes	
Match Network	Yes	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes ²			
Match NextHop		Yes	Yes	Yes		

	Redistribute				Accept	
	Direct	Static	RIP	BGP	IS-IS	OSPF
Match Interface			Yes			
Match Route Type					Yes ³	
Match Metric	Yes	Yes	Yes	Yes	Yes	Yes
MatchAs Path						
Match Community						
Match Community Exact						
MatchTag				Yes		
Set NSSA Bit	Yes	Yes	Yes	Yes	Yes	
SetRoute Preference						
SetMetric TypeInternal						
SetMetric	Yes	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	Yes	
SetNextHop				Yes		
Set Inject NetList	Yes	Yes	Yes	Yes	Yes	Yes
SetMask						
SetAsPath						
SetAsPath Mode						
Set Automatic Tag						
Set CommunityNumber						
Set CommunityMode						
SetOrigin						
SetLocal Pref						
SetOrigin EgpAs						
SetTag						
SetWeight						

Table 33: Protocol route policy table for IS-IS

		Redistribute			Accept
	Direct	Static	RIP	BGP	OSPF
Match Protocol	Yes	Yes	Yes	Yes	Yes
Match Network	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source			Yes		
Match NextHop		Yes	Yes	Yes	Yes
Match Interface			Yes		

	Redistribute				
	Direct	Static	RIP	BGP	OSPF
Match Route Type					Yes ³
Match Metric	Yes	Yes	Yes	Yes	Yes
MatchAs Path					
Match Community					
Match Community Exact					
MatchTag				Yes	
Set NSSA Bit					
SetRoute Preference					
SetMetric Type Internal					
SetMetric	Yes	Yes	Yes	Yes	Yes
SetMetric Type	Yes	Yes	Yes	Yes	Yes
SetNextHop				Yes	
SetInject NetList	Yes	Yes	Yes	Yes	Yes
SetMask					
SetAsPath					
SetAsPath Mode					
Set Automatic Tag					
Set CommunityNumber					
Set CommunityMode					
SetOrigin					
SetLocal Pref					
SetOrigin EgpAs					
SetTag					
SetWeight					

Table 34: Protocol route policy table for BGP

		Redistribute			Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
Match as-path				Yes	Yes
Match community	Yes	Yes	Yes	Yes	Yes
Match community-exact				Yes	Yes
Match extcommunity				Yes	Yes
Match interface					
Match local-preference					

		Redistribute		Accept	Announce
	IPv6 Direct	IPv6 Static	OSPFv3	BGP	BGP
Match metric	Yes	Yes	Yes	Yes	Yes
Match network	Yes	Yes	Yes	Yes	Yes
Match next-hop		Yes	Yes	Yes	Yes
Match protocol					
Match route-source				Yes	
Match route-type			Yes		Yes
Match tag					
Match vrf					
Match vrfids					
Set as-path				Yes	Yes
Set as-path-mode				Yes	Yes
Set automatic-tag					
Set community				Yes	Yes
Set community-mode				Yes	Yes
Set injectlist	Yes	Yes	Yes		
Set ip-preference					
Set local-preference				Yes	Yes
Set mask					
Set metric	Yes	Yes	Yes	Yes	Yes
Set metric-type					
Set metric-type-internal					
Set next-hop				Yes	Yes
Set nssa-pbit					
Set origin					Yes
Set origin-egp-as					
Set Tag					
Set Weight				Yes	

IP policy configuration using the CLI

Configure IP policies to form a unified database of route policies that Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) can use for filtering tasks.

A policy is identified by a name or an ID. Under a given policy you can have several sequence numbers, each of which is equal to one policy in the old convention. Each policy sequence number

contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a configured set-preference field, use only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce and redistribute purposes.

You can apply one policy for one purpose, for example, RIP announce on a RIP interface. All sequence numbers under the given policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

Configuring prefix lists

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

About this task



When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure a prefix list:

```
ip prefix-list WORD < 1-64 > \{A.B.C.D/X\} [ge < 0-32 >] [id < 1-2147483647 >] [le < 0-32 >]
```

3. (Optional) Rename an existing prefix list:

```
ip prefix-list WORD<1-64> name WORD<1-64>
```

4. Display the prefix list:

```
show ip prefix-list [prefix {A.B.C.D}] [vrf WORD < 1-16 >] [vrfids WORD < 0-512 >] [WORD < 1-64 >]
```

Example

Configure a prefix-list. Display the prefix list.

```
Switch> enable
Switch# configure terminal
Switch(config)# ip prefix-list LIST1 192.0.2.1/255.255.255.0
Switch(config)# show ip prefix-list LIST1
```

```
PREFIX MASKLEN FROM TO

List 1 LIST1:

192.0.1.2.1 24 24
1 Total Prefix List entries configured

Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
```

Variable definitions

Use the data in the following table to use the ip prefix-list command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats:
	• a.b.c.d/x
	• a.b.c.d/x.x.x.x
	default
ge <0-32>	Specifies the minimum length to match.
	Lower bound and higher bound mask lengths together can define a range of networks.
id <1-2147483647>	Specifies the Prefix list ID.
le <0–32>	Specifies the maximum length to match.
	Lower bound and higher bound mask lengths together can define a range of networks.
name WORD<1-64>	Renames the specified prefix list. The name length is 1–64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

Use the data in the following table to use the **show** ip **prefix-list** command.

Variable	Value
{A.B.C.D}	Specifies the prefix to include in the command output.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Use the following table to use the **show** ip **prefix-list** command output.

Variable	Value
PREFIX	Indicates the member of a specific prefix list.
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
ТО	Indicates the prefix mask endpoint in bits.

Configuring IP Route Policies

Configure a route policy so that the device can control routes that certain packets can take. For example, you can use a route policy to deny certain Border Gateway Protocol (BGP) routes.

The route policy defines the matching criteria and the actions taken if the policy matches.

About this task

After you create and enable the policy, you can apply it to an interface. You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In this case, all sequence numbers under the given policy apply to that filter.

Create and enable the policy for IS-IS accept policies for Fabric Connect for Layer 3 Virtual Services Networks (VSNs) and IP Shortcuts, then apply the IS-IS accept policy filters. For more information on IS-IS accept policy filters, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Note:

After you configure route-map in Global Configuration mode or VRF Router Configuration mode, the device enters Route-Map Configuration mode, where you configure the action the policy takes, and define other fields the policy enforces.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Note:

You cannot configure IPv4 and IPv6 route-maps on the same match statement.

Procedure

1. Enter Route-Map Configuration mode:

enable

```
configure terminal
route-map WORD<1-64> <1-65535>
```

2. At the route-map prompt, define the match criteria for the policy:

```
match {as-path WORD < 0-256 > | community WORD < 0-256 > | community-exact enable | extcommunity WORD < 0-1027 > | interface WORD < 0-259 > | local-preference < 0-2147483647 > |metric < 0-65535 > | metric-type-isis < any | internal|external> | network WORD < 0-259 > | next-hop WORD < 0-259 > | protocol WORD < 0-60 > | route-source WORD < 0-259 > | route-type < any | | local | internal | external | external - 1 | external - 2 > | tag | WORD < 0-256 > | vrf | WORD < 1-16 > | vrfids | WORD < 0-512 > }
```

- 3. Define the action the policy takes:
 - a. Allow the route:

```
permit
```

OR

b. Ignore the route:

no permit

4. Define the set criteria for the policy:

```
set {as-path WORD < 0-256 > | as-path-mode <tag|preprend> | automatic-tag enable | community WORD < 0-256 > | community-mode <additive|none|unchanged>| injectlist WORD < 0-1027 > | ip-preference <0-255> | local-preference <0-2147483647> | mask <A.B.C.D> | metric <0-65535> | metric-type <type1|type2> | metric-type-internal <0-1> | metric-type-isis <none|internal|external>| metric-type-live-metric | next-hop WORD < 0-256 > | nssa-pbit enable | origin <igp|egp|incomplete> | origin-egp-as <0-65535>| tag WORD < 0-256 > | weight <0-65535> }
```

5. Display current information about the IP route policy:

```
show route-map [WORD<1-64>] [<1-65535>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Enter Route-Map Configuration mode. At the route-map prompt, define the fields the policy enforces. Define the action the policy takes. Display current information about the IP route policy.

RedisStatic 1 PRMT DIS

Variable Definitions

Use the data in the following table to use the match command.

Variable	Value
as-path WORD<0-256>	Configures the device to match the as-path attribute of the Border Gateway Protocol (BGP) routes against the contents of the specified AS-lists. This field is used only for BGP routes and ignored for all other route types.
	WORD <0-256> specifies the list IDs of up to four AS-lists, separated by a comma.
	Use the no operator to disable match as-path: no match as-path WORD<0-256>
community WORD<0-256>	Configures the device to match the community attribute of the BGP routes against the contents of the specified community lists. This field is used only for BGP routes and ignored for all other route types.
	WORD <0-256> specifies the list IDs of up to four defined community lists, separated by a comma.
	Use the no operator to disable match community: no match community WORD<0-256>
community-exact enable	When disabled, configures the device so match community- exact results in a match when the community attribute of the BGP routes match an entry of a community-list specified in match-community.
	When enabled, configures the device so match-community- exact results in a match when the community attribute of the BGP routes matches all of the entries of all the community lists specified in match-community.
	enable enables match community-exact.
	Use the no operator to disable match community-exact: no match community-exact enable
extcommunity WORD <0-1027>	Configures the device to match the extended community.
	WORD<0-1027> specifies an integer value from 1–1027 that represents the community list ID you want to create or modify.
interface WORD <0-259>	If configured, configures the device to match the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other route types.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.

Variable	Value
	Use the no operator to disable match-interface: no match interface WORD <0-259>
local-preference <0-2147483647>	Configures the device to match the local preference, applicable to all protocols.
	<0-2147483647> specifies the preference value.
metric <0-65535>	Configures the device to match the metric of the incoming advertisement or existing route against the specified value. If 0, this field is ignored.
	<0-65535> specifies the metric value. The default is 0.
network WORD <0-259>	Configures the device to match the destination network against the contents of the specified prefix lists.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match network: no match network WORD <0-259>
next-hop WORD<0-259>	Configures the device to match the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match next hop: no match next-hop WORD<0-259>
protocol WORD<0-60>	Configures the device to match the protocol through which the route is learned.
	WORD <0-60> is xxx, where xxx is local, ospf, ebgp, ibgp,isis, rip, static, or a combination separated by ,
	Use the no operator to disable match protocol: no match protocol WORD<0-60>
route-source WORD<0-259>	Configures the system to match the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
	WORD <0-259> specifies the name of up to four defined prefix lists, separated by a comma.
	Use the no operator to disable match route source: no match route-source WORD<0-259>
route-type {any local internal external external-1 external-2}	Configures a specific route type to match (applies only to OSPF routes).

Variable	Value
	any local internal external external-1 external-2 specifies OSPF routes of the specified type only (External-1 or External-2). Another value is ignored.
tag WORD<0-256>	Specifies a list of tags used during the match criteria process. Contains one or more tag values.
	WORD<0-256> is a value from 0-256.
[vrf WORD<1-16>] [vrfids WORD<0-512>]	Configures a specific VRF to match (applies only to RIP routes).

Use the data in the following table to use the set command.

Variable	Value
as-path WORD<0-256>	Configures the device to add the AS number of the AS-list to the BGP routes that match this policy.
	WORD<0-256> specifies the list ID of up to four defined AS-lists separated by a comma.
	Use the no operator to delete the AS number: no set aspath WORD<0-256>
as-path-mode <tag prepend></tag prepend>	Configures the AS path mode.
	Prepend is the default configuration. The device prepends the AS number of the AS-list specified in set-as-path to the old aspath attribute of the BGP routes that match this policy.
	Note:
	Prepend is not applicable to an internal BGP (iBGP) peer with outbound route policy. For more information about iBGP, see Configuring BGP Services for VOSS .
automatic-tag enable	Configures the tag automatically. Used for BGP routes only.
	Use the no operator to disable the tag: no set automatic-tag enable
community WORD<0-256>	Configures the device to add the community number of the community list to the BGP routes that match this policy.
	WORD <0-256> specifies the list ID of up to four defined community lists separated by a comma.
	Use the no operator to delete the community number: no set community WORD<0-256>
community-mode <additive none < td=""><td>Configures the community mode.</td></additive none <>	Configures the community mode.
unchanged>	additive—the device prepends the community number of the community list specified in set-community to the old community path attribute of the BGP routes that match this policy.

Variable	Value
	none—the device removes the community path attribute of the BGP routes that match this policy to the specified value.
injectlist WORD<0-1027>	Configures the device to replace the destination network of the route that matches this policy with the contents of the specified prefix list.
	WORD<0-1027> specifies one prefix list by name.
	Use the no operator to disable set injectlist: no set injectlist
ip-preference <0-255>	Configures the preference. This applies to accept policies only.
	<0-255> is the range you can assign to the routes.
local-preference <0-65535>	Configures the device to match the local preference, applicable to all protocols. <0–655356> specifies the preference value.
mask <a.b.c.d></a.b.c.d>	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.
	A.B.C.D is a valid contiguous IP mask.
	Use the no operator to disable set mask: no set mask
metric <0-65535>	Configures the metric value for the route while announcing a redistribution. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF for RIP, the original cost of the route or default-import-metric is used (applies to IS-IS routes also).
metric-type {type1 type2}	Configures the metric type for the routes to announce into the OSPF domain that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.
metric-type-internal <0-1>	Configures the MED value for routes advertised to ebgp nbrs to the IGP metric value.
	<0-1> specifies the metric type internal.
metric-type-isis <none external="" internal="" =""></none>	Configures the metric type for IS-IS routes. The default is none. This field is applicable only for IS_IS policies.
metric-type-live-metric	Configures the metric type for BGP routes. The default is disabled. This field is applicable only for BGP policies.
next-hop WORD <1-256>	Specifies the IP address of the next-hop router. Both IPv4 and IPv6 addresses are supported.
	Use the no operator to disable set next-hop: no set next-hop
nssa-pbit enable	Configures the not so stubby area (NSSA) translation P bit. Applicable to OSPF announce policies only.
	Use the no operator to disable set nssa-pbit: no set nssa-pbit enable

Variable	Value
origin {igp egp incomplete}	Configures the device to change the origin path attribute of the BGP routes that match this policy to the specified value.
origin-egp-as <0-65535>	Indicates the remote autonomous system number. Applicable to BGP only.
tag <0-65535>	Configures the tag of the destination routing protocol. If not specified, the device forwards the tag value in the source routing protocol. A value of 0 indicates that this parameter is not configured.
	Note:
	This parameter is not supported on all hardware platforms.
weight <0-65535>	Configures the weight value for the routing table. For BGP, this value overrides the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. Used for BGP only. A value of 0 indicates that this parameter is not configured.

Use the data in the following table to use the name command.

Variable	Value
WORD<1-64>	Renames a policy and changes the name field for all
	sequence numbers under the given policy.

Job aid

Use the data in the following table to use the <code>show route-map</code> command output.

Table 35: Variable definitions

Variable	Value
NAME	Indicates the name of the route policy.
SEQ	Indicates the second index used to identify a specific policy within the route policy group (grouped by ID). Use this field to specify different match and set parameters and an action.
MODE	Indicates the action to take when this policy is selected for a specific route. Options are permit, deny, or continue. Permit indicates to allow the route. Deny indicates to ignore the route. Continue means continue checking the next match criteria configured in the next policy sequence; if none, take the default action in the given context.
EN	Indicates whether this policy is enabled. If disabled, the policy is not used.

Configuring a policy to accept external routes from a router

Perform this procedure to configure a policy to accept external routes from a specified advertising router.

For more information on IS-IS accept policy filters for Fabric Connect for Layer 3 VSNs and IP Shortcuts, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter OSPF Router Configuration mode:

```
enable
configure terminal
router ospf
```

2. Create a policy to accept external routes from a specified advertising route:

```
accept adv-rtr <A.B.C.D>
```

3. Exit to the Privileged EXEC mode.

exit

4. Apply the OSPF accept policy change:

```
ip ospf apply accept adv-rtr <A.B.C.D>
```

5. Confirm your configuration:

```
show ip ospf accept
```

Example

Log on to the OSPF Router Configuration mode:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router ospf
```

Create a policy to accept external routes from a specified advertising route:

Switch:1(config-ospf) #accept adv-rtr 192.0.2.122

Enable an OSPF accept entry for a specified advertising route:

Switch:1(config-ospf) #accept adv-rtr 192.0.2.122 enable

Exit to the Privileged EXEC mode:

```
Switch:1(config-ospf) #exit
Switch:1(config) #exit
```

Apply the OSPF accept policy change and confirm your configuration:

```
Switch:1#ip ospf apply accept adv-rtr 192.0.2.122
Switch:1#show ip ospf accept

Ospf Accept - GlobalRouter

ADV_RTR MET_TYPE ENABLE POLICY

192.0.2.122 - FALSE
```

Variable definitions

Use the data in the following table to use the accept adv-rtr command.

Variable	Value
<a.b.c.d></a.b.c.d>	Specifies the IP address.
enable	Enables an OSPF accept entry for a specified advertising router.
	Use the no operator to disable an OSPF accept entry: no accept adv-rtr <a.b.c.d> enable</a.b.c.d>
metric-type {type1 type2 }	Indicates the OSPF external type. This parameter describes which types of OSPF external routes match this entry.
	means match all external routes.
	type1 means match external type 1 only.
	type2 means match external type 2 only.
	Use the no operator to disable metric-type: no ip ospf accept adv-rtr <a.b.c.d> metric-type</a.b.c.d>
route-map WORD<0-64>	Specifies the name of the route policy to use for filtering external routes advertised by the specified advertising router before accepting into the routing table.

Applying OSPF accept policy changes

Apply OSPF accept policy changes to allow the configuration changes in the policy to take effect in an OSPF Accept context (and to prevent the device from attempting to apply the changes one by one after each configuration change).

For more information on IS-IS accept policy filters for Fabric Connect for Layer 3 VSNs and IP Shortcuts, see Configuring Fabric Basics and Layer 2 Services for VOSS.

