



ExtremeSwitching™

Configuring VLANs, Spanning Tree, and NLB on VSP Operating System Software

NN47227-500
Issue 11.01
September 2017

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

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Extreme Networks, Inc. assumes no liability for any errors. Extreme Networks, Inc. reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published in varying mediums which may include product information, operating instructions and performance specifications that are generally made available to users of products. Documentation does not include marketing materials. Extreme Networks shall not be responsible for any modifications, additions, or deletions to the original published version of Documentation unless such modifications, additions, or deletions were performed by or on the express behalf of Extreme Networks. End User agrees to indemnify and hold harmless Extreme Networks, Extreme Networks' agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Extreme Networks is not responsible for the contents or reliability of any linked websites referenced within this site or Documentation provided by Extreme Networks. Extreme Networks is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Extreme Networks does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Extreme Networks provides a limited warranty on Extreme Networks hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Extreme Networks' standard warranty language, as well as information regarding support for this product while under warranty is available to Extreme Networks customers and other parties through the Extreme Networks Support website: <http://www.extremenetworks.com/support> under the link "Policies" or such successor site as designated by Extreme Networks. Please note that if You acquired the product(s) from an authorized Extreme Networks Channel Partner outside of the United States and Canada, the warranty is provided to You by said Extreme Networks Channel Partner and not by Extreme Networks.

"Hosted Service" means an Extreme Networks hosted service subscription that You acquire from either Extreme Networks or an authorized Extreme Networks Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Extreme Networks or Extreme Networks Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN EXTREME NETWORKS HOSTED SERVICE SUBSCRIPTION FROM EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE EXTREME NETWORKS WEBSITE, <https://extremeportal.force.com> OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU

REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE EXTREME NETWORKS WEBSITE, <https://extremeportal.force.com> OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS EXTREME NETWORKS SOFTWARE, PURCHASED FROM EXTREME NETWORKS, INC., ANY EXTREME NETWORKS AFFILIATE, OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY EXTREME NETWORKS IN WRITING, EXTREME NETWORKS DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN EXTREME NETWORKS, AN EXTREME NETWORKS AFFILIATE OR AN EXTREME NETWORKS CHANNEL PARTNER; EXTREME NETWORKS RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND EXTREME NETWORKS, INC. OR THE APPLICABLE EXTREME NETWORKS AFFILIATE ("EXTREME NETWORKS").

Extreme Networks grants You a license within the scope of the license types described below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License as set forth below in the Designated System(s) License (DS) section as applicable. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Extreme Networks or an Extreme Networks Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a set of Designated Processors that hosts (physically or virtually) a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

License type(s)

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only: 1) on a number of Designated Processors up to the number indicated in the order; or 2) up to the number of Instances of the Software as indicated in the order, Documentation, or as authorized by Extreme Networks in writing. Extreme Networks may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Extreme Networks through electronic means established by Extreme Networks specifically for this purpose.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Extreme Networks. All content on this site, the documentation, Hosted Service, and the product provided by Extreme Networks including the selection, arrangement and design of the content is owned either by Extreme Networks or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part,

including any code and software unless expressly authorized by Extreme Networks. Unauthorized reproduction, transmission, dissemination, storage, and/or use without the express written consent of Extreme Networks can be a criminal, as well as a civil offense under the applicable law.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note, unless otherwise stated, that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Extreme Networks Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Extreme Networks' website at: <http://www.extremenetworks.com/support/policies/software-licensing> or such successor site as designated by Extreme Networks. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE [HTTP://WWW.MPEGLA.COM](http://WWW.MPEGLA.COM).

Service Provider

THE FOLLOWING APPLIES TO EXTREME NETWORKS CHANNEL PARTNER'S HOSTING OF EXTREME NETWORKS PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN EXTREME NETWORKS CHANNEL PARTNER'S HOSTING OF EXTREME NETWORKS PRODUCTS MUST BE AUTHORIZED IN WRITING BY EXTREME NETWORKS AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE EXTREME NETWORKS CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS, AT THE EXTREME NETWORKS CHANNEL PARTNER'S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER.

WITH RESPECT TO CODECS, IF THE EXTREME NETWORKS CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE EXTREME NETWORKS CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE EXTREME NETWORKS CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE WWW.SIPRO.COM/CONTACT.HTML. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR

THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE [HTTP://WWW.MPEGLA.COM](http://WWW.MPEGLA.COM).

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Extreme Networks product is used.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Security Vulnerabilities

Information about Extreme Networks' security support policies can be found in the Global Technical Assistance Center Knowledgebase at <https://gtacknowledge.extremenetworks.com/>.

Downloading Documentation

For the most current versions of Documentation, see the Extreme Networks Support website: <http://documentation.extremenetworks.com>, or such successor site as designated by Extreme Networks.

Contact Extreme Networks Support

See the Extreme Networks Support website: <http