About this task



Changing OSPF Accept contexts is a process-oriented operation that can affect system performance and network accessibility while you perform the procedures. If you want to change the default preferences for an OSPF Accept or a prefix-list configuration (as opposed to the default preference), do so before enabling the protocols.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Apply an OSPF accept policy change:

```
ip ospf apply accept [vrf WORD<1-16>]
```

3. Display information about the configured OSPF entries:

```
show ip ospf accept [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

Apply the OSPF accept policy and confirm the configuration:

```
Switch:1>enable
Switch:1#ip ospf apply accept
Switch:1#show ip ospf accept

Ospf Accept - GlobalRouter

ADV_RTR MET_TYPE ENABLE POLICY

192.0.2.122 - TRUE
```

Variable definitions

Use the data in the following table to use the ip ospf apply accept adv-rtr command.

Table 36: Variable definitions

Variable	Value
adv-rtr	Commits entered changes. Issue this command after you modify a policy configuration that affects an OSPF accept policy.
vrf WORD<1–16>	Specifies the name of the VRF.

Use the data in the following table to use the show ip ospf accept command output.

Table 37: Variable definitions

Variable	Value
ADV_RTR	Indicates the router advancing the packets.
MET_TYPE	Indicates the metric type for the routes to import into OSPF routing protocol, which passed the matching criteria configured in this route policy. Options include: local, internal, external, externaltype1, and externaltype2.
ENABLE	Indicates if the policy is enabled.
POLICY	Indicates the type of policy.

Configuring inter-VRF redistribution policies

Configure redistribution entries to allow a protocol to announce routes of a certain source type, for example, static, RIP, or direct.

For more information on IS-IS redistribution, see <u>Configuring Fabric Basics and Layer 2 Services for VOSS</u>.

Before you begin

- Ensure the routing protocols are globally enabled.
- You must configure the route policy, if required.
- · Ensure the VRFs exist.
- You must create the route policy and prefix list under the source VRF context.

Note:

The route policies treat permit and deny rules differently for inbound and outbound traffic.

- For an in-policy (RIP, BGP) or an accept policy (OSPF) using a route-map, if a particular route is not explicitly denied in the accept policy or in-policy with the route-map, then the route is implicitly allowed.
- For an out-policy (RIP, BGP) or a redistribute policy (RIP, OSPF, BGP) using a route-map, even if a particular route is not explicitly allowed in the redistribution policy or out-policy with the route-map, then the route is implicitly denied.
- In order to permit or deny only explicit routes, configure a policy with additional sequences, where, the last sequence permits all routes that are not explicitly permitted or denied.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create the redistribution instance:

ip <bgp|ospf|rip> redistribute <bgp|direct|ipv6-direct|ipv6-isis|
ipv6-static|isis|ospf|ospfv3|rip|static|dvr>

3. Apply a route policy if required:

ip <bgp|ospf|rip> redistribute <bgp|direct|ipv6-direct|ipv6-isis|
ipv6-static|isis|ospf|ospfv3|rip|static|dvr> route-map <WORD 0-64>
[vrf-src <WORD 1-16>]

- 4. Use the following variable definitions table to configure other parameters as required.
- 5. Enable the redistribution:

ip <bgp|ospf|rip> redistribute <bgp|direct|ipv6-direct|ipv6-isis|
ipv6-static|isis|ospf|ospfv3|rip|static|dvr> enable [vrf-src <WORD
1-16>]

6. Ensure that the configuration is correct:

show ip $\langle bgp|ospf|rip \rangle$ redistribute [vrf WORD $\langle 1-16 \rangle$] [vrfids $WORD \langle 0-512 \rangle$]

For RIPng, use show ipv6 rip redistribute.

7. Apply the redistribution:

ip <bgp|ospf|rip> apply redistribute <bgp|direct|ipv6-direct|ipv6isis|ipv6-static|isis|ospf|ospfv3|rip|static|dvr> [vrf WORD<1-16>]
[vrf-src WORD<1-16>]

Example

Switch:1>enable Switch:1#config terminal

Log on to the VRF Router Configuration mode:

Switch:1(config) #router vrf test

Create the redistribution instance:

Switch:1(router-vrf) #ip rip redistribute ospf

Enable the redistribution

Switch:1(router-vrf) #ip rip redistribute ospf enable

Ensure that the configuration is correct:

Switch:1(router-vrf) #show ip rip redistribute

Exit to Global Configuration mode:

Switch:1(router-vrf)#exit

Apply the redistribution:

Switch:1(config) #ip rip apply redistribute ospf

Variable definitions

Use the data in the following table to use the redistribution commands.

Variable	Value
 static dvr>	Specifies the type of routes to redistribute—the protocol source.
vrf WORD<1-16>	Specifies the VRF instance.
vrfids WORD<0-512>	Specifies a list of VRF IDs.
vrf-src WORD<1-16>	Specifies the source VRF instance. This parameter is not required for redistribution within the same VRF.

Use the data in the following table to use the ip <bgp|ospf|rip> redistribute <bgp| direct|isis|ospf|rip|static|dvr> command.

Variable	Value
apply [vrf-src WORD<1–16>]	Applies the redistribution configuration.
enable [vrf-src WORD<1-16>]	Enables the OSPF route redistribution instance.
metric <metric-value> [vrf-src WORD<1–16>]</metric-value>	Configures the metric to apply to redistributed routes.
metric-type <type1 type2> [vrf-src WORD<1–16>]</type1 type2>	Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone.
route-map <word 0-64=""> [vrf-src WORD<1-16>]</word>	Configures the route policy to apply to redistributed routes.
subnets <allow suppress> [vrf-src WORD<1–16>]</allow suppress>	Allows or suppresses external subnet route advertisements when routes are redistributed into an OSPF domain.

IP Policy Configuration using Enterprise Device Manager

You can form a unified database of route policies that the protocols (RIP, OSPF or Border Gateway Protocol [BGP]) can use for any type of filtering task.

For information about configuring a prefix list, community list, or AS path list, see <u>Configuring BGP Services for VOSS</u>.

A name or an ID identifies a policy. Under a policy you can have several sequence numbers, each of which is equal to one policy in the old convention. If a field in a policy is not configured, it appears as 0 or any when it appears in Enterprise Device Manager (EDM). This means that the field is

ignored in the match criteria. You can use the clear option to remove existing configurations for any field.

Each policy sequence number contains a set of fields. Only a subset of those fields is used when the policy is applied in a certain context. For example, if a policy has a set-preference field set, it is used only when the policy is applied for accept purposes. This field is ignored when the policy is applied for announce and redistribute purposes.

You can apply only one policy for one purpose (for example, RIP Announce on a given RIP interface). In that example, all sequence numbers under the given policy are applicable for that filter. A sequence number also acts as an implicit preference: a lower sequence number is preferred.

Configuring a prefix list

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or non-contiguous routes. Reference prefix lists by name from within a routing policy.

Before you begin

• Change the VRF instance as required to configure a prefix list on a specific VRF instance.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click Policy.
- 3. Click the Prefix List tab.
- 4. Click Insert.
- 5. In the **Id** box, type an ID for the prefix list.
- 6. In the **Prefix** box, type an IP address for the route.
- 7. In the **PrefixMaskLength** box, type the length of the prefix mask.
- 8. Configure the remaining parameters as required.
- 9. Click Insert.

Prefix List field descriptions

Use the data in the following table to use the **Prefix List** tab.

Name	Description
ld	Configures the list identifier.
Prefix	Configures the IP address of the route.
PrefixMaskLen	Configures the specified length of the prefix mask.

Name	Description
	You must enter the full 32-bit mask to exact a full match of a specific IP address, for example, if you create a policy to match on the next hop.
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name length can use from 1 to 64 characters.
MaskLenFrom	Configures the lower bound of the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Configures the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Configure a Route Policy

Configure a route policy so that all protocols use them for In, Out, and Redistribute purposes.

For more information on IS-IS accept policy filters for Fabric Connect for Layer 3 VSNs and IP Shortcuts, see Configuring Fabric Layer 3 Services for VOSS.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click **Policy**.
- 3. Click the Route Policy tab.
- 4. Click Insert.
- 5. Enter the appropriate information for your configuration in the Insert Route Policy dialog box.
- 6. Click Insert.

Route Policy Field Descriptions

Use the data in the following table to use the **Route Policy** tab.

Name	Description
Id	Specifies the ID of an entry in the Prefix list table.
SequenceNumber	Specifies a policy within a route policy group.
Name	Specifies the name of the policy. This command changes the name field for all sequence numbers under the given policy.
Enable	Indicates whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored. The default is disabled.
Mode	Specifies the action to take when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route). The default is permit.

Name	Description
MatchProtocol	Selects the appropriate protocol. If configured, matches the protocol through which the route is learned. This field is used only for RIP Announce purposes. The default is to enable all match protocols.
MatchNetwork	Specifies if the system matches the destination network against the contents of the specified prefix list.
MatchipRouteSource	Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
	Click the ellipsis button and choose from the list in the Match Route Source dialog box. You can select up to four entries. To clear an entry, use the ALT key.
	You can also change this field in the Route Policy tab of the Policy dialog box.
MatchipRouteDest	Specifies if the system matches the next-hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
MatchNextHop	Specifies if the system matches the next-hop IP address of the route against the contents of the specified prefix list. This field applies only to nonlocal routes.
	Click the ellipsis button and choose from the list in the Match Next Hop dialog box. You can select up to four entries. To clear an entry, use the ALT key.
MatchInterface	Specifies if the system matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route.
	Click the ellipsis button and choose from the list in the Match Interface dialog box. You can select up to four entries. To clear an entry, use the ALT key.
MatchRouteType	Configures a specific route type to match (applies only to OSPF routes).
	Externaltype1 and Externaltype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes. The default is any.
MatchMetric	Specifies if the system matches the metric of the incoming advertisement or existing route against the specified value (1 to 65535). If 0, this field is ignored. The default is 0.
MatchMetricTypelsis	Specifies the match metric type field in the incoming ISIS routes in accept policy.

Name	Description
MatchAsPath	Configures if the system matches the BGP autonomus system path. Applicable to BGP only. This overrides the BGP neighbor filter list information.
MatchCommunity	Filters incoming and outgoing updates based on a Community List. Applicable to BGP only. The default is disable.
MatchCommunityExact	Indicates if the match must be exact (that is, all of the communities specified in the path must match). Applicable to BGP only. The default is disabled.
MatchTag	Specifies a list of tags used during the match criteria process. Applicable to BGP only. It contains one or more tag values.
MatchVrf	Identifies the source VRFs that leaks routes to the local VRF (applies only to RIP routes).
MatchLocalPref	Specifies the local preference value to be matched.
NssaPbit	Configures or resets the P bit in specified type 7 link state advertisements (LSA). By default, the Pbit is always configured in case the user configures the Pbit to a disable state for a particular route policy other than all type 7. LSAs associated with that route policy have the Pbit cleared. With this intact, not so stubby area (NSSA) area border router (ABR) does not perform translation of these LSAs to type 5. The default is enable.
SetRoutePreference	Configures a value from 0 to 255. The default value is 0. If the default is configured, the global preference value is used.
	When configured to a value greater than zero, specifies the route preference value assigned to the routes that matches the policy. This feature applies to accept policies only.
SetMetricTypeInternal	Identifies the MED value for routes advertised to BGP numbers to the Interior Gateway Protocol (IGP) metric value. The value must be 0 or 1. The default is 0.
SetMetricTypeIsis	Sets the metric type IS-IS.
SetMetric	Configures the system to use the metric value for the route while announcing or redistributing. The default-import-metric is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or the default value is used (applies to IS-IS routes also). The default is 0.
SetMetricType	Configures the metric type for the routes to announce into the OSPF routing protocol that matches this policy. Applicable to OSPF protocol only. The default is type 2. This field is applicable only for OSPF announce policies. The default is type2.
SetNextHop	Configures the IP address of the next-hop router. Applicable to BGP only. The default is 0.0.0.0.

Name	Description
SetInjectNetList	Configures the destination network of the route that matches this policy with the contents of the specified prefix list. Click the ellipsis button and choose from the list in the Set Inject NetList dialog box.
SetMask	Configures the mask of the route that matches this policy. This applies only to RIP accept policies.
SetAsPath	Indicates the AS path value to use whether the SetAsPathMode field is Tag or Prepend. Applicable to BGP only.
	Note:
	Prepend is not applicable to an internal BGP (iBGP) peer with outbound route policy. For more information about iBGP, see Configuring BGP Services for VOSS .
SetAsPathMode	Configures if the system converts the tag of a route into an AS path. Applicable to BGP protocol only. The mode is either Tag or Prepend tag. The value is applicable only while redistributing routes to BGP The default is prepend.
	Note:
	Prepend is not applicable to an iBGP peer with outbound route policy. For more information about iBGP, see Configuring BGP Services for VOSS .
SetAutomaticTag	Enables the automatic tag feature. Applicable to BGP protocol only. The default is disable.
SetCommunityNumber	Configures the community number for BGP advertisements. This value can be a number (1 to 42949672000) or no-export or no-advertise.
SetCommunityMode	Configures the community mode for the BGP protocol. This value can be either append, none, or unchanged. The default is unchanged.
	Unchanged—keeps the community attribute in the route path as it is.
	None—removes the community in the route path additive.
	Append—adds the community number specified in SetCommunityNumber to the community list attribute.
SetExtCommunity	Configures a BGP community. The values are 0 to 256.
SetExtCommunityMode	Configures the extended-community mode. The value can be append, unchanged, or overwrite. The default value is unchanged.
	append — creates another community string.
	unchanged — keeps the community attribute as it is.

Name	Description
	overwrite — changes the current value.
SetOrigin	Configures the origin for the BGP protocol to IGP, EGP, incomplete, or unchanged. If not configured, the system uses the route origin from the IP routing table (protocol). The default is unchanged.
SetLocalPref	Configures the local preference for the BGP protocol only. The system uses this value during the route decision process for the BGP protocol. The default is 0.
SetOriginEgpAs	Indicates the remote autonomous system number for the BGP protocol. The default is 0.
SetWeight	Configures the weight value for the routing table for the BGP protocol. This field must be used with the match as-path condition. For BGP, this value overrides the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. The default is 0.
SetTag	Configures the list of tags used during the match criteria process for the BGP protocol. The default is 0.
Ipv6SetNextHop	Specifies the address of the IPv6 next hop router.

Apply a Route Policy

Apply route policies to define route behavior.

For more information on IS-IS accept policy filters for Fabric Connect for Layer 3 VSNs and IP Shortcuts, see Configuring Fabric Layer 3 Services for VOSS.

About this task



! Important:

Changing route policies or prefix lists that affect OSPF accept or redistribute is a processoriented operation that can affect system performance and network reachability while you perform the procedures. Therefore, if you want to change a prefix list or a routing protocol, you configure all route policies and prefix lists before enabling the protocols.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
- 2. Click Policy
- 3. Click the **Applying Policy** tab.
- 4. Select the type of policy to apply.
- 5. Click Apply.

Applying Policy field descriptions

Use the data in the following table to use the **Applying Policy** tab.

Name	Description
RoutePolicyApply	Specifies that configuration changes in the policy take effect in an OSPF route policy context. This prevents the system from attempting to apply the changes one by one after each configuration change. The default is enabled.
RedistributeApply	Specifies that configuration changes in the policy take effectfor an OSPF Redistribute context. This prevents the system from attemptingto apply the changes one-by-one after each configuration change. The default is enabled.
OspfInFilterApply	Specifies that configuration changes in a route policy or a prefix list take effect in an OSPF Accept context. This prevents the system from attempting to apply the changes one by one after each configuration change. The default is enabled.
	Note:
	This field does not appear on all hardware platforms. If you do not see this command in EDM, the feature is not supported on your hardware.

Viewing IP routes

View IP routes learned on the device.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click IP.
- 3. Click the **Routes** tab to view IP routes learned on the device.
- 4. If you want to limit the routes displayed, click **Filter** to show a smaller subset of the learned routes.
- 5. In the Filter dialog box, select an option, or options, and enter information to limit the routes to display in the Routes table.
- 6. Click **Filter** and the Routes table displays only the routes that match the options and information that you enter.

Routes field descriptions

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

Name	Description
Dest	Specifies the destination IP network of this route. An entry with a value of 0.0.0.0 is a default route. Multiple routes to a single destination can appear in the table, but access to multiple entries depends on the table

Name	Description
	access mechanisms defined by the network management protocol in use.
Mask	Indicates the network mask to logically add with the destination address before comparison to the destination IP network.
NextHop	Specifies the IP address of the next hop of this route.
AltSequence	Indicates the alternative route sequence. The value of 0 denotes the best route.
NextHopId	Displays the MAC address or hostname of the next hop.
HopOrMetric	Displays the primary routing metric for this route. The semantics of this metric are specific to different routing protocols.
Interface	Specifies the router interface for this route.
	Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation.
	Brouter interfaces are identified by the slot and port number of the brouter port.
Proto	Specifies the routing mechanism through which this route was learned:
	other—none of the following
	local—nonprotocol information, for example, manually configured entries
	• static
	• ICMP
	• EGP
	• GGP
	• Hello
	• RIP
	• IS-IS
	• ES-IS
	Cisco IGRP
	• bbnSpflgp
	• OSPF
	• BGP
	Inter-VRF Redistributed Route
Age	Displays the number of seconds since this route was last updated or otherwise determined to be correct.

Name	Description
PathType	Indicates the route type, which is a combination of direct, indirect, best, alternative, and ECMP paths.
	iA indicates Indirect Alternative route without an ECMP path
	iAE indicates Indirect Alternative ECMP path
	iB indicates Indirect Best route without ECMP path
	iBE indicates Indirect Best ECMP path
	dB indicates Direct Best route
	iAN indicates Indirect Alternative route not in hardware
	iAEN indicates Indirect Alternative ECMP route not in hardware
	iBN indicates Indirect Best route not in hardware
	iBEN indicates Indirect Best ECMP route not in hardware
	dBN indicates Direct Best route not in hardware
	iAU indicates Indirect Alternative Route Unresolved
	iAEU indicates Indirect Alternative ECMP Unresolved
	iBU indicates Indirect Best Route Unresolved
	iBEU indicates Indirect Best ECMP Unresolved
	dBU indicates Direct Best Route Unresolved
	iBF indicates Indirect Best route replaced by FTN
	iBEF indicates Indirect Best ECMP route replaced by FTN
	iBV indicates Indirect best IPVPN route
	iBEV indicates Indirect best ECMP IP VPN route
	iBVN indicates Indirect best IP VPN route not in hardware
	iBEVN indicates Indirect best ECMP IP VPN route not in hardware
Pref	Displays the preference.
NextHopVrfld	Specifies the VRF ID of the next-hop address.

Configure an OSPF Accept Policy

Perform the following procedure to create or configure an OSPF accept policy.

For more information on IS-IS accept policy filters for Fabric Connect for Layer 3 VSNs and IP Shortcuts, see Configuring Fabric Layer 3 Services for VOSS.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click Policy.

- 3. Click the **OSPF Accept** tab.
- 4. Click Insert.
- 5. Configure the parameters as required.
- 6. Click Insert.

OSPF Accept field descriptions

Use the data in the following table to use the **OSPF Accept** tab.

Name	Description
AdvertisingRtr	Specifies the routing ID of the advertising router.
Enable	Enables or disables the advertising router.
	You can also enable or disable advertising in the OSPF Accept tab of the Policy dialog box by clicking in the field and selecting enable or disable from the menu. The default is disable.
MetricType	Specifies the OSPF external type. This parameter describes which types of OSPF ASE routes match this entry.
	Any means match either ASE type 1 or 2
	Type1 means match any external type 1
	Type2 means match any external type 2
	You can also select your entry in the OSPF Accept tab of the Policy dialog box by clicking in the field and selecting any, type1, or type2 from the menu. The default is any.
PolicyName	Specifies the name of the OSPF in filter policy.
	Click the ellipsis button and choose from the list in the Policy Name dialog box. To clear an entry, use the ALT key.

Configuring inbound/outbound filtering policies on a RIP interface

About this task

Configure inbound filtering on a RIP interface to determine whether to learn a route on a specified interface and to specify the parameters of the route when it is added to the routing table. Configure outbound filtering on a RIP interface to determine whether to advertise a route from the routing table on a specified interface and to specify the parameters of the advertisement.

The port on which the multimedia filter is enabled becomes a DIFFSERV access port.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click Policy.
- 3. Click the RIP In/Out Policy tab.

- 4. In the desired row, double-click the **InPolicy** or **OutPolicy** column.
- 5. Select a preconfigured In/Out policy and click **OK**.

RIP In/Out Policy field descriptions

Use the data in the following table to use the RIP In/Out Policy tab.

Name	Description	
Address	Specifies the IP address of the RIP interface.	
Interface	Specifies the internal index of the RIP interface.	
InPolicy	Specifies the policy name used for inbound filtering on this RIP interface. This policy determines whether to learn a route on this interface and specifies the parameters of the route when it is added to the routing table.	
OutPolicy	Specifies the policy name used for outbound filtering on this RIP interface. This policy determines whether to advertise a route from the routing table on this interface and specifies the parameters of the advertisement.	

Deleting inbound/outbound filtering policies on a RIP interface

About this task

Delete a RIP In/Out policy when you no longer want to learn a route on a specified interface or advertise a route from the routing table on a specified interface.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click Policy.
- 3. Click the RIP In/Out Policy tab.
- 4. In the desired row, double-click the **InPolicy** or **OutPolicy** column for the policy you want to delete.
- 5. In the **InPolicy** or **OutPolicy** dialog box, press CTRL and then, click the policy you want to delete.
- 6. Click OK.

The policy is deleted and you are returned to the RIP In/Out Policy tab.

7. Click Apply.

Chapter 8: Routed Split MultiLink Trunking

This section provides conceptual information and procedures to configure Routed Split MultiLink Trunking (RSMLT) using Command Line Interface (CLI) and Enterprise Device Manager (EDM).

RSMLT

Table 38: RSMLT for IPv4 product support

Feature	Product	Release introduced	
For configuration details, see Configuring IPv4 Routing for VOSS.			
IPv4 RSMLT	VSP 4450 Series	VOSS 4.1	
	VSP 4900 Series	VOSS 8.1	
	VSP 7200 Series	VOSS 4.2.1	
	VSP 7400 Series	VOSS 8.0	
	VSP 8200 Series	VSP 8200 4.0	
	VSP 8400 Series	VOSS 4.2	
	VSP 8600 Series	VSP 8600 4.5	
	XA1400 Series	Not Supported	

In many cases, core network convergence time depends on the length of time a routing protocol requires to successfully converge. Depending on the specific routing protocol, this convergence time can cause network interruptions that range from seconds to minutes.

Routed Split MultiLink Trunking (RSMLT) permits rapid failover for core topologies by providing an active-active router concept to core SMLT networks.

RSMLT scenarios include SMLT triangles, squares, and SMLT full-mesh topologies, with routing enabled on the core VLANs.

Routing protocols include the following:

- IP Unicast Static Routes
- RIP1
- RIP2
- OSPF

BGP

In the event of core router failures, RSMLT manages packet forwarding, thus eliminating dropped packets during the routing protocol convergence.

SMLT and RSMLT Operation in Layer 3 Environments

<u>Figure 18: SMLT and RSMLT in Layer 3 environments</u> on page 250 shows a typical redundant network example with user aggregation, core, and server access layers. To minimize the creation of many IP subnets, one VLAN (VLAN 1, IP subnet A) spans all wiring closets.

SMLT provides the loop-free topology and forwards all links for VLAN 1, IP subnet A.

The aggregation layer switches are configured with routing enabled and provide active-active default gateway functionality through RSMLT.

After you enable RSMLT on a VLAN (on both aggregation devices), the cluster devices simply inform each other (over vIST messaging) of their physical IP and MAC on that VLAN. Thereafter, the two cluster devices take mutual ownership of their IP addresses on that VLAN. This action means each cluster device routes IP traffic that is directed to the physical MAC of the IP or the physical MAC of the peer IP on that VLAN, and when one of them is down the other cluster device:

- Replies to ARP requests for both the IP and the peer IP on that VLAN
- Replies to pings to the IP and the peer IP on that VLAN

In this case, routers R1 and R2 forward traffic for IP subnet A. RSMLT provides both router failover and link failover. For example, if the Split MultiLink Trunk link between R2 and R4 is broken, the traffic fails over to R1 as well.

For IP subnet A, VRRP with a backup master can provide the same functionality as RSMLT, as long as no additional router is connected to IP subnet A.

RSMLT provides superior router redundancy in core networks (IP subnet B), where OSPF is used for the routing protocol. Routers R1 and R2 provide router backup for each other, not only for the edge IP subnet A, but also for the core IP subnet B. Similarly routers R3 and R4 provide router redundancy for IP subnet C and also for core IP subnet B.

Router R1 Failure

The following figure shows SMLT and RSMLT in Layer 3 environments.

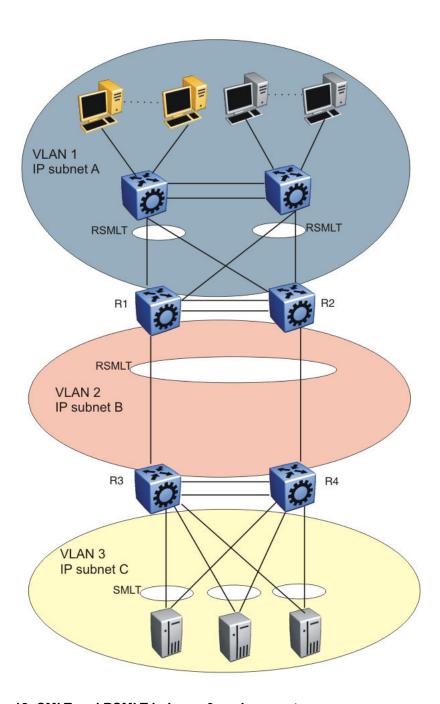


Figure 18: SMLT and RSMLT in Layer 3 environments

R3 and R4 both use R1 as their next hop to reach IP subnet A. Even though R4 sends the packets to R2, they are routed directly at R2 into subnet A. R3 sends its packets to R1 and they are also sent directly into subnet A. After R1 fails, all packets are directed to R2, with SMLT. R2 still routes for R2 and R1. After OSPF convergence, the routing tables in R3 and R4 change their next hop to R2 to reach IP subnet A. You can configure the hold-up timer (that is, for the amount of time R2 routes for R1 in a failure) for a time period greater than the routing protocol convergence, you can configure it as indefinite (that is, the members of the pair always route for each other).

Use an indefinite hold-up timer value for applications that use RSMLT at the edge instead of VRRP.

Router R1 Recovery

When R1 restarts after a failure, it does not route traffic for R2, nor does it provide backup for R2, until the hold-down timer expires. Similar to VRRP, the hold-down timer value must be greater than the time the routing protocol requires to converge its tables.

During the hold-down interval, R1 routes traffic if the destination MAC of the packet is its own routable VLAN MAC. R1 bridges incoming traffic to R2 if the destination MAC of the packet is the routable VLAN MAC of R2. A temporary default route (one having a route preference equal to 4) that points to R2 is installed on R1. R1 uses this temporary route to forward traffic to R2 that it cannot route itself because of the incomplete routing table; the default route is not saved in the configuration file.

After the hold-down timer expires, the temporary default route that points to R2 is deleted; from this moment on, in addition to routing packets destined to itself, R1 starts routing packets for R2, as well.

RSMLT Network Design and Configuration

Because RSMLT is based on SMLT, all SMLT configuration rules apply. In addition, RSMLT is enabled on the SMLT aggregation switches for each VLAN. The VLAN must be a member of SMLT links and vIST's L2VSN. For more information about how to configure SMLT in a Layer 2 environment, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS.

The VLAN also must be routable (IP address configured) and you must configure an Interior Gateway Protocol (IGP) such as OSPF on all four routers, although it is independent of RSMLT. All routing protocols, even static routes, work with RSMLT.

The RSMLT pair switches provide backup for each other. As long as one of the two routers of an vIST pair is active, traffic forwarding is available for both next hops R1/R2 and R3/R4.

RSMLT Edge Support

The switch stores the peer MAC and IP address pair in its local configuration file and restores the configuration if the peer does not restore after a simultaneous restart of both RSMLT-peer switches.

The RSMLT edge support feature adds an enhancement whereby the peer MAC (for the IP on the VLAN) is committed to the config.cfg file after you use the save config command. If you power off both devices, and then power up only one of them, that single device can still take ownership of its peer IP on that VLAN even if it has not seen that peer switch since it started. This enhancement is necessary if you configure the peer (the device which is still down) IP as the default gateway in end stations.

If you enable RSMLT edge support, you must also ensure that the hold-up timer for RSMLT on those edge VLANs equals infinity (9999). This timer value ensures that if one cluster device fails, the remaining cluster device maintains ownership of the failed peer IP indefinitely.

The edge VLAN can be tagged over SMLT links, single attached links, or more SMLT links.

Important:

If you clear the peer information the device can stop forwarding for the peer.

RSMLT implementation does not use a virtual IP address but instead uses physical IP addresses for redundancy. At the same time, you can deploy RSMLT in either routed configurations, or edge configurations, where you previously used VRRP (and back-up master). Previously, if a power outage occurred or a shutdown of both switches within a dual core vIST pair, only one device came back up. Clients using the powered-off device IP/MAC as the default gateway lost connectivity to the network. In such a scenario, even with RSMLT enabled on the device, it cannot act as a backup for the peer as it was unaware of the peer IP or MAC address.

After both the dual core vIST switches come back, the vIST is operational. If an RSMLT peerenabled message is received from the peer, normal RSMLT operation occurs.

If the peer has either an IP or MAC change, you must save the configuration for the RSMLT edge support to operate correctly. However, if the vIST peer up message is not received (for example, if you do not enable RSMLT properly), and you enable the RSMLT edge support flag, the RSMLT hold-down timer starts and permits routing protocols to converge; during this time user operation can be affected. After the hold-down timer expires, saved peer information is picked up and the device starts to act as backup for the peer by adding the previously saved MAC and ARP records.

The hold-up timer starts and after this timer expires the previously added MAC and ARP records are deleted and the device stops acting as backup for the peer, as the peer is not running proper RSMLT for the VLAN. The RSMLT is a parameter for each VLAN, and therefore all affects are on an individual VLAN basis, not necessarily a global device. Edge support mode uses the local values of the hold-down timer (default value of 60 seconds) and hold-up timer (default value of 180 seconds).

RSMLT configuration using the CLI

Routed Split MultiLink Trunking (RSMLT) forwards packets in the event of core router failures, thus eliminating dropped packets during the routing protocol convergence.

Configuring RSMLT on a VLAN

Perform this procedure to configure RSMLT on each IP VLAN interface.

Before you begin

- You must enable the IP routing protocol on VLAN Layer 3 interfaces.
- VLANs with Layer 3 interfaces must also participate in Split MultiLink Trunking (SMLT).

About this task

Use the no operator to disable RSMLT: no ip rsmlt

To configure this value to the default value, use the default operator with this command.

Procedure

1. Enter VLAN Interface Configuration mode:

```
enable
configure terminal
interface vlan <1-4059>
```

2. Enable RSMLT on a VLAN:

ip rsmlt

Example

Enable RSMLT on a VLAN:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface VLAN 100
Switch:1(config-if)#ip rsmlt
```

Variable definitions

Use the data in the following table to use the ip rsmlt command.

Table 39: Variable definitions

Variable	Value
holddown-timer <0-3600>	Configures how long the RSMLT device does not participate in Layer 3 forwarding.
	The value is in the range 0 to 3600 seconds.
	To configure this value to the default value, use the default operator with this command. The default value is 60 seconds.
	Configure this value to be longer than the anticipated routing protocol convergence.
holdup-timer <0-3600 9999>	Configures how long the RSMLT device maintains forwarding for its peer.
	The value is in the range 0 to 3600 seconds or 9999. 9999 means infinity.
	To configure this value to the default value, use the default operator with this command. The default is 1800 seconds.

Showing IP RSMLT information

Show IP RSMLT information to view data about all RSMLT interfaces.

About this task



If you use the <code>show ip rsmlt</code> command after you delete an RSMLT, the RSMLT still shows until you restart the switch.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display RSMLT information using the following command:

show ip rsmlt {edge-support] [<local|peer>] [vrf WORD < 1-16>] [vrfids WORD < 0-512>]

Example

	Switch:1>enable Switch:1#show ip rsmlt					
		Ip Rsmlt Local In	fo - Glo			
VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
	150.0.0.12	b0:ad:aa:40:05:25 b0:ad:aa:40:05:28 b0:ad:aa:40:05:01	Enable	Up	60	180 180 180
VID	SMLT ID					
1000 1500 3000	50 50					
VID	IPv6	MAC	ADMIN	OPER	HDTMR	HUTMR
1000	100:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0		Enable	Up	60	180
1500	150:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:		Enable	Up	60	180
3000	30:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:0:	b0:ad:aa:40:05:01 :0/64	Enable	Up	60	180
VID	SMLT ID					
1000 1500 3000	50 50					
====		Ip Rsmlt Peer Inf				
VID	IP	MAC	ADMIN	OPER	HDTMR	HUTMR
1000	100.0.0.206	b0:ad:aa:41:7d:23	Enable	Up	60	180
1500	150.0.0.206	b0:ad:aa:41:7d:24	Enable	Up	60	180

VID	HDT REMAIN	HUT REMAIN	SMLT ID				
1000 1500	60 60		50 50				
VID	IPv6	MAC		ADMIN	OPER	HDTMR	HUTMR
1000	100:0:0:0:0				Up	60	180
1500	150:0:0:0:0 150:0:0:0:0 fe80:0:0:0:	:0:0:0/64	a:41:7d:24 41:7d24/128		Up	60	180
VID	HDT REMAIN	HUT REMAIN	SMLT ID				
1000	60	180	50				
1500	60	180	50				
Switc	h:1#						

Variable definitions

Use the information in the following command to use the **show ip rsmlt** command.

Table 40: Variable definitions

Variable	Value
edge-support	Displays the RSMLT edge-support and peer information
<local peer></local peer>	Specifies values for the local or peer device.
vrf WORD<1-16>	Displays IP routing for a VRF.
vrfids WORD<0-512>	Displays IP routing for a range of VRFs.

Use the following table to use the show ip rsmlt [<local|peer>] command output.

Table 41: Variable definitions

Variable	Value
VID	Indicates the VLAN ID.
IP	Indicates the IP address of the VLAN.
MAC	Indicates the MAC address assigned.
ADMIN	Indicates the administrative status of RSMLT on the VLAN.
OPER	Indicates the operational status of RSMLT on the VLAN.
HDTMR	Indicates the hold-down timer value in the range of 0 to 3600 seconds.
HUTMR	Indicates the hold-up timer value in the range of 0 to 3600 seconds or 9999. 9999 means infinity.

Variable	Value
HDT REMAIN	Indicates the time remaining of the hold-down timer.
HUT REMAIN	Indicates the time remaining of the hold-up timer.
SMLT ID	Indicates the Split MultiLink Trunk ID.

Configuring RSMLT edge support

Configure RSMLT edge support to store the RSMLT peer MAC/IP address-pair in its local config file, and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT-peer systems. If enabled, all peer MAC/IP information for all RSMLT-enabled VLANs are saved during next the save config command.

About this task

RSMLT edge support is disabled by default.



If you use the show ip rsmlt command after you delete an RSMLT, the RSMLT still displays until you restart the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable RSMLT-edge:

```
ip rsmlt edge-support
```

Use the no operator to disable RSMLT-edge: no ip rsmlt edge-support

3. Clear RSMLT peer information, and then delete the RSMLT peer address:

```
no ip rsmlt peer-address <1-4059>
```

4. Display RSMLT-edge status information:

```
show ip rsmlt edge-support
```

Example

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

Enable RSMLT-edge:

Switch:1(config) #ip rsmlt edge-support

Display RSMLT-edge status information:

Switch:1(config) #show ip rsmlt edge-support

Variable definitions

Use the data in the following table to use the no ip rsmlt peer-address command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.

RSMLT configuration using Enterprise Device Manager

Routed Split MultiLink Trunking (RSMLT) forwards packets in the event of core router failures, thus eliminating dropped packets during the routing protocol convergence.

Configuring RSMLT on a VLAN

Configure RSMLT on a VLAN to exchange Layer 3 information between peer nodes in a switch cluster.

Before you begin

- Enable an IP routing protocol on VLAN Layer 3 interfaces.
- Ensure VLANs with Layer 3 interfaces participate in Split MultiLink Trunking (SMLT).

About this task

Use the following procedure to configure RSMLT using EDM

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANs.
- Click the Basic tab.
- 4. Select a VLAN.
- 5. Click IP.
- 6. Click the **RSMLT** tab.
- 7. Select Enable.
- 8. In the **HoldDownTimer** field, type a hold-down timer value.
- 9. In the **HoldUpTimer** field, type a holdup timer value.

10. Click Apply.

RSMLT field descriptions

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

Name	Description
Enable	Enables RSMLT. The default is disabled.
HoldDownTimer	Defines how long the recovering or restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address.
	The range of this value is from 0 to 3600 seconds. The default is 60.
	If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer if the peer is down. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.
	If you disable RSMLT on a VLAN, non default values for this field do not save across restarts.

Viewing and editing RSMLT local information

About this task

Perform the following procedure to view and edit RSMLT local VLAN information.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click RSMLT.
- 3. Click the Local tab.
- 4. Configure the parameters as required.
- 5. Click Apply.

Local field descriptions

Use the data in the following table to use the Local tab.

Name	Description
IfIndex	IP interface identification.
VlanId	Specifies the VLAN ID of the chosen VLAN.
lpAddr	Specifies the IP address on the RSMLT VLAN.
MacAddr	Specifies the MAC address of the selected VLAN.

Name	Description
Enable	Displays the RSMLT operating status as enabled or disabled.
OperStatus	Displays the RSMLT operating status as either up or down. The default is down.
HoldDownTimer	Defines how long the recovering/restarting system remains in a non- Layer 3 forwarding mode for the peer router MAC address.
	The range of this value is from 0 to 3600 seconds. The default is 60.
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 1800.
Smltld	Specifies the ID range for the SMLT. A valid range is 1 to 512.
Vrfld	Identifies the VRF.
VrfName	Indicates the VRF name.

Viewing RSMLT peer information

About this task

Perform this procedure to view and edit RSMLT peer information.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click RSMLT.
- 3. Click the **Peer** tab.

Peer field descriptions

Use the following table to use the **Peer** tab.

Name	Description
IfIndex	IP interface identification.
VlanId	Specifies the VLAN ID of the chosen VLAN.
lpAddr	Specifies the IP address on the RSMLT VLAN.
MacAddr	Specifies the MAC address of the selected VLAN.
Enable	Displays the RSMLT operating status as enabled or disabled.
OperStatus	Displays the RSMLT operating status as either up or down. The default is down.
HoldDownTimer	Defines how long the recovering/restarting system remains in a non-Layer 3 forwarding mode for the peer router MAC address.
	The range of this value is from 0 to 3600 seconds. The default is 0.

Name	Description
HoldUpTimer	Defines how long the RSMLT system maintains forwarding for its peer.
	The value is a range from 0 to 3600 seconds or 9999. 9999 means infinity. The default is 0.
HoldDownTimeRemaining	Displays the time remaining of the HoldDownTimer. The default is 0.
HoldUpTimeRemaining	Displays the time remaining of the HoldUpTimer. The default is 0.
Smitld	Specifies the ID range for the Split MultiLink Trunk. A valid range is 1 to 32.
Vrfld	Identifies the VRF.
VrfName	Indicates the VRF name.

Enabling RSMLT Edge support

Enable RSMLT Edge support to store the RSMLT peer MAC and IP address-pair in the local configuration file and restore the configuration if the peer does not restore after a simultaneous restart of both RSMLT peer systems.

The default is disabled.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click RSMLT.
- Click the Globals tab.
- 4. Select EdgeSupportEnable.
- 5. Click Apply.

Viewing RSMLT edge support information

About this task

View RSMLT edge support information to verify the RSMLT peer MAC/IP address-pair in its local configuration file and restore the configuration if the peer does not restore it after a simultaneous restart of both RSMLT-peer systems.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click RSMLT.
- 3. Click the **Edge Peers** tab.

Edge Peers field descriptions

Use the data in the following table to use the **Edge Peers** tab fields.

Name	Description
VlanId	Specifies the VLAN ID of the chosen VLAN.
PeerlpAddress	Specifies the peer IP address.
PeerMacAddress	Specifies the peer MAC address.
PeerVrfld	Identifies the Peer VRF.
PeerVrfName	Specifies the Peer VRF name.

Chapter 9: Virtual Router Redundancy Protocol

Table 42: Virtual Router Redundancy Protocol product support

Feature	Product	Release introduced			
For configuration details, see Configuring IPv4 Routing for VOSS.					
Virtual Router Redundancy	VSP 4450 Series	VSP 4000 4.0			
Protocol (VRRP)	VSP 4900 Series	VOSS 8.1			
	VSP 7200 Series	VOSS 4.2.1			
	VSP 7400 Series	VOSS 8.0			
	VSP 8200 Series	VSP 8200 4.0			
	VSP 8400 Series	VOSS 4.2			
	VSP 8600 Series	VSP 8600 4.5			
	XA1400 Series	VOSS 8.0.50			
VRRPv3 for IPv4	VSP 4450 Series	VOSS 5.1			
	VSP 4900 Series	VOSS 8.1			
	VSP 7200 Series	VOSS 5.1			
	VSP 7400 Series	VOSS 8.0			
	VSP 8200 Series	VOSS 5.1			
	VSP 8400 Series	VOSS 5.1			
	VSP 8600 Series	VSP 8600 6.1			
	XA1400 Series	VOSS 8.0.50			

VRRP Fundamentals

Because end stations often use a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks.

The Virtual Router Redundancy Protocol (VRRP) (RFC 2338) eliminates the single point of failure that can occur when the single static default gateway router for an end station is lost. VRRP introduces a virtual IP address (transparent to users) shared between two or more routers that

connect the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides dynamic default gateway redundancy in the event of failover.

The VRRP router that controls the IP addresses associated with a virtual router is the primary router and it forwards packets to these IP addresses. The election process provides a dynamic transition of forwarding responsibility if the primary router becomes unavailable.

Note:

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

In the following figure, the first three hosts install a default route to the R1 (virtual router 1) IP address and the other three hosts install a default route to the R2 (virtual router 2) IP address.

This configuration not only shares the load of the outgoing traffic, but it also provides full redundancy. If either router fails, the other router assumes responsibility for both addresses.

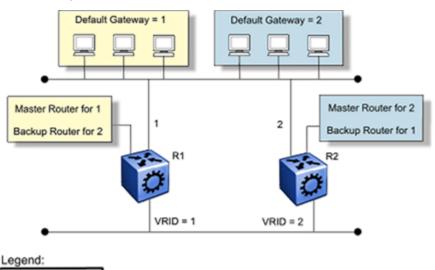


Figure 19: Virtual Router Redundancy Protocol configuration

For information about the number of supported VRRP interfaces, see the scaling information in Release Notes for VOSS.

The following terms are specific to VRRP:

Switch

VRRP router	a router running the VRRP protocol
-------------	------------------------------------

Virtual router an abstract object acting as the default router for one or more hosts, consisting

of a virtual router ID and a set of addresses

Primary IP an IP address selected from the real addresses and used as the source address address of packets sent from the router interface (The virtual primary router sends VRRP advertisements using this IP address as the source.)

Virtual primary router

the router that assumes responsibility to forward packets sent to the IP address

associated with the virtual router and answer ARP requests for these IP

addresses

Virtual router backup

the virtual router that becomes the primary router if the current primary router

When a VRRP router is initialized it sends a VRRP advertisement. The VRRP router also broadcasts a gratuitous ARP request that contains the virtual router MAC address for each IP address associated with the virtual router. The VRRP router then transitions to the controlling state.

In the controlling state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. The VRRP router responds to ARP requests for these IP addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts only packets addressed to IP addresses associated with the virtual router, the router transitions to the backup state to ensure that all Layer 2 switches in the downstream path relearn the new origin of the VRRP MAC addresses.

In the backup state, a VRRP router monitors the availability and state of the primary router. The backup router does not respond to ARP requests and must discard packets with a MAC address equal to the virtual router MAC address. The backup router does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, the backup router transitions back to the initialize state. If the primary router goes down, the backup router sends the VRRP advertisement and ARP request described in the preceding paragraph and transitions to the controlling state.

If an advertisement timer becomes active, the router sends an advertisement. If an advertisement is received with a 0 priority, the router sends an advertisement. The router transitions to the backup state in the following situations:

- If the priority is greater than the local priority
- If the priority is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address

Otherwise, the router discards the advertisement. If a shutdown occurs, the primary router sends a VRRP advertisement with a priority of 0 and transitions to the initialize state.

Critical IP Address

Within a VRRP VLAN, one link can go down while the remaining links in the VLAN remain operational. Because the VRRP VLAN continues to function, a virtual router associated with that VLAN does not register a master router failure.

As a result, if the local router IP interface connecting the virtual router to the external network fails, this does not automatically trigger a master router failover.



Note:

In this context, local implies an address from the same VRF as the IP interface where VRRP is being configured.

The critical IP address resolves this issue. If the critical IP address fails, it triggers a failover of the master router.

You can specify the local router IP interface uplink from the VRRP router to the network as the critical IP address. This ensures that, if the local uplink interface fails, VRRP initiates a master router failover to one of the backup routers.

In VRRP, the local network uplink interface on router 1 is shown as the critical IP address for router 1. As well, the same network uplink is shown as the critical IP address for router 2. Router 2 also requires a critical IP address for cases in which it assumes the role of the master router.

With the support of VRRP and the critical IP interface linked to VRRP, you can build reliable small core networks that provide support for converged applications, such as voice and multimedia.



A Brouter port with a VLACP Critical IP address in a VRRP is supported.

VRRP and SMLT

The standard implementation of VRRP supports only one active master device for each IP subnet. All other VRRP interfaces in a network are in backup mode.

A deficiency occurs when VRRP-enabled switches use Split MultiLink Trunking (SMLT). If VRRP switches are aggregated into two Split MultiLink Trunk switches, the end host traffic is load-shared on all uplinks to the aggregation switches (based on the Multilink Trunk traffic distribution algorithm).

However, VRRP usually has only one active routing interface enabled. All other VRRP routers are in backup mode. Therefore, all traffic that reaches the backup VRRP router is forwarded over the vIST towards the master VRRP router. In this case, the vIST does not have enough bandwidth to carry all the aggregated traffic.

To resolve this issue, assign the backup router as the backup master router. The backup master router can actively load-share the routing traffic with a master router.

When the backup master router is enabled, the incoming host traffic is forwarded over the SMLT links as usual. When the backup master router is configured along with the critical IP interface and the critical IP interface goes down, the VRRP router transitions to be the backup router with the backup master state down. In this state, the VRRP router does not forward traffic.

VRRP Fast Hello Timers

You can configure the advertisement time interval (in seconds) between sending advertisement messages. This interval permits fast network convergence with standardized VRRP failover. However, losing connections to servers for more than a second can result in missing critical failures. Customer network uptime in many cases requires faster network convergence, which means network problems must be detected within hundreds of milliseconds. Fast Advertisement Enable and Fast Advertisement Interval meet these requirements

Fast Advertisement Enable acts like a toggle device for the Advertisement Interval and the Fast Advertisement Interval. When Fast Advertisement Enable is enabled, the Fast Advertisement Interval is used instead of the Advertisement Interval.

The Fast Advertisement Interval is similar to the Advertisement Interval parameter except for the unit of measure and the range. The Fast Advertisement Interval is expressed in milliseconds and the range is from 200 to 1000 milliseconds. This unit of measure must be in multiples of 200 milliseconds, otherwise an error appears.

When you enable the fast advertisement interval, VRRP can communicate with other switch ports and networking products that have the same configuration.

Handling of IPv4 Layer 2 Unicast Packets at VRRP Backup Master

For VSP 8600 Series 6.2 and VOSSVOSS 7.1 releases and later, the handling of IPv4 Layer 2 unicast packets (for example, ARP Request/Reply) with the destination MAC as VRRP MAC has been modified on Backup-Master. These packets are now handled by VRRP Master only.

The Backup-Master forwards all IPv4 Layer 2 unicast packets to the Master and the Master VRRP sends an ARP reply only.

Processing of IP unicast packets (for example, ICMP packets to VRRP IP) or IPv4 routed packets (with destination MAC as VRRP MAC) on VRRP Backup-Master stays the same. For example, the VRRP Backup-Master replies to ICMP requests and routes Layer 3 routed packets to the destination and does not forward these packets to the Master when they arrive at the Backup-Master.

To reflect the above changes, the VRRP MAC entry on the Backup-Master now points to the Master instead of itself, and the ARP entry for VRRP IP on the backup-master points to local.

VRRP guidelines

VRRP guidelines

VRRP provides another layer of resiliency to your network design by providing default gateway redundancy for end users. If a VRRP-enabled router that connects to the default gateway fails, failover to the VRRP backup router ensures no interruption for end users who attempt to route from their local subnet.

Only the VRRP Master router forwards traffic for a given subnet. The backup VRRP router does not route traffic destined for the default gateway.

To allow both VRRP switches to route traffic, the switch software has an extension to VRRP, the BackupMaster, that creates an active-active environment for routing. If you enable BackupMaster on the backup router, the backup router no longer switches traffic to the VRRP Master. Instead the BackupMaster routes all traffic received on the BackupMaster IP interface according to the switch routing table.

Figure 20: VRRP with BackupMaster

Stagger VRRP instances on a network or subnet basis. The following figure shows the VRRP Masters and BackupMasters for two subnets. For more information about how to configure VRRP using the Command Line Interface (CLI) and Enterprise Device Manager (EDM), see VRRP configuration using the CLI on page 270 and VRRP configuration using EDM on page 285.

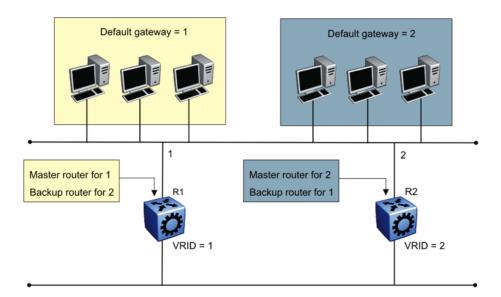


Figure 21: VRRP network configuration

The VRRP BackupMaster uses the VRRP standardized backup switch state machine. Thus, VRRP BackupMaster is compatible with standard VRRP.

Use the following best practices to implement VRRP:

- Do not configure the virtual address as a physical interface that is used on the routing switches. Instead, use a third address, for example:
 - Interface IP address of VLAN A on Switch 1 = x.x.x.2
 - Interface IP address of VLAN A on Switch 2 = x.x.x.3
 - Virtual IP address of VLAN A = x.x.x.1

Note:

The switch software does not support a VRRP virtual IP address that is the same as the local physical address of the device.

- Configure the VRRP holddown timer with enough time that the Interior Gateway Protocol (IGP)
 routing protocol has time to update the routing table. In some cases, configuring the VRRP
 holddown timer to a minimum of 1.5 times the IGP convergence time is sufficient. For OSPF, it
 is recommended that you use a value of 90 seconds if you use the default OSPF timers.
- Implement VRRP BackupMaster for an active-active configuration (BackupMaster works across multiple switches that participate in the same VRRP domain).

- Configure VRRP priority as 200 to configure VRRP Master.
- Stagger VRRP Masters between switches in the core to balance the load between switches.
- If you implement VRRP Fast, you create additional control traffic on the network and also create a greater load on the CPU. To reduce the convergence time of VRRP, the VRRP Fast feature allows the modification of VRRP timers to achieve subsecond failover of VRRP. Without VRRP Fast, normal convergence time is approximately 3 seconds.
- Do not use VRRP BackupMaster and critical IP at the same time. Use one or the other.

VRRP and spanning tree

The switch can use one of two spanning tree protocols: Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP).

VRRP protects clients and servers from link or aggregation switch failures. Configure the network to limit the amount of time a link is out of service during VRRP convergence. The following figure shows two possible configurations of VRRP and spanning tree; configuration A is optimal and configuration B is not.

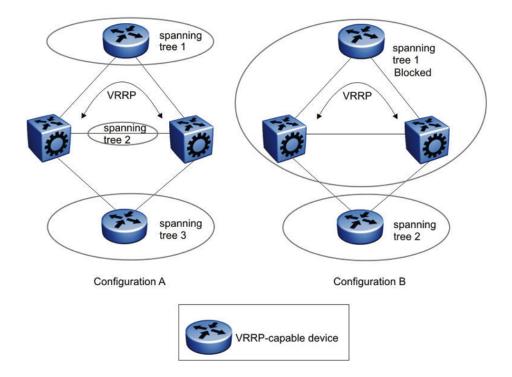


Figure 22: VRRP and STG configurations

In this figure, configuration A is optimal because VRRP convergence occurs within 2 to 3 seconds. In configuration A, three spanning tree instances exist and VRRP runs on the link between the two routers. Spanning tree instance 2 exists on the link between the two routers, which separates the link between the two routers from the spanning tree instances found on the other devices. All uplinks are active.

In configuration B, VRRP convergence takes between 30 and 45 seconds because it depends on spanning tree convergence. After initial convergence, spanning tree blocks one link (an uplink), so

only one uplink is used. If an error occurs on the uplink, spanning tree reconverges, which can take up to 45 seconds. After spanning tree reconvergence, VRRP can take a few more seconds to fail over.

VRRPv3

VRRPv3 is a combined protocol for both IPv4 and IPv6. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router that controls the IPv4 or IPv6 addresses associated with a virtual router is called the Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Backups wait for a Master and take ownership when the Master is no longer detected.

The election protocol provides dynamic failover in the forwarding responsibility when the Master is unavailable. VRRP for IPv4 gains a higher-availability default path without configuring dynamic routing or router discovery protocols on every end-host. VRRP for IPv6 gains a quick switch-over to Backup routers compared to the standard IPv6 Neighbor Discovery mechanisms.

Note:

The VRRP IPv6 link-local address must be the same for all VRRP routers sharing the same link and the same virtual router ID, that is, the same VRRP instance. It is the address the VRRP advertisements are sent from. Also, the Router Advertisement packets from the VRRP interface are transmitted using this address, so, when IPv6 Stateless Address Autoconfiguration is used, this address is added to host Default Router List and is used as a gateway.

The software supports VRRPv3 for IPv4 and VRRPv3 for IPv6. VRRPv3 for IPv6 is compliant to RFC 5798. The software also supports VRRPv2 for IPv4.

VRRPv3 guidelines

The switch also supports VRRPv2 for IPv4. If you configure VRRP IPv6 on an interface, it runs independently of the IPv4 version. Configure the version of the VRRP IPv4 on the interface before you configure any other IPv4 VRRP attributes. By default, the version is not configured to a particular value. However, when sourcing older configuration files that do not have the version saved, the router configures the version to VRRPv2 by default. If you change the version, all IPv4 configuration under that interface is automatically removed, and you are prompted for a confirmation before this operation.

Perform the CLI configuration through ip vrrp or ipv6 vrrp nodes; CLI commands for IPv4 are common for version 2 and version 3.

The following list identifies the features that make both IPv4 and IPv6 VRRPv3 features compliant to RFC 5798:

 Advertisement vs Fast-advertisement — Prior to RFC 5798, the minimum advertisement interval was 1 second, with Fast-advertisement a sub-second interval could be configured. When this feature is enabled, the VRRP ADVERTISEMENT packets are sent with type 7 instead of 1. With RFC 5798 the sub-second interval is standardised, and the switch sends all packets for VRRPv3 with type 1. The use of Fast-advertisement remains the same. VRRPv2 packets send with type 7, if Fast-advertisement is enabled.

- Add Master-advertisement-interval Prior to RFC 5798 compliance, all virtual routers on the same VLAN had the same Advertisement-Interval configured. RFC 5798 states that you can use different Advertisement Intervals on the Master and Backup. On the Master, the Masteradvertisement-interval and the Advertisement-Interval have the same value. On the Backup. the Master-advertisement-interval is used to calculate the timers, and the locally configured Advertisement-Interval is ignored until the Backup transitions to Master. The Masteradvertisement-interval value is put in the advertisement packet type sent by the Master
- Transition to master as specified in RFC 5798 Prior to RFC 5798, if a Backup receives an advertisement with a lower priority (or same priority but lower IP), it immediately sends its own advertisement and transitions to Master, However, RFC 5798 states that such packets must be discarded, which means it will transition to Master after the Master_Down_Timer expires
- Add skew-time RFC 5798 states that skew-time is calculated depending on the priority, and Master-advertisement-interval assures that the Backup with highest priority sends the first advertisement when the Master goes down

```
Skew time is calculated using the formula: (((256 - priority) *
Master Adver Interval) / 256).
```

 Add preempt-mode — Preempt-mode is different from the ipv6 vrrp <vrid> action preempt command, which is an operational command issued when you want to stop the holddown timer. RFC 5798 states that preempt-mode should be set to false when you do not want a higher priority Backup to transition to Master. By default, it is set to true



Note:

Accept-mode is not fully implemented for IPv4 VRRPv3. You can only ping the virtual IP address, the same way as it is for IPv4 VRRPv2.

VRRP configuration using the CLI

One active master switch exists for each IP subnet. All other VRRP interfaces in a network are in backup mode.

Virtual Router Redundancy Protocol (VRRP) eliminates the single point of failure that can occur after the single static default gateway router for an end station is lost. VRRP introduces the concept of a virtual IP address shared between two or more routers connecting the common subnet to the enterprise network.



Note:

The VRRP virtual IP address cannot be same as the local IP address of the port or VLAN on which VRRP is enabled.

Important:

The switch, when it acts as a VRRP master, does not reply to Simple Network Management Protocol (SNMP) Get requests to the VRRP virtual interface address. However, if the switch acts as a VRRP master, and receives SNMP Get requests to its physical IP address, then it does respond.

The Simple Network Management Protocol (SNMP) provides facilities to manage and monitor network resources. An SNMP manager and agent communicate through the SNMP protocol. The manager sends queries and the agent responds. An SNMP Get request is a message that requests the values of one or more objects.

Note:

The VRRP IP address responds only to ICMP-based traceroute requests. It does not respond to UDP-based traceroute requests.

When you use the fast advertisement interval option to configure a master and backup device, you must enable the fast advertisement interval option on both systems for VRRP to work correctly. If you configure one device with the regular advertisement interval, and the other device with the fast advertisement interval, it causes an unstable state and drops advertisements.

Configuring VRRP on a port or a VLAN

About this task

Configure VRRP on a port or a VLAN to forward packets to the virtual IP addresses associated with the virtual router and customize the VRRP configuration.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
```

Note:

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

2. Configure a backup VRRP address:

```
ip vrrp address <1-255> <A.B.C.D>
```

3. Configure VRRP on a port:

```
ip vrrp <1-255> enable
```

4. Show the global VRRP configuration:

show ip vrrp

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Configure a backup VRRP address:

Switch:1(config-if)# ip vrrp address 28 192.0.2.1

Configure VRRP on a port:

Switch:1(config-if) # ip vrrp 28 enable

Show the global VRRP configuration:

Switch:1(config-if)# show ip vrrp

Variable definitions

Use the data in the following table to use the ip vrrp command.

Variable	Value
1-255	Specifies the number of the VRRP to create or modify.
action {none preempt}	Causes the virtual router to disregard the timer and transition to Master state immediately, provided the hold-down timer is running.
	Note:
	You can use this parameter only if the hold-down timer is active.
	To set this option to the default value, use the default operator with this command.
action {none preempt}	Enables the choice option to manually override the hold-down timer and force preemption.
	You can configurenone preempt to preempt the timer or configure it as none to allow the timer to keep working.
	To configure this option to the default value, use the default operator with this command.
address <1-255> <a.b.c.d></a.b.c.d>	Configures the IP address of the VRRP physical interface that forwards packets to the virtual IP addresses associated with the virtual router.
	A.B.C.D is the IP address of the master VRRP.
	Use the no operator to remove the IP address of the VRRP physical interface:no ip vrrp address <1-255> <a.b.c.d>. To configure this option to the default value, use the default operator with this command.</a.b.c.d>
adver-int <1-255>	Configures the time interval between sending VRRP advertisement messages. The range is between 1 and 255

Variable	Value
	seconds. This value must be the same on all participating routers. The default is 1.
	To configure this option to the default value, use the default operator with this command.
backup-master enable	Enables the VRRP backup master.
	Use the no operator to disable the VRRP backup master: no ip vrrp <1-255> backup-master enable. To configure this option to the default value, use the default operator with this command.
	When backup master functionality is enabled, the VRRP router will IP-forward packets destined to the VRRP MAC even when the router is not the VRRP Master.
	Important:
	Do not enable backup master if you enable critical IP.
critical-ip-addr <a.b.c.d></a.b.c.d>	Configures the critical IP address for VRRP.
	A.B.C.D is the IP address on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.
	Note:
	In this context, <i>local</i> implies an address from the same VRF as the IP interface where VRRP is being configured.
critical-ip enable	Enables the critical IP address option.
	Use the no operator to disable the critical IP address option: no ip vrrp <1-255> critical-ip enable. To configure this option to the default value, use the default operator with this command.
	Important:
	Do not enable Critical IP if backup master is enabled.
enable	Enables VRRP on the port.
	Use the no operator to disable VRRP on the port: no ip vrrp <1-255> enable. To configure this option to the default value, use the default operator with this command.
fast-adv enable	Enables the Fast Advertisement Interval. The default is disabled.
	Use the no operator to disable VRRP on the port: no ip vrrp <1-255> fast-adv enable. To configure this option to the default value, use the default operator with this command.
fast-adv-int <200-1000>	Configures the Fast Advertisement Interval, the time interval between sending VRRP advertisement messages.

Variable	Value
	200-1000 is the range in milliseconds, and must be the same on all participating routers. The default is 200. You must enter values in multiples of 200 milliseconds.
	To configure this option to the default value, use the default operator with this command.
holddown-timer <0-21600>	Specifies the time interval (in seconds) for which the transition of virtual router to Master state is delayed in case of the following conditions:
	The VRRP hold-down timer runs only when the VRRP virtual router transitions from initialization to backup to master. This occurs only on a system startup.
	The VRRP hold-down timer does not run if the amount of time passed since VRRP virtual router initialization is greater than preset hold-down time. In such a case, VRRP virtual router transitions to Master happens irrespective of the hold-down timer.
	The VRRP hold-down timer also applies to the VRRP BackupMaster feature.
	0-21600 is the time interval range (in seconds). To configure this option to the default value, use the default operator with this command. The default value for hold-down timer is 0, that is, the timer is disabled by default.
priority <1-255>	Configures the port VRRP priority.
	1-255 is the value used by the VRRP router. The default is 100. Assign the value 255 to the router that owns the IP address associated with the virtual router.
	To configure this option to the default value, use the default operator with this command.

Showing VRRP information

About this task

Show VRRP port or VLAN information to view configuration details and operational status.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display basic VRRP configuration information about the specified port, all ports, or the VLAN:

```
show ip vrrp address [vrid <1-255>] [addr <A.B.C.D>] [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

3. Displaying the VRRPv3 configuration:

```
show ip vrrp address version <2-3>
```

4. Displaying version based VRRP configuration for the specified VRF:

```
show ip vrrp address vrf WORD<1-16> version <2-3>
```

5. Displaying version based VRRP configuration for the specified VRF ID:

show ip vrrp address vrfids WORD<0-512> version <2-3>

Example

Switch:1	Switch:1#show ip vrrp address										
		VRRP	Info - GlobalRouter		=====		=====	=====	====	=====	====
======			MA C	===:		===	=====		7 DT	VEDCT	====
			00:00:5e:00:01:03 00:00:5e:00:01:02								
2 out of	2 Total	Num of VRRP Addr	ess Entries display	ed.							
	D /11	MA GEED				T-73-7	OD THE	031 TD	(=====	D	
VRRP ID	P/V 	MASTER 	UP TIME		 HLD D		CRITI	CAL IP	(ENA	RTED)	VERSION
3 2	3 1/1	30.30.30.18 20.20.20.18	0 day(s), 00:08:53 0 day(s), 00:02:01		0		0.0.0	.0		(No) (No)	2 3
2 out of	2 Total	Num of VRRP Addr	ess Entries display	ed.							
VRRP ID	P/V	BACKUP MASTER	BACKUP MASTER STAT	E	FAST	ADV	(ENAB	LED)	V	ERSION	
3 2		disable disable			200 200		(NO) (NO)			2	
2 out of	2 Total	Num of VRRP Addr	ess Entries display	ed.							

Variable definitions

Use the data in the following table to use the show ip vrrp address command.

Variable	Value
addr <a.b.c.d></a.b.c.d>	Specifies the physical local address of the master VRRP.
vrf WORD<1-16>	Specifies the name of the VRF.
vrid <1-255>	Specifies a unique integer value that represents the virtual router ID in the range 1–255. The virtual router acts as the default router for one or more assigned addresses.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
version <2–3>	Specifies the VRRP version (2 or 3) to be shown.

Use the data in the following table to interpret the show ip vrrp address command output.

Table 43: Field descriptions

Name	Description		
ADV	Indicates the Advertisement Interval, in seconds, between sending advertisement messages.		
BACKUP MASTER	Indicates if the Backup-Master feature is disabled or enabled.		
BACKUP MASTER STATE	Indicates if the Backup-Master is up. If the switch is in Master state but Backup-Master is enabled, then the BACKUP MASTER STATE will be down.		
CONTROL	Indicates the virtual router function. Configure the value to enabled to transition the state of the router from initialize to backup. Configure the value to disabled to transition the router from master or backup to initialize.		
CRITICAL IP	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.		
CRITICAL IP (ENABLED)	Indicates if the critical IP feature is enabled.		
FAST ADV	Indicates the Fast Advertisement Interval, in milliseconds, between sending advertisement messages. When the Fast Advertisement Interval is enabled, the Fast Advertisement Interval is used instead of the regular advertisement interval.		
FAST ADV (ENABLED)	Indicates the state of fast advertisement.		
HLD DWN	Specifies the time interval (in seconds) the Hold-down timer has until it expires. If the value is 0, it means the Hold-down timer is not running. This timer will delay the transition from Backup to Master only on a system startup (the VRRP comes from INIT to Backup and determines it should become Master).		
	The VRRP hold-down timer runs when the system transitions from initialization to backup to master. This occurs only on a system startup		
	 The VRRP hold-down timer does not run under the following condition: In a nonstartup condition, the backup system becomes master after the Master Downtime Interval (3 * hello interval), if the master virtual router goes down 		
	The VRRP hold-down timer also applies to the VRRP BackupMaster feature		
IP	Indicates the assigned IP addresses that a virtual router backs up.		
MAC	Indicates the virtual MAC address of the virtual router in the format 00-00-5E-00-01- <vrrpid>, where the first three octets consist of the IANA OUI; the next two octets indicate the address block of the VRRP protocol; and the remaining octets consist of the vrrpid.</vrrpid>		
MASTER	Indicates the master router real (primary) IP address.		

Name	Description
PRIO	Indicates the priority for the virtual router with respect to other virtual routers that are backing up one or more associated IP addresses. Higher values indicate higher priority.
	A priority of 255 cannot be configured and it is set for the VRRP router that has the same IP as the physical IP addresses (is Address Owner).
P/V	Indicates the P(ort)/V(lan) on which the VRRP was configured.
STATE	Indicates the current state of the virtual router.
	initialize—waiting for a startup event
	backup—monitoring the state or availability of the master router
	master—forwarding IP addresses associated with this virtual router.
UP TIME	Indicates the time interval since this virtual router exited the INIT state.
VRRP ID	Indicates the virtual router ID on a VRRP router.
VERSION	Indicates the VRRP version.

Showing extended VLAN VRRP

Perform this procedure to display the extended VRRP configuration for all VLANs or a specified VLAN on the device.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Show the extended VRRP configuration for all VLANs on the device or for the specified VLAN:

show ip vrrp interface vlan [<1-4059>] [portList] verbose [vrf WORD<1-16>] [vrfids WORD<0-512>]

Example

Variable definitions

Use the data in the following table to use the show ip vrrp interface vlan command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
portList	Specifies the slot or port number of a range of ports.
vrf WORD<1-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.

Use the data in the following table to use the show ip vrrp interface vlan [<1-4059>] [portList] verbose [vrf WORD<1-16>] [vrfids WORD<0-512>] command output.

Variable	Value
VLAN ID	Indicates the VLAN ID.
STATE	Indicates the current state of the virtual router.
	initialize—waiting for a startup event
	backup—monitoring the state or availability of the master router
	master—forwarding IP addresses associated with this virtual router
CONTROL	Indicates the virtual router function. Configure the value to enabled to transition the state of the router from initialize to backup. Configure the value to disabled to transition the router from master or backup to initialize.
PRIORITY	Indicates the priority for the virtual router (for example, master election) with respect to other virtual routers that are backing up one or more associated IP addresses. Higher values indicate higher priority.
	A priority of 0, which you cannot configure, indicates that this router ceased to participate in VRRP and a backup virtual router transitions to become a new master.
	Use a priority of 255 for the router that owns the associated IP addresses.
MASTER IPDDR	Indicates the master router real (primary) IP address. The master IP address is listed as the source in the VRRP advertisement last received by this virtual router.
ADVERTISE INTERVAL	Indicates the time interval, in seconds, between sending advertisement messages. Only the master router sends VRRP advertisements.
CRITICAL IPADDR	Indicates the IP address of the interface that causes a shutdown event.
HOLDDOWN_TIME	Indicates the configured time (in seconds) that the system waits before it preempts the current VRRP master.

Variable	Value
ACTION	Indicates the trigger for an action on this VRRP interface. Options include none and preemptHoldDownTimer.
CRITICAL IP ENABLE	Indicates that a user-defined critical IP address is enabled. No indicates the use of the default IP address (0.0.0.0).
BACKUP MASTER	Indicates the state of designating a backup master router.
BACKUP MASTER STATE	Indicates the state of the backup master router.
FAST ADV INTERVAL	Indicates the time interval, in milliseconds, between sending Fast Advertisement messages. When the Fast Advertisement Interval is enabled, the Fast Advertisement Interval is used instead of the regular advertisement interval.
FAST ADV ENABLE	Indicates the Fast Advertisement Interval status.

Showing VRRP interface information

About this task

If you enter a virtual router ID or an IP address when showing VRRP interface information, the information appears only for that virtual router ID or for that interface.

Procedure

1. Enter Privileged EXEC mode:

enable

2. Display VRRPv3 information about the specified interface:

```
show ip vrrp interface version <2-3>
```

3. Display additional VRRPv3 information about the specified interface:

```
show ip vrrp interface verbose version <2-3>
```

4. Display VRRPv3 information for the specified VRF:

```
show ip vrrp interface vrf WORD<1-16> version <2-3>
```

5. Display VRRPv3 information for the specified virtual router:

```
show ip vrrp interface vrfids WORD<0-512> [version <2-3>]
```

Example

Swite	ch:1#show ip vrrp	inte	rface			
			Vlan Vr	rp		
VLAN ID	VRF NAME	VRRP ID	IP ADDRESS	VIRTUAL MAC ADDRESS	VERSION	
3	GlobalRouter	3	30.30.30.99	00:00:5e:00:01:03	2	
All 1	All 1 out of 1 Total Num of Vlan Vrrp displayed					

					Port	Vrr								
PORT	VRF NAME		WRRP	TP			VTR	TUAL ADDRESS		VEI	RSION	=====		=====
	Glok	oalRouter	2	20.20.2	20.99)	00:		:01:0					
Swit	ch:1#s	show ip vrrp	inter	face ve	rbose	:								
	=====		====			p Ext			====	====			======	=====
VLAN	VRRP	VRF NAME STAT					MAS	TER		ADVEI	RTISE	CRITI	CAL	VERSION
10 20	1 2	Global~ init Global~ init	d d	isable :	100		0.0	.0.0		1 1		0.0.0	.0	3 3
All	2 out	of 2 Total N	um of	Vlan V	rp E	Exten	ded 1	Entries	displ	ayed				
ID	ID	VRF NAME		TIM	C	IP ENAB	LE	MASTER	MAST STAT	ER I E	INTERV	/AL	ENABLE	
		GlobalRouter GlobalRouter												
All	2 out	of 2 Vlan Vr	rp Ex	tended 1	Entri	es d	ispl	ayed						
VLAN	VRRP	VRF	MAST	ER ADV	PRE	EMPT	PS:	EUDO-HEA	DER V	ERSI	ON			
ID	ID	NAME	INTE	RVAL (ms	MOI	Œ	CH	ECKSUM						
10 20	1 2	GlobalRouter GlobalRouter	1000		ena	bled	en.	abled	3 3					
All	2 out	of 2 Vlan Vr	rp Ex	tended 1	Entri	es d	ispl	ayed						
	=====		====			p Ext			====			====		=====
		P VRF NAME STA	TE C	ONTROL 1	PRIOF	RITY	MAS'	TER DDR		ADVEI INTEI				VERSION
1/2	3	Global~ ini												
NUM	ID	P VRF NAME		IIT	1E	IP ENAI	BLE	MASTER	MAS STA	TER TE	INTER	RVAL	ENABLE	
		GlobalRoute											disable	
PORT	VRRI TD	P VRF NAME	MAS	TER ADV	PF	REEMP'	r P	SEUDO-HE	ADER	VERS	ION			

Variable definitions

Use the data in the following table to use the show ip vrrp interface command.

Variable	Value
gigabitethernet {slot/port[-slot/port][,]}	Specifies to show the VRRP information of which interface.
verbose	Specifies to show all available information about the VRRP interfaces.
vlan	Specifies the VLAN that contains the VRRP.
vrf WORD<1-16>	Specifies the name of the VRF.
vrid <1-255>	Specifies a unique integer value that represents the virtual router ID in the range 1–255. The virtual router acts as the default router for one or more assigned addresses.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
version<2-3>	Specifies the VRRP version (2 or 3) configured.

Enabling ping to a virtual IP address

Use the following procedure to enable ping to a virtual IP address. The default is enabled.

Procedure

1. Enter VRRP Router Configuration mode:

```
enable
configure terminal
router vrrp
```

2. Enable ping to a virtual IP address:

```
ping-virtual—address enable [vrf WORD<1-16>]
default ping-virtual—address enable [vrf WORD<1-16>]
```

3. Disable ping to a virtual IP address:

```
no ping-virtual—address enable [vrf WORD<1-16>]
```

4. Display the configuration:

```
show ip vrrp [vrf WORD<1-16>]
```

Example

Variable definitions

Use the data in the following table to use the ping-virtual—address enable and show ip vrrp commands.

Variable	Value
enable	Enables ping to a virtual IP address.
vrf WORD<1–16>	Specifies the VRF.

Configuring VRRP notification control

Use the following procedure to enable VRRP notification control. The generation of SNMP traps for VRRP events is enabled, by default.

About this task

You can configure traps by creating SNMPv3 trap notifications, creating a target address to send the notifications, and specify target parameters. For more information about how to configure trap notifications, see <u>Troubleshooting VOSS</u>.

Procedure

1. Enter VRRP Router Configuration mode:

```
enable
configure terminal
router vrrp
```

2. Enable a trap for VRRP events:

```
send-trap enable [vrf WORD<1-16>]
```

3. Disable a trap for VRRP events:

```
no send-trap enable [vrf WORD<1-16>]
```

4. Configure a trap for VRRP events to the default:

```
default send-trap enable [vrf WORD<1-16>]
```

5. Display the configuration:

```
show ip vrrp [vrf WORD<1-16>]
```

Example

```
ping-virtual-address : enabled
send-trap : enabled
```

Variable definitions

Use the data in the following table to use the send-trap and show ip vrrp commands.

Variable	Value
enable	Enables generation of SNMP traps.
vrf WORD<1–16>	Configures the send-trap for a particular VRF.

Configuring VRRP version on an interface

About this task

Use the following command to configure the VRRP version on an interface.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```

Note:

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

2. Use the following command to configure the VRRP version:

```
ip vrrp version <2-3>
```

3. Use the following command to set the VRRP version to default:

```
default ip vrrp version
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Configure VRRP version for the specified interface:

```
Switch:1(config-if) # ip vrrp version 3
```

Variable definitions

Use the data in the following table to use the ip vrrp version command.

Variable	Value
version <2–3>	Configures the VRRP version (2 or 3) on the specified interface

Enabling IPv4 VRRP preempt-mode

You can configure VRRP to preempt the existing router. If a new VRRP router is added to the network with a higher priority than the existing routers, then the new router becomes the master. If preempt-mode is disabled, then the new router does not become a master, it transitions to master only when the current master is down, that is when it does not receive any advertisement packets from the current master. By default, preempt-mode is enabled.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```



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

2. Enter the following command:

```
ip vrrp <vrid> preempt-mode enable
```

3. Use the following command to set the preempt-mode to its default value:

```
default ip vrrp <vrid> preempt-mode
```

4. Use the following command to disable the preempt-mode:

```
no ip vrrp <vrid> preempt-mode enable
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface gigabitethernet 1/2
```

Enabling preempt-mode on interface 1/2:

```
Switch:1(config-if) # ip vrrp 1 preempt-mode enable
```

Variable definitions

Use the data in the following table to use the ip vrrp <vrid> command.

Variable	Value
preempt-mode enable	Enables preempt-mode for VRRPv3 for IPv4.
default ip vrrp <vrid> preempt-mode</vrid>	Sets the default preempt-mode value for VRRPv3 for IPv4.
no ip vrrp <vrid> preempt-mode enable</vrid>	Disables preempt-mode for VRRPv3 for IPv4.

VRRP configuration using **EDM**

One active master switch exists for each IP subnet. All other VRRP interfaces in a network are in backup mode.

If you have VRRP and IP routing protocols configured on the same IP physical interface, you cannot select the interface address as the VRRP virtual IP address (logical IP address). Use a separate dedicated IP address for VRRP.

To modify the behavior of the VRRP failover mechanism, use the hold-down timer to allow the router enough time to detect and update routes. The timer delays the preemption of the master over the backup, when the master becomes available. The hold-down timer has a default value of 0 seconds. Configure all of your routers to the identical number of seconds for the hold-down timer. In addition, you can manually force the preemption of the master over the backup before the delay timer expires.

Note:

The VRRP virtual IP address cannot be the same as the local IP address of the port or VLAN on which VRRP is enabled.

Important:

The switch, when it acts as a VRRP master, does not reply to Simple Network Management Protocol (SNMP) Get requests to the VRRP virtual interface address. However, if the switch acts as a VRRP master, and receives SNMP Get requests to its physical IP address, then it does respond.

The Simple Network Management Protocol (SNMP) provides facilities to manage and monitor network resources. An SNMP manager and agent communicate through the SNMP protocol. The manager sends queries and the agent responds. An SNMP Get request is a message that requests the values of one or more objects.

Note:

The VRRP IP address responds only to ICMP-based traceroute requests. It does not respond to UDP-based traceroute requests.

Before you begin

- Assign an IP address to the interface.
- · Enable VRRP globally.

Enabling VRRP global variables

About this task

Enable VRRP global variables to enable the VRRP function.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click VRRP.
- 3. Click the Globals tab.
- 4. Configure the required features.
- 5. Click Apply.

Globals field descriptions

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

Name	Description		
NotificationCntl	Indicates whether the VRRP-enabled router generates SNMP traps for events.		
	enabled—SNMP traps are generated		
	disabled—no SNMP traps are sent		
	The default is enabled.		
PingVirtualAddrEnable	Configures whether this device responds to pings directed to a virtual router IP address. The default is enabled.		

Modifying VRRP parameters for an interface

Before you begin

• You must enable VRRP on a brouter port or VLAN.

About this task

You can manage and configure VRRP parameters for the routing interface.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click VRRP.
- 3. Click the **Interface** tab.
- 4. Double-click the **HoldDownTimer** field, and enter the number of seconds for the timer.

The **HoldDownState** field displays active when the hold-down timer is counting down and preemption occurs. The field displays dormant when preemption is not pending. When the hold-down timer is active, the **HoldDownTimeRemaining** field displays the seconds remaining before preemption.

- 5. In the **Action** check box, select an option.
- 6. Click Apply.

Interface field descriptions

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

Name	Description
IfIndex	Specifies the index value that uniquely identifies the interface to which this entry is applicable.
Vrld	Specifies a number that uniquely identifies a virtual router on a VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
IpAddr	Specifies the assigned IP addresses that a virtual router is responsible for backing up.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Specifies the state of the virtual router interface:
	Initialize—waiting for a startup event
	Backup—monitoring availability and state of the master router
	Master—functioning as the forwarding router for the virtual router IP addresses.
Control	Specifies whether VRRP is enabled or disabled for the port (or VLAN). The default is enabled.
Priority	Specifies the priority value used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
AdvertisementInterval	Specifies the time interval (in seconds) between sending advertisement messages. The range is 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements. The default is 1.
MasterlpAddr	Specifies the IP address of the physical interface of the master virtual router that forwards packets sent to the virtual IP addresses associated with the virtual router.

Name	Description
VirtualRouterUpTime	Specifies the time interval (in hundredths of a second) since the virtual router was initialized.
Action	Lists options to override the delay timer manually and force preemption:
	none does not override the timer
	preemptHoldDownTimer preempts the timer
HoldDownTimer	Configures the amount of time (in seconds) to wait before preempting the current VRRP master.
HoldDownState	Indicates the hold-down state of this VRRP interface. If the hold-down timer is operational, this variable is set to active; otherwise, this variable is set to dormant.
HoldDownTimeRemaining	Indicates the amount of time (in seconds) left before the HoldDownTimer expires.
CriticallpAddr	Configures the critical IP address for VRRP.
	This command specifies an IP interface on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.
	Note:
	In this context, <i>local</i> implies an address from the same VRF as the IP interface where VRRP is being configured.
CriticallpAddrEnable	Configures the IP interface on the local router to enable or disable the backup. The default is disabled.
BackUpMaster	Enables the backup VRRP system traffic forwarding. The default is disabled.
BackUpMasterState	Indicates whether the backup VRRP system traffic forwarding is enabled or disabled. The default is disabled.
FasterAdvinterval	Configures the Fast Advertisement Interval between sending VRRP advertisement messages. The interval is between 200 and 1000 milliseconds, and you must enter the same value on all participating routers. The default is 200. You must enter the values in multiples of 200 milliseconds.
FasterAdvintervalEnable	Enables or disables the Fast Advertisement Interval. When disabled, the regular advertisement interval is used. The default is disable.

Configuring VRRP on a V3 interface

Perform this procedure to configure VRRP on a V3 interface on either a brouter port or a VLAN.

Before you begin

- · Assign an IPv4 address to the interface
- Enable routing globally
- Do not configure RSMLT on the VLAN

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click VRRP.
- 3. Click the V3 Interface tab.
- 4. Click Insert.
- 5. Beside the IfIndex field, click Port or VLAN.
- 6. Select a port or VLAN.
- 7. Click OK.
- 8. Type the virtual router ID.
- 9. Type the primary IP address.
- 10. Type the advertisement interval.
- 11. Click Insert.

Interface field descriptions

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

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
InetAddrType	Specifies the source network INET Address Type.
Vrld	Specifies a number that uniquely identifies a virtual router on a VRRP router.
PrimarylpAddr	Specifies the virtual address assigned to the VRRP.
VirtualMacAddr	Specifies the MAC address of the virtual router interface.
State	Shows the state of the virtual router interface. The possible states are • initialize—waiting for a startup event

Name	Description
	backup—monitoring availability and state of the master router
	master—functioning as the forwarding router for the virtual router IP addresses
Control	Displays whether VRRP is enabled or disabled for the port or VLAN.
Priority	Specifies the priority value used by this VRRP router. The value 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
Advinterval	Specifies the time interval, in seconds, between sending advertisement messages. The default is 1 second.
MasterlpAddr	Specifies the IP address of the physical interface of the Master's virtual router.
UpTime	Indicates the time interval since this virtual router exited the INIT state.
CriticallpAddr	Indicates the IP address of the interface that is critical to VRRP. If that IP interface is down, the VRRP state will transition to Backup, even if it has higher priority.
CriticallpAddrEnabled	Enables or disables the use of critical IP. When disabled, the VRRP ignores the availability of the address configured as critical IP. This address must be a local address. The default is disabled.
BackUpMaster	Uses the backup VRRP switch for traffic forwarding. This option reduces the traffic on the vIST. The default is disabled.
BackUpMasterState	Indicates if the Backup-Master is operational up. If the switch is in Master state but the Backup-Master is enabled, then the BACKUP MASTER STATE will be down.
FasterAdvIntervalEnabled	Enables or disables the fast advertisement interval. When disabled, the regular advertisement interval is used. The default is disable.
FasterAdvInterval	Configures the interval between VRRP advertisement messages. The default is 200.
	Enter the values in multiples of 200 milliseconds.
PreemptMode	Issued to preempt the existing router. If a new router is added to the network with its priority higher than the existing routers, then the new router becomes the master. The default is enabled.

Name	Description
Action	Lists options to override the hold-down timer manually and force preemption:
	none does not override the timer.
	preemptHoldDownTimer preempts the timer.
	This parameter applies only if the holddown timer is active.
HoldDownTimer	Configures the amount of time, in seconds, to wait before preempting the current VRRP master. The default is 0.
HoldDownTimeRemaining	Indicates the amount of time, in seconds, left before the HoldDownTimer expires.
MasterAdvinterval	On the VRRPv3 master, the master advertisement interval is same as the advertisement interval. On the VRRPv3 Backup, the master advertisement interval is set to the Advertisement configured on the Master (received in the packet).

Configuring VRRPv3 Checksum

Perform this procedure to configure VRRPv3 checksum on either a brouter port or a VLAN.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > IP**.
- 2. Click VRRP.
- 3. Click the V3 Checksum tab.
- 4. Click Insert.
- 5. In the **Interface** field, click **Port** or **Vlan**.
- 6. Select a type of checksum computation.
- 7. Select a VRRP version.
- 8. Click Insert.

V3 Checksum field descriptions

Use the data in the following table to use the V3 Checksum tab.

Name	Description
Interface	Shows VRRPv3 information about a specified interface.

Name	Description
rclpConflfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
ChkSumComputation	Specifies the type of checksum computation, with Pseudo Header or without Pseudo Header.
VrrpVersion	Specifies the VRRP version as unspecified, v2, or v3.

Configuring Fast Advertisement Interval on a port or a VRF instance

About this task

Configure the Fast Advertisement Interval to send VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.

Procedure

- 1. In the Device Physical View tab, select a port.
- 2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
- 3. Click IP.
- 4. Click the **VRRP** tab.
- 5. Click Insert.
- 6. In the Insert VRRP dialog box, enable FasterAdvIntervalEnable.
- 7. In the **FasterAdvinterval** field, enter a value. You must set this value using multiples of 200 milliseconds.
- 8. Click Insert.

Configuring Fast Advertisement Interval on a VLAN or a VRF instance

About this task

Configure the Fast Advertisement Interval to send VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. Enter the values in multiples of 200 milliseconds.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
- 2. Click VLANs > Basic.

- 3. Select a VLAN.
- 4. Click IP.
- 5. Click the **VRRP** tab.
- 6. Click Insert.
- 7. In the IP, VLAN, Insert VRRP dialog box, click the **FasterAdvIntervalEnable** enable option.
- 8. In the **FasterAdvinterval**, box, enter a value. You must set the value using multiples of 200 milliseconds.
- 9. Click Insert.

Chapter 10: VRF Lite

Table 44: Virtual Routing and Forwarding product support

Feature	Product	Release introduced
Virtualization with IPv4 Virtual	VSP 4450 Series	VSP 4000 4.0
Routing and Forwarding (VRF)	VSP 4900 Series	VOSS 8.1
• ARP	VSP 7200 Series	VOSS 4.2.1
DHCP RelayInter-VRF Routing (static,	VSP 7400 Series	VOSS 8.0
dynamic, and policy)	VSP 8200 Series	VSP 8200 4.0
Local routing	VSP 8400 Series	VOSS 4.2
• OSPFv2	VSP 8600 Series	VSP 8600 4.5
RIPv1 and v2Route policiesStatic routingVRRP	XA1400 Series	VOSS 8.0.50
For configuration details, see Configuring IPv4 Routing for VOSS.		
Increased VRF and Layer 3 VSN	VSP 4450 Series	VOSS 6.0
scaling	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 6.0
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 6.0
	VSP 8400 Series	VOSS 6.0
	VSP 8600 Series	VSP 8600 8.0
	XA1400 Series	Not Supported
IBGP over user-created VRFs	VSP 4450 Series	VOSS 8.1
For configuration details, see	VSP 4900 Series	VOSS 8.1
Configuring BGP Services for VOSS.	VSP 7200 Series	VOSS 8.1
<u>voss</u> .	VSP 7400 Series	VOSS 8.1
	VSP 8200 Series	VOSS 8.1
		Table continues

Feature	Product	Release introduced
	VSP 8400 Series	VOSS 8.1
	VSP 8600 Series	Not Supported
	XA1400 Series	Not Supported

VRF Lite provides secure customer data isolation.

VRF Lite Fundamentals

The switch supports what is termed as VRF Lite. Lite conveys the fact that the switch does not use Multiprotocol Label Switching (MPLS) for VRF; VRF Lite is a device virtualization feature, not a network-wide virtualization feature.

Use VRF Lite to offer networking capabilities and traffic isolation to customers that operate over the same node (router). Each virtual router emulates the behavior of a dedicated hardware router; the network treats each virtual router as a separate physical router. In effect, you can perform the functions of many routers using a single platform that runs VRF Lite. The result is a substantial reduction in the cost associated with providing routing and traffic isolation for multiple clients.

With multicast virtualization for IPv4, the switch can function as multiple virtual multicast routers.

The following figure shows one platform acting as multiple virtual routers, each serving a different customer network.

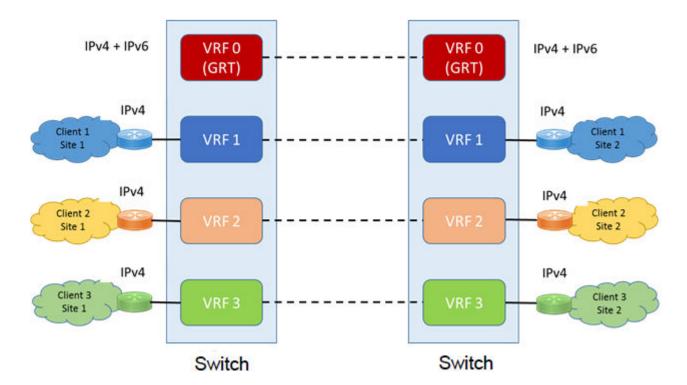


Figure 23: Multiple virtual routers in one system

A switch can support many virtual routers. Each virtual router instance is called a VRF instance. A VRF represents a single instance of a virtual router. Each instance maintains its own routing table. The term Multiple Virtual Router (MVR) is sometimes used to represent a router that contains many VRF instances.

The IPv6 Virtualization functionality adds IPv6 support on VRFs and Layer 3 VSNs. Each VRF instance has its own IPv6 interfaces, IPv6 address space, IPv6 routing table, and IPv6 global parameters. For more information on Layer 3 VSN, see Configuring Fabric Layer 3 Services for VOSS.

The Global Router, VRF 0, is the first instance of the router. When the system starts, it creates VRF 0 by default. VRF 0 provides all non-virtual and traditional routing services. You cannot delete this instance. You can create and configure other VRF instances, if required.

VRF 0 is the only VRF that you can log into through CLI. CLI requires you to specify the VRF when you enter commands.

You can associate one VRF instance with many IPv4 or IPv6 interfaces. These interfaces are unique for each VRF instance. An interface is an entity with an IPv4 or IPv6 address that has the following characteristics:

- A unique association with a VLAN.
- A unique association with a brouter, if not associated with a VLAN
- A unique association with a circuit

A VLAN can only be associated with a single VRF instance.

Note:

- You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IPv4
 or IPv6 address. You must first associate the port and VRF instance and then you can
 configure the IPv4 or IPv6 address.
- Use command boot config flag vrf-scaling to increase total VRFs. You must have a premier license to increase the total VRF count on the switch. For more information on route scaling, see <u>Release Notes for VOSS</u>.

VRF Lite capability and functionality

On a VRF instance, VRF Lite supports the following protocols:

- Border Gateway Protocol (BGP)
- IP
- Internet Control Message Protocol (ICMP)
- Address Resolution Protocol (ARP)
- · Static routes
- · Default routes
- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- · Route policies
- Virtual Router Redundancy Protocol (VRRP)
- Dynamic Host Configuration Protocol (DHCP), and BootStrap Protocol relay agent
- · User Datagram Protocol (UDP) forwarding
- Protocol Independent Multicast Sparse Mode (PIM-SM)
- Protocol Independent Multicast Source Specific Multicast (PIM-SSM)
- Internet Group Management Protocol (IGMP)
- Intermediate-System-to-Intermediate-System (IS-IS)

The switch uses VRF Lite to perform the following actions:

- Partition traffic and data and represent an independent router in the network
- Provide virtual routers that are transparent to end-users
- Support addresses that are not restricted to the assigned address space provided by host Internet Service Providers (ISP)
- Support overlapping IP address spaces in separate VRF instances

Note:

If you enable multicast route redistribution between two VRFs, the switch does not support IP addresses that overlap within the two VRFs. The device does not generate an error if addresses overlap. You must avoid this situation.

VRF Lite interoperates with RFC 4364, Layer 3 VPNs. Split MultiLink Trunking (SMLT) and Routed SMLT (RSMLT) are also supported for VRF instances.

Although customer data separation into Layer 3 virtual routing domains is usually a requirement, sometimes customers must access a common network infrastructure. For example, they want to access the Internet, data storage, Voice over IP (VoIP)-public switched telephone network (PSTN), or call signaling services. To interconnect VRF instances, you can use an external firewall that supports virtualization, or use inter-VRF forwarding for specific services. With the inter-VRF solution, you can use routing policies and static routes to inject IP subnets from one VRF instance to another, and you can use filters to restrict access to certain protocols. The following figure depicts inter-VRF forwarding by the switch.

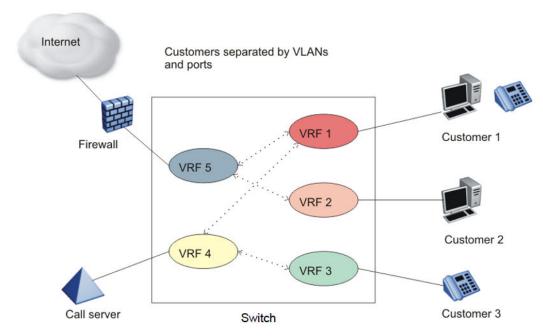


Figure 24: Inter-VRF forwarding

For more information about the latest VRF Lite scalability, see Release Notes for VOSS.

For configuration information about multicast virtualization, see <u>Configuring IP Multicast Routing Protocols for VOSS</u>.

VRF Lite and inter-VRF route redistribution

The switch supports three route redistribution functions:

- Intra-VRF inter-protocol route redistribution (redistribution within the same VRF instance), for example, redistribute RIP to OSPF.
- Inter-VRF inter-protocol redistribution (redistribution between two VRF instances), for example, redistribute RIP in VRF 2 to OSPF in VRF 4.
- Inter-VRF static routes (for example, a static route in a given VRF instance) configured as a typical static route but with the added parameter of a next-hop-vrf (the next-hop IP address is found in the next-hop-vrf instance).

With inter-VRF route redistribution, a user in one VRF instance can access route data in other VRF instances. You can redistribute routes within a VRF instance or between VRF instances; for example, one VRF instance can redistribute routes to all other VRF instances. You can redistribute Local, static, OSPF, RIP, and BGP routes and both dynamic (OSPF, BGP, and RIP) and static route redistribution is supported.

More than one routing protocol can be present in each VRF instance. Route redistribution can occur either between different protocol types, or between the same protocol types on different VRF instances.

An interface uses redistribution to announce routes that are learned by other protocols (OSPF or BGP, for example). Control route redistribution by using route policies. When you associate routing policies with route redistribution, the policy is checked before the target protocol is updated. Across VRF instances, the policy is checked at the source VRF instance, so only qualified routes are added to the routing table.

You can use static route commands to inject one specific route (including a default route) from one VRF instance to another. The route is added to the target VRF instance, while the next hop is resolved by the next-hop VRF instance.

Static routes are used to direct packets from a given source using a next-hop IP address. The next-hop-vrf option in a static route permits this path to proceed from one VRF to another. Overlapping IP addresses are supported within VRFs, thus it is possible for two VRFs to have identical IP addresses.

The following list describes interVRF route redistribution:

- Redistributed routes are added to the target VRF instance, and their next hop remains in the source VRF instance.
- If either the source or destination VRF instance is deleted, the redistribution configuration is automatically deleted.
- Redistributed routes are not further redistributed to another VRF instance.
- Route redistribution is unidirectional. You must configure route redistribution for the reverse direction if you require it. You can configure different route policies for each direction.
- After you configure interVRF route redistribution between two VRF instances, you must avoid using overlapping IP addresses between these two VRF instances.

Avoid overlapping addresses; the device does not generate an error if addresses overlap.

- Intra-VRF routes take precedence over inter-VRF routes.
- You can physically connect two VRF instances to distribute route across VRF instances (in this
 case, you do not need to configure route redistribution).

Route Redistribution Operation

To perform redistribution, the device maintains a route change list. The change list contains all the best routes that are either added to or deleted from the forwarding table. When a best route is added to or deleted from the forwarding table, the change list is updated to reflect the change and notify registered protocols. The registered protocols pick up the change from the change list when it becomes available.

An example scenario of interVRF redistribution follows. To redistribute OSPF routes in VRF 1 to RIP in VRF 0:

- Create, enable, and apply a RIP redistribution instance. The source protocol is OSPF and the VRF source is VRF 1.
- When an OSPF route is added in VRF 1, the Routing Table Manager (RTM) in VRF 1 puts the new route into the change list.
- The device notifies RIP in VRF 0, because RIP is registered with the RTM of VRF 1 for OSPF route changes.
- To send OSPF routes from VRF 1 through the RIP interface in VRF 0, the interface uses a route policy with match VRF criterion of VRF 1.

The switch also supports inter-domain multicast routing. For more information, see <u>Configuring IP Multicast Routing Protocols for VOSS</u>.

Port parameters and VRF Lite management

You can configure each VRF instance as a separate router, this means that you can configure different routing protocols and associated parameters for each instance. You can associate non0 VRF instances with ports.

The port parameters that you can edit for a VRF instance depend on whether the port belongs to only one, or more than one, VRF instance. For example, if a port belongs to only one VRF, you can edit the port parameters of the VRF. If a port belongs to more than one VRF instance, you cannot edit the port parameters of that port unless you are accessing the port through the Global Router with read-write-all access. If you do not have read-write-all access, you can only edit the GlobalRouter port parameters. If a port belongs to a single non0 VRF, the port parameters can be changed by this VRF. If a port belongs to multiple VRF instances, only a user with read-write-all access who is accessing the port through the Global Router can change this port configuration.

Management VRF

The following sections detail Management VRF features.

Management Port

The management port is a 10/100/1000 Mb/s Ethernet port that you can use for an out-of-band management connection to the switch.

Note:

Not all hardware platforms provide a dedicated, physical management interface. For more information, see your hardware documentation.

Management Router VRF

Note:

MgmtRouter is only supported on VSP 8600 Series.

The switch has a separate VRF called Management Router (MgmtRouter) reserved for the management port and the Virtual Management IP address. The configured IP subnet has to be globally unique because the management protocols, for example, SNMP, Telnet, and FTP, can go through in-band or out-of-band ports. The VRF ID for the Management Router is 512.

The switch never switches or routes transit packets between the Management Router VRF port and the Global Router VRF, or between the Management Router VRF and other VRF ports.

The switch honors the VRF of the ingress packet; however, in no circumstance does the switch allow routing between the Management VRF and Global Router VRF. The switch does not support the configuration if you have an out-of-band management network with access to the same networks present in the GRT routing table.

Note:

IPv6 is not supported on MgmtRouter.

Non Virtualized Client Management Applications

Note:

This section only applies to VSP 8600 Series.

Ensure that you do not define a default route in the Management Router VRF. A route used for non-virtualized client management applications, such as Telnet, Secure Shell (SSH), and FTP, originating from the switch, will always match a default route defined in the Management Router VRF.

If you want out-of-band management, it is recommended that you define a specific static route in the Management Router VRF to the IP subnet where your management application resides.

When you specify a static route in the Management Router VRF, it enables the client management applications originating from the switch to perform out-of-band management without affecting inband management. This enables in-band management applications to operate in the Global Router VRF.

Non-virtualized client management applications originating from the switch, such as Telnet, SSH, and FTP, follow the behavior listed below:

- 1. Look at the Management Router VRF route table.
- 2. If no route is found, the applications will proceed to look in the Global Router VRF table.

Non-virtualized client management applications include:

- DNS
- · FTP client with the copy command
- NTP
- rlogin



Note:

Rlogin is only supported on VSP 8600 Series.

- RADIUS authentication and accounting
- SSH
- SNMP clients in the form of traps
- SYSLOG
- TACACS+
- Telnet
- TFTP client

For management applications that originate outside the switch, the initial incoming packets establish a VRF context that limits the return path to the same VRF context.

Virtualized Management Applications

Virtualized management applications, such as ping and traceroute, operate using the specified VRF context. To operate ping or traceroute you must specify the desired VRF context. If not specified, ping defaults to the Global Router VRF. For example, if you want to ping a device through the out-ofband management port you must select the Management Router VRF. For example, if you want to ping a device through the out-of-band management port you must select the Management Router VRF



Note:

Ipv6 is not supported on MgmtRouter.

```
Switch:1(config) #ping 192.0.2.1 vrf MgmtRouter
192.0.2.1 is alive
```

Ping test for IPv6:

```
Switch:1(config) #ping 2001:db8::1 vrf vrfRED
2001:db8::1 is alive
```

Traceroute test for IPv4:

Switch:1#traceroute 192.0.2.1 vrf MgmtRouter

Traceroute test for IPv6:

Switch:1#traceroute 2001:db8::1 vrf vrfRED

VRF Lite configuration rules

You must select the VRF for global IPv4 or IPv6 options before entering commands.

Not all Global Router parameters are configurable on other VRF instances.

For instructions about how to configure a VRF instance, see the following paragraphs.

Layer 1 and Layer 2 information (including VLAN information) is global and is not maintained for each VRF instance. However, you can associate a set of VLANs with a VRF instance.

One VLAN cannot belong to more than one VRF instance at one time. When you create a VLAN, more than one physical port can belong to it. You can associate a VRF instance with more than one IPv4 or IPv6 interface (a physical Ethernet port or a VLAN).

Perform physical port assignment at the VLAN and brouter port level. A VRF instance inherits all the ports assigned to its VLANs and brouter ports. You cannot directly assign a physical port to a VRF instance, but it is implicitly assigned when you associate the VRF with VLANs or brouter ports.

For IPv4, after you configure interVRF route redistribution between two VRF instances, avoid overlapping IP addresses between these two VRF instances.

When you configure VRF Lite, remember the following rules:

- You can connect two VRFs from the same system with an external cable.
- An IPv4 or IPv6 routable VLAN can become a member of a VRF.
- · An IPv4 or IPv6 interface can belong to only one VRF.
- A VRF can exist even if no interfaces are assigned to it.
- Routing policies apply to VRFs on an individual basis.
- Multiple VRFs on the same node can function in different autonomous systems.

Following rules apply to IPv4 interfaces specifically:

- If you configure an IPv4 interface without specifying the VRF instance, it is mapped to VRF 0 by default.
- VRF Lite supports SMLT and RSMLT
- VRF Lite supports RIP in and out policies
- VRF Lite supports OSPF in and out (accept and redistribute) policies
- Before you delete a VRF instance, disable OSPF. Deleting a VRF instance deletes the OSPF instance if OSPF is disabled
- When you create a VRF instance, an OSPF instance is not automatically created. To activate OSPF on a VRF instance, first create an OSPF instance, and then enable OSPF

- You can configure a VRF so it can have IP interfaces with OSPF, RIP, static routes, and policies simultaneously
- Every IPv4 interface is a member of VRF 0 unless explicitly defined to belong to another VRF.

Virtualized Protocols

VRF Lite supports virtualization of the following IPv4 and IPv6 protocols and features. Use this table to find applicable VRF command and procedure information.

Table 45: Virtualized IPv4 Protocols and Documentation

Virtualized IPv4 protocol or feature	Where to find information
ARP	Configuring IPv4 Routing for VOSS
BGP	Configuring BGP Services for VOSS
Circuitless IP	Configuring IPv4 Routing for VOSS
DHCP	Configuring IPv4 Routing for VOSS
IGMP	Configuring IP Multicast Routing Protocols for VOSS
OSPF	Configuring OSPF and RIP for VOSS
RIP	Configuring OSPF and RIP for VOSS
Route policies	Configuring IPv4 Routing for VOSS
Route preferences	Configuring IPv4 Routing for VOSS
Router Discovery	Configuring IPv4 Routing for VOSS
Static routes	Configuring IPv4 Routing for VOSS
User Datagram Protocol (UDP)	Configuring IPv4 Routing for VOSS
VLAN	Configuring VLANs, Spanning Tree, and NLB for VOSS.
VRRP	Configuring IPv4 Routing for VOSS

Table 46: Virtualized IPv6 Protocols and Documentation

Virtualized IPv6 protocol or feature	Where to find information
BGP	Configuring BGP Services for VOSS
IPv6 Interfaces and IPv6 Static Routes	Configuring IPv6 Routing for VOSS
ECMP and Alternative Route	Configuring IPv6 Routing for VOSS
OSPF	Configuring IPv6 Routing for VOSS
Route redistribution for static and direct routes	Configuring IPv6 Routing for VOSS

Virtualized IPv6 protocol or feature	Where to find information
VRRPv3	Configuring IPv6 Routing for VOSS
DHCP Relay	Configuring IPv6 Routing for VOSS
IPv6 Reverse Path Forwarding	Configuring IPv6 Routing for VOSS
ICMP Ping & Traceroute	Configuring IPv6 Routing for VOSS
ISIS Accept Policies	Configuring Fabric Layer 3 Services for VOSS

VRF Lite configuration using the CLI

Use Virtual Router and Forwarding (VRF) Lite to provide many virtual routers using one switch.

This section shows you how to configure a VRF instance and how to associate ports and VLANs with VRF instances.

The following task flow shows you the sequence of procedures you perform to configure VRF Lite.

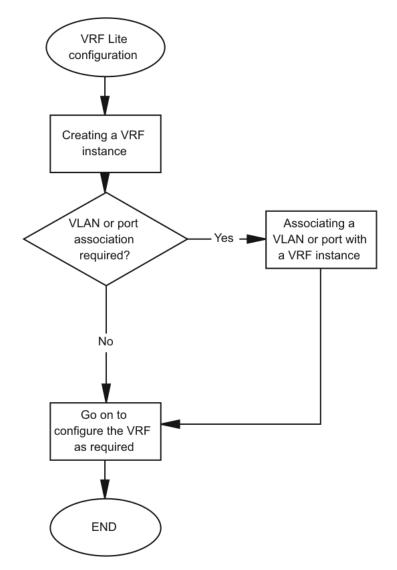


Figure 25: VRF Lite configuration procedures

Creating a VRF Instance

About this task

Create a VRF instance to provide a virtual routing interface for a user.

For more information on route scaling, see Release Notes for VOSS.



Note:

Different hardware platforms support different parameter ranges. Use the CLI Help to see the available range.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a VRF instance and specify a VRF name:

```
ip vrf WORD<1-16>
```

- 3. Configure the maximum number of routes:
 - For IPv4:

```
ip vrf WORD<1-16> max-routes <0-15744 | 0-15488 | 0-256000>
```

For IPv6:

```
ip vrf WORD<1-16> ipv6-max-routes <0-7744 | 0-7872>
```



The maximum routes of IPv4 and IPv6 for Global Router (GRT) are not configurable and fixed at the system limits.

- 4. Enable max-routes traps:
 - For IPv4:

```
ip vrf WORD<1-16> max-routes-trap enable
```

• For IPv6:

```
ip vrf WORD<1-16> ipv6-max-routes-trap enable
```

Note:

Maximum Route traps are not generated for GRT. For non-default VRFs, the permitted maximum routes can be lower than system limits and traps generate when the limit is exceeded.

5. Enter VRF Router Configuration mode:

```
router vrf WORD<1-16>
```

6. Configure the IP routing protocol triggers for the VRF:

Use one of the following commands on your switch:

• ip bgp

ip bgp creates both ipv4 and ipv6 instances.

• ip ospf

Use ipv6 ospf to create an OSPFv3 instance.

• ip rip

RIPng is not virtualized, hence the IPv6 configuration is not applicable here.

Note:

You cannot configure BGP, OSPF, or RIP on a VRF instance unless you first configure the routing protocol trigger.

7. Ensure that the instance is configured correctly:

```
show ip vrf WORD<1-16>
```

Example

Create a VRF instance:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip vrf vrfRED
```

Configure the maximum number of IPv4 routes and enable max-routes traps.

```
Switch:1(config) #ip vrf vrfRED max-routes 12000
Switch:1(config) #ip vrf vrfRED max-routes-trap enable
```

Enter Router Configuration mode and configure the routing protocol triggers for the VRF:

```
Switch:1(config) #router vrf vrfRED
Switch:1(router-vrf) #ip bgp
Switch:1(router-vrf) #ip ospf
Switch:1(router-vrf) #ip rip
```

To Configure OSPFv3 instance for the VRF:

```
Switch:1(config) #router vrf vrfRED
Switch:1(router-vrf) #ipv6 ospf
```

Exit to Global configuration mode:

Switch:1(router-vrf)#exit

Configure the maximum number of IPv6 routes and enable IPv6 max-routes traps.

```
Switch:1(config) #ip vrf vrfRED ipv6-max-routes 7700
Switch:1(config) #ip vrf vrfRED ipv6-max-routes-trap enable
```

Ensure that the instance is configured correctly:

			7	RF INFOR	RMATION				
VRF COUNT	OSPF COUNT	RIP COU	NT	BGP COUNT	PIM COUNT	ARI COU		PIM6 COUNT	OSPFv3 COUNT
 1	1	1		1	1	7		1	1
VRF NAME	VRF ID	OSPF	RIP	BGP	PIM	VLAN COUNT	ARP COUNT	PIM6	OSPFv3
rfRED	3	TRUE	TRUE	TRUE	TRUE	0	7	TRUE	FALSE

Variable definitions

Use the data in the following table to use the ip vrf command.

Note:

Different hardware platforms support different parameter ranges. Use the CLI Help to see the available range.

Table 47: Variable definitions

Variable	Value
Depending on your hardware platform:	Configures the maximum number of IPv4 routes allowed for the VRF, which is 15488, 15744, or 252000, depending on your hardware platform.
max-routes <0-15488 0-15744 0-252000>	The default value is 10000, except for the GlobalRouter, which is 15488, 15744, or 252000, depending on your hardware platform.
ipv6-max-routes <0-7744 0-7872>	Configures the maximum number of IPv6 routes allowed for the VRF, which is 7744 or 7872, depending on your hardware platform.
	The default value is 5000.
max-routes-trap enable	Enables SNMP traps after the maximum number of IPv4 routes are reached.
ipv6-max-routes-trap enable	Enables SNMP trap generation based on the configured number of maximum IPv6 routes. The default is enabled.
name WORD<0-16>	Renames the VRF instance.
vrf-trap	Enables the device to send VRF-related traps.

Use the data in the following table to use the show ip vrf command.

Table 48: Variable definitions

Variable	Value
max-routes [vrfids WORD<0-512>]	Displays the maximum number of routes for the specified VRFs.
[WORD <1–16>]	vrfids WORD<0-512> specifies a list of VRFs by VRF IDs.
	WORD<1-16> specifies a VRF by name.
vrfids WORD<0-512>	Specifies a list of VRFs by VRF IDs.
WORD<1-16>	Specifies a VRF by name.

Associating a VLAN or port with a VRF instance

You can assign a VRF instance to a port or VLAN. You cannot associate a VLAN or port and a VRF instance if the VLAN or port has an IP address. You can configure the IP address after you associate the port and VRF instance.

Before you begin

• Ensure the VRF is already configured.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]} Of interface vlan <1-4059>
```



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

2. Associate the port or VLAN with a VRF instance:

```
vrf WORD<1-16>
```

Example

Switch:1> enable
Switch:1# configure terminal

Create a VRF named Two:

Switch:1(config-if) # ip vrf Two

Create a VLAN of type byport:

Switch:1(config-if) # vlan create 33 name vlan-30 type port-mstprstp 0

Enter VLAN Interface Configuration mode:

Switch:1(config-if) # interface vlan 33

Assign the VLAN to VRF Two:

Switch:1(config-if) # vrf Two

Give the VLAN an IP address:

Switch:1(config-if) # ip address 192.0.2.1 255.255.255.0

Enter VRF configuration mode:

Switch:1(config-if) # router vrf Two

Variable definitions

Use the data in the following table to use the vrf command.

Table 49: Variable definitions

Variable	Value
vrf WORD<0-16>	Specifies the VRF instance by name.

Creating an IP VPN instance on a VRF

Create an IP VPN instance to advertise IP routes from one VRF to another across a Shortest Path Bridging MAC (SPBM) network.

For more information about Layer 3 Virtual Services Networks (VSN) and SPBM, see <u>Configuring Fabric Basics and Layer 2 Services for VOSS</u>.

Before you begin

· The VRF must exist.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create an IP VPN instance on the VRF:

```
ipvpn
```

3. Assign a service instance identifier (I-SID) to the IP VPN:

```
i-sid <0-16777215>
```

4. Enable IP VPN on the VRF:

```
ipvpn enable
```

By default, a new IP VPN instance is disabled.

5. Display all IP VPNs:

```
show ip ipvpn [vrf WORD<1-16>] [vrfids WORD<0-512>]
```

Example

From Global Configuration mode, log on to Router VRF Configuration mode:

```
Switch:1(config) # router vrf red
```

Create the IP VPN instance:

```
Switch:1(router-vrf) # ipvpn
```

Enable IP VPN:

```
Switch:1(router-vrf) # i-sid 100
```

Enable IP VPN:

1 out of 2 Total Num of VRF Entries displayed.

Variable definitions

Use the data in the following table to use the **show ip ipvpn** command.

Variable	Value
vrf WORD<1–16>	Specifies the VRF name.
vrfids WORD<0-512>	Specifies the VRF ID.

Use the data in the following table to use the i-sid command.

Variable	Value
i-sid <0-16777215>	Assigns an I-SID to the VRF to configure. Use the no or default option to remove the I-SID to VRF allocation for this VRF.

Configure the Maximum Number of VRFs

Perform this procedure to change the maximum number of VRFs and Layer 3 VSNs that the switch supports. By default, the switch supports 24 VRFs and Layer 3 VSNs. Increasing the number of VRFs or Layer 3 VSNs can be useful in a WAN scenario or other large network.

The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see Release Notes for VOSS.

About this task



If you enable this boot config flag, and the switch operates in SPBM mode (default configuration), the switch reduces the number of configurable VLANs. In such a configuration, the switch reserves VLANs 3500 to 3998 for internal use. You cannot use these VLANs as either platform VLANs or B-VLANs. You can still use the reserved VLAN range for customer VLANs (C-VLAN) on Flex UNI and B-VLANs on FE-VID.

Enabling the boot config flag to use more than 24 VRFs requires a Premier or Premier + MACsec license.

Before you begin

- If the switch operates in SPBM mode, before you enable the boot config flag, perform the following actions:
 - Check in-VLAN filters. If a filter references a VLAN in the 3500 to 3998 range, you must delete the filter or the filter configuration fails when you restart the switch.
 - Delete VLANs in the 3500 to 3998 range.
- Before you disable the boot config flag, delete additional VRFs if more than 24 exist.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Increase the maximum number of VRFs and Layer 3 VSNs:

```
boot config flag vrf-scaling
```

OR

3. Return to the default of 24 VRFs and Layer 3 VSNs:

```
no boot config flag vrf-scaling
or
default boot config flag vrf-scaling
```

4. Verify the configuration:

```
show boot config flags
```

5. Save the configuration:

```
save config
```

6. Restart the switch for the change to take effect:

reset

Example

Enable the boot config flag to increase the maximum number of VRFs and Layer 3 VSNs. In the following example, the switch operates in SPBM mode and reserves the VLAN ID range of 3500 to 3999. If the switch does not operate in SPBM mode, the VLAN warning message does not appear when you enable VRF scaling.

```
Switch:1>en
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flag vrf-scaling
Warning: Vlan 3500 to 3999 will be reserved for internal use.

Warning: Please save the configuration and reboot the switch
for this configuration to take effect.
```

Note:

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
```

```
flags ftpd true
flags ha-cpu true
flags hsecure false
flags ipv6-egress-filter true
flags ipv6-mode false
flags linerate-directed-broadcast false
flags logging true
flags nni-mstp false
flags reboot true
flags rlogind false
flags savetostandby true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

The following example shows the message that appears if you try to enable the boot config flag and configured VLANs use IDs between 3500 and 3999.

```
Switch:1>en
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flag vrf-scaling

Error: Delete all configured platform vlans between 3500 and 3999 to enable vrf-scaling.
```

VRF Lite configuration using Enterprise Device Manager

Use VRF Lite to provide many virtual routers using a single switch.

Configuring a VRF instance

About this task

Configure a VRF instance to provide a virtual routing interface for a user.



Note:

The maximum routes of IPv4 and IPv6 for Global Router (GRT) are non-configurable and fixed at the system limits.

Maximum route traps are not generated on GRT. For non-default VRFs, the permitted maximum routes can be lower than system limits and traps generate when the limit is exceeded.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VRF**.
- 2. Click VRF.

- 3. Click the **VRF** tab.
- 4. Click Insert.
- 5. Specify the VRF ID.
- 6. Name the VRF instance.
- 7. Configure VRF Lite-related traps.
- 8. Configure the other parameters as required.
- 9. Click Insert.

VRF field descriptions

Use the data in the following table to help you use the VRF tab.

Name	Description
Id	Specifies the ID number of the VRF instance. VRF ID 0 is reserved for the GlobalRouter.
Name	Names the VRF instance.
ContextName	Identifies the VRF. The SNMPv2 Community String or SNMPv3 contextName denotes the VRF context and is used to logically separate the MIB port management.
TrapEnable	Enables the VRF to send VRF Lite-related traps (VrfUp and VrfDown). The default is enabled.
MaxRoutes	Configures the maximum number of routes allowed for the VRF, which is 15488 or 15744, depending on your hardware platform.
	The default value is 10000, except for the GlobalRouter, which is 15488 or 15744, depending on your hardware platform.
RpTrigger	Specifies the Routing Protocol (RP) triggers for the VRF. The triggers are used to initiate or shutdown routing protocols on a VRF. You can act upon multiple RPs simultaneously. You can also use this option to bring individual RPs up in steps.
MaxRoutesTrapEnable	Enables the generation of the VRF Max Routes Exceeded traps. The default is enabled.
Ipv6MaxRoutes	Configures the maximum number of IPv6 routes allowed for the VRF, which is 7744 or 7872, depending on your hardware platform.
	The default value is 5000.
Ipv6MaxRoutesTrapEnable	Enables SNMP trap generation after the maximum number of IPv6 routes are reached.
	The default is enabled.

Associating a port to a VRF instance

About this task

Associate a port to a Virtual Router Forwarding (VRF) instance so that the port becomes a member of the VRF instance.

You can assign a VRF instance to a port after you configure the VRF. The system assigns ports to the GlobalRouter, VRF 0, by default.

Procedure

- 1. In the **Device Physical** View tab, select a port.
- 2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
- 3. Click General.
- 4. Click the VRF tab.
- 5. To the right of the **BrouterVrfld** box, click the ellipsis (...) button.
- 6. In the BrouterVrfld dialog box, select the required VRF.
- 7. Click OK.
- 8. Click Apply.

VRF field descriptions

Use the data in the following table to use the VRF tab.

Name	Description
Vrflds	Shows the VRF ID.
VrfNames	Shows the VRF name.
VrfCount	Shows the number of VRFs to which the port is associated.
BrouterVrfld	Specifies the VRF ID for a brouter port.
BrouterVrfName	Shows the VRF name for a brouter port.

Associating an Extreme Integrated Application Hosting Port to a VRF Instance

About this task

Perform this procedure to associate an Extreme Integrated Application Hosting (IAH) port to a Virtual Router Forwarding (VRF) instance.

Note:

You can associate a VRF instance to an IAH port after you configure the VRF. By default, the IAH ports are associated to the GlobalRouter.

Procedure

- 1. In the navigation pane, expand the following folders: **Configuration > Edit > Insight Port**.
- 2. Select the IAH port you want to configure.
- 3. Select the VRF tab.
- 4. **(Optional)** In the **VrfNames** field, select the **ShowAll** button to view the VRF instances the IAH Port is associated with.
- 5. In the **BrouterVrfld** field, select the ellipsis (...) button, and select the required VRF instance(s).
- 6. Select Ok.
- 7. Select Apply.

VRF Field Descriptions

Use data in the following table to use the VRF tab.

Name	Description
Vrflds	Shows the VRF ID.
VrfNames	Shows the VRF name.
VrfCount	Shows the number of VRFs to which the Extreme Integrated Application Hosting (IAH) port is associated with.
BrouterVrfld	Specifies the VRF ID for a brouter port.
BrouterVrfName	Shows the VRF name for a brouter port.

Configuring interVRF route redistribution policies

Before you begin

- · Ensure VRF instances exist.
- Configure route policies, if required.
- Change the VRF instance as required.

About this task

Configure inter-VRF route redistribution so that a VRF interface can announce routes that other protocols learn, for example, OSPF or BGP. Use a route policy to control the redistribution of routes.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > IP**.

- 2. Click Policy.
- 3. Click the Route Redistribution tab.
- 4. Click Insert.
- 5. Choose the source and destination VRF IDs.
- 6. Choose the protocol and route source.
- 7. Select Enable.
- 8. Choose the route policy to apply to the redistributed routes.
- 9. Configure other parameters as required.
- 10. Click Insert.
- 11. Click the **Applying Policy** tab.
- 12. Select RedistributeApply, and then click Apply.

Route Redistribution field descriptions

Use the data in the following table to use the Route Redistribution tab.

Name	Description
DstVrfld	Specifies the destination VRF ID to use in the redistribution.
Protocol	Specifies the protocols for which you want to receive external routing information.
SrcVrfld	Specifies the source VRF ID to use in the redistribution.
RouteSource	Indicates if the protocol receives notification about the routes this source learns. The route source is equivalent to the owner in the routing table.
Enable	Enables or disables route redistribution. The default is disabled.
RoutePolicy	Specifies the route policies to apply to the redistributed routes from the source VRF. Use the route policy to determine whether the system advertises a specific route to the specified protocol.
Metric	Specifies the metric announced in advertisements. The default is 0.
MetricType	Specifies the metric type (applies to OSPF and BGP only). Specifies a type 1 or a type 2 metric. For metric type 1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type 2, the cost of the external routes is equal to the external cost alone. The default is type2.
Subnets	Indicates that all the subnets must be advertised individually. The values are allow(1), and suppress(2). The default value is allow. This variable applies to OSPF only.

Viewing brouter port and VRF associations

About this task

You can view each port and associated VRFs. You can also change the VRFs associated with the port if the port has no IP address.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VRF**.
- 2. Click VRF.
- 3. Click the VRF-Ports tab.
- 4. To display the VRF names associated with a port, click a cell in one of the table rows and, on the toolbar, click the **ShowVRFNames** button.
- 5. To change the VRF, double-click the **BrouterVrfld** field for the port.
 - Tip:

You can associate a port with more than one VRF.

- 6. Choose the required VRFs, and then click **Ok**.
- 7. Click Apply.

VRF-Ports field descriptions

Use the data in the following table to use the **VRF-Ports** tab.

Name	Description
Index	Specifies the slot and port.
Туре	Specifies the port type.
Vrflds	Identifies the set of VRF IDs to which this port belongs.
VrfCount	Shows the number of VRF instances associated with this port.
BrouterVrfld	Shows the VRF ID for this brouter port.
BrouterVrfName	Shows the VRF name for this brouter port.
Show VrfNames	You can use this toolbar button to identify the set of VRF names to which a port belongs.

Use the data in the following table to use the **Show VrfNames** button.

Name	Description
Index	Specifies the slot and port.
VrfNames	Shows the VRF name for this brouter port.

Viewing global VRF status information

About this task

View global VRF status information to determine the number of VRFs that are configured and active.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VRF**.
- 2. Click VRF.
- 3. Click the Global Status tab.

Global Status field descriptions

Use the data in the following table to use the Global Status tab.

Name	Description
ConfigNextAvailableVrfld	Specifies the number of the next available Virtual Router ID (index).
ConfiguredVRFs	Specifies the number of VRFs configured on this network element.
ActiveVRFs	Specifies the number of VRFs that are active on the network element. These are VRFs for which the OperStatus is up.

Viewing VRF instance statistics and status information

About this task

View VRF instance status information to determine the operational status of each VRF, as well as other operational parameters.

Procedure

- 1. In the navigation tree, expand the following folders: **Configuration > VRF**.
- 2. Click VRF.
- 3. Click the VRF Stats tab.

VRF Stats field descriptions

Use the data in the following table to use the VRF Stats tab.

Name	Description
Id	Specifies the ID number of the VRF instance.
StatRouteEntries	Specifies the total number of routes for this VRF.
StatFIBEntries	Specifies the total number of Forwarding Information Base (FIB) entries for this VRF.

Name	Description
StatUpTime	Specifies the time in (in hundredths of a second) since this VRF entry has been operational.
OperStatus	Shows the operational status of the Virtual Router.
RouterAddressType	Specifies the router address type of this VRF.
Router Address	Specifies the router address of this VRF, derived from one of the interfaces. If a loopback interface is present, you can use the loopback interface address.

Viewing Statistics for a VRF

About this task

View VRF statistics to ensure the instance is performing as expected.

Procedure

- 1. In the navigation pane, expand the **Configuration** > **VRF** folders.
- 2. Click VRF.
- 3. Click the **VRF** tab.
- 4. Select a VRF.
- 5. Click the **Stats** button.

Stats Field Descriptions

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

Name	Description
StatRouteEntries	Specifies the number of routes for this VRF.
StatFIBEntries	Specifies the number of Forwarding Information Base (FIB) entries for this VRF.

Selecting and launching a VRF context view

About this task

Use this procedure to switch to another VRF context view when you use the embedded EDM. GlobalRouter is the default view at log in. You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view. You can open only five tabs for each EDM session.

Important:

If you log out from the GRT view, the system generates a warning: all tabs close and your session terminates. If you close a VRF view tab, you close only that view.

Note:

The Set VRF Context view function is not available to users in a service provider deployment where only a tenant VRF view is assigned. If you use a tenant VRF view, it is recommended to use the applicable EDM plugin with COM to access EDM. COM provides VRF mapping and Role-Based Access Control.

Procedure

- 1. In the navigation pane, expand the following folders: Configuration > VRF Context View.
- 2. Click Set VRF Context View.
- 3. Click the **VRF** tab.
- 4. Select a context to view.
- Click Launch VRF Context view.

A new browser tab opens containing the selected VRF view

VRF field descriptions

Use the descriptions in the following table to use the **VRF** tab.

Name	Description		
Id	Shows the unique VRF ID.		
Name	Shows the name of the virtual router.		
ContextName	Shows the SNMPv3 context name that denotes the VRF context and logically separates the MIB port management.		

Create an IP VPN Instance on a VRF

Create an IP VPN instance to advertise IP routes from one VRF to another across a Shortest Path Bridging MAC (SPBM) network.

For more information about Layer 3 Virtual Services Networks (VSNs) and SPBM, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- The VRF must exist.

Procedure

- 1. In the navigation tree, expand **Configuration > IP** folders.
- 2. Click IP-VPN.
- Click the VPN tab.
- 4. Click Insert.

- 5. Click [...] and select a VRF from the list.
- 6. Click OK.
- 7. Click Insert.

By default, the new IP VPN instance is disabled.

- 8. In the **Isid Number** column, double-click the **0** value, and then enter the service instance identifier (I-SID) to assign to the IP-VPN.
- 9. In the **Enable** column, double-click the **disable** value.
- 10. Click the arrow to view a list of choices, and then choose **enable**.
- 11. Click Apply.

VPN Field Descriptions

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

Name	Description		
Vrfld	Specifies the ID of the VRF to configure.		
Enable	Enables or disables the IP VPN instance on the VRF. The default is disabled.		
Isid Number	Specifies the I-SID to associate with the VPN. By default, no I-SID is assigned.		
Isid Name	Specifies the name of the I-SID associated with the		
Note:	VPN.		
This field does not apply to all hardware platforms.			

Configure the maximum number of VRFs

Perform this procedure to change the maximum number of VRFs and Layer 3 VSNs that the switch supports. By default, the switch supports 24 VRFs and Layer 3 VSNs. Increasing the number of VRFs or Layer 3 VSNs can be useful in a WAN scenario or other large network.

The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see Release Notes for VOSS.

About this task



If you enable this boot config flag, and the switch operates in SPBM mode (default configuration), the switch reduces the number of configurable VLANs. In such a configuration, the switch reserves VLANs 3500 to 3998 for internal use. You cannot use these VLANs as either platform VLANs or B-VLANs. You can still use the reserved VLAN range for customer VLANs (C-VLAN) on Flex UNI and B-VLANs on FE-VID.

Enabling the boot config flag to use more than 24 VRFs requires a Premier or Premier + MACsec license.

Before you begin

- If the switch operates in SPBM mode, before you enable this boot config flag, perform the following actions:
 - Check in-VLAN filters. If a filter references a VLAN in the 3500 to 3998 range, you must delete the filter or the filter configuration fails when you restart the switch.
 - Delete VLANs in the 3500 to 3998 range.
- Before you disable this boot config flag, delete additional VRFs if more than 24 exist.

Procedure

- 1. In the navigation pane, expand the **Configuration > Edit** folders.
- 2. Select Chassis.
- 3. Select the **Boot Config** tab.
- 4. Perform one of the following actions:
 - a. To enable VRF scaling, select the **EnablevrfScaling** check box.
 - b. To disable VRF scaling, clear the **EnablevrfScaling** check box.
 - Note:

This field does not apply to all hardware platforms.

- 5. Select Apply.
- 6. Restart the switch for the change to take effect.

Glossary

Address Resolution Maps an IP address to a physical machine address, for example, maps an IP address to an Ethernet media access control (MAC) address. Protocol (ARP) A prefix length that is formed by combining several specific prefixes. The aggregate resulting prefix is used to combine blocks of address space into a single routing announcement. **Autonomous System** A set of routers under a single technical administration, using a single IGP (AS) and common metrics to route packets within the Autonomous System, and using an EGP to route packets to other Autonomous Systems. **Autonomous System** A two-byte number that is used to identify a specific AS. Number (ASN) **Bootstrap Protocol** A User Datagram Protocol (UDP)/Internet Protocol (IP)-based protocol that (BootP) a booting host uses to configure itself dynamically and without user supervision. A dynamically elected Protocol Independent Multicast (PIM) router that bootstrap router (BSR) collects information about potential Rendezvous Point routers and distributes the information to all PIM routers in the domain. **Bridge Protocol Data** A data frame used to exchange information among the bridges in local or Unit (BPDU) wide area networks for network topology maintenance. candidate bootstrap Provides backup protection in case the primary rendezvous point (RP) or bootstrap router (BSR) fails. Protocol Independent Multicast (PIM) uses the router (C-BSR) BSR and C-BSR. Circuitless IP (CLIP) A CLIP is often called a loopback and is a virtual interface that does not map to any physical interface. classless The protocol defined in RFCs 1517 and 1518 for using subnetwork masks, other than the defaults for IP address classes. interdomain routing (CIDR) **Dynamic Random** A read-write random-access memory, in which the digital information is **Access Memory** represented by charges stored on the capacitors and must be repeatedly

replenished to retain the information.

(DRAM)

Enterpri	se	Dev	ice
Manage	r (E	EDM)

A web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.

equal cost multipath (ECMP)

Distributes routing traffic among multiple equal-cost routes.

Global routing engine (GRE)

The base router or routing instance 0 in the Virtual Routing and Forwarding (VRF).

Institute of Electrical and Electronics Engineers (IEEE)

An international professional society that issues standards and is a member of the American National Standards Institute, the International Standards Institute, and the International Standards Organization.

Interior Gateway Protocol (IGP)

Distributes routing information between routers that belong to a single Autonomous System (AS).

Internet Assigned Numbers Authority (IANA)

The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.

Internet Control Message Protocol (ICMP)

A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.

Internet Protocol version 4 (IPv4)

The protocol used to format packets for the Internet and many enterprise networks. IPv4 provides packet routing and reassembly.

Layer 1

Layer 1 is the Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interacts with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.

Layer 2

Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.

Layer 3

Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).

Layer 3 Virtual Services Network

The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN).

link-state advertisement (LSA)

Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.

management information base (MIB)

The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).

mask

A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.

maximum transmission unit (MTU) The largest number of bytes in a packet—the maximum transmission unit of the port.

media

A substance that transmits data between ports; usually fiber optic cables or category 5 unshielded twisted pair (UTP) copper wires.

Media Access Control (MAC)

Arbitrates access to and from a shared medium.

MultiLink Trunking (MLT)

A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.

multiplexing

Carriage of multiple channels over a single transmission medium; a process where a dedicated circuit is shared by multiple users. Typically, data streams intersperse on a bit or byte basis (time division), or separate by different carrier frequencies (frequency division).

Network Basic Input/ Output System (NetBIOS) An application programming interface (API) that augments the DOS BIOS by adding special functions for Local Area Networks (LAN).

next hop

The next hop to which a packet can be sent to advance the packet to the destination.

operation, administration, and maintenance (OA&M) All the tasks necessary for providing, maintaining, or modifying switching system services.

Packet Capture Tool (PCAP)

A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.

port

A physical interface that transmits and receives data.

prefix

A group of contiguous bits, from 0 to 32 bits in length, that defines a set of addresses.

Protocol Data Units (PDUs)

A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.

Protocol Independent Multicast, Sparse Mode (PIM-SM) PIM-SM is a multicast routing protocol for IP networks. PIM-SM provides multicast routing for multicast groups that can span wide-area and interdomain networks, where receivers are not densely populated. PIM-SM sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.

remote monitoring (RMON)

A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.

Reverse Address Resolution Protocol (RARP) A protocol that maintains a database of mappings between physical hardware addresses and IP addresses.

reverse path checking (RPC)

Prevents packet forwarding for incoming IP packets with incorrect or forged (spoofed) IP addresses.

route flapping

An instability that is associated with a prefix, where the associated prefix routes can exhibit frequent changes in availability over a period of time.

route table manager (RTM)

Determines the best route to a destination based on reachability, route preference, and cost.

Routed Split MultiLink Trunking (RSMLT) Provides full router redundancy and rapid failover in routed core SMLT networks and as RSMLT-edge in routed SMLT edge applications; eliminating routing protocol timer dependencies when network failures occur.

Routing Information Protocol (RIP)

A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.

routing policy

A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.

Service Instance Identifier (I-SID)

The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 24 bits. SPBM uses this I-SID to identify and transmit any

virtualized traffic in an encapsulated SPBM frame. SPBM uses I-SIDs to virtualize VLANs (Layer 2 Virtual Services Network [VSN]) or VRFs (Layer 3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone. With Layer 2 VSNs, you associate the I-SID with a customer VLAN, which is then virtualized across the backbone. With Layer 3 VSNs, you associate the I-SID with a customer VRF, which is also virtualized across the backbone.

Shortest Path Bridging (SPB)

Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.

Shortest Path Bridging MAC (SPBM)

Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.

Simple Network Management Protocol (SNMP)

SNMP administratively monitors network performance through agents and management stations.

SMLT aggregation switch

One of two IST peer switches that form a split link aggregation group. It connects to multiple wiring closet switches, edge switches, or customer premise equipment (CPE) devices.

spanning tree

A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.

Spanning Tree Group (STG)

A collection of ports in one spanning-tree instance.

time-to-live (TTL)

The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.

Trivial File Transfer Protocol (TFTP)

A protocol that governs transferring files between nodes without protection

against packet loss.

trunk

A logical group of ports that behaves like a single large port.

Universal/Local (U/L)

Determines global and local link addresses; used with the Extended Unique

Identifier (EUI).

User Datagram Protocol (UDP)

In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application

programs.

variable-length subnet masking (VLSM) Allocating IP addressing resources to subnets according to their individual

need rather than some general network-wide rule.

virtual router

An abstract object managed by the Virtual Router Redundancy Protocol

(VRRP) that acts as a default router for hosts on a shared LAN.

virtual router forwarding (VRF)

Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF

as a separate physical router.

Virtual Router Redundancy Protocol (VRRP) A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup

router is quickly available to take its place.

Voice over IP (VOIP)

The technology that delivers voice information in digital form in discrete packets using the Internet Protocol (IP) rather than the traditional circuit-committed protocols of the public switched telephone network (PSTN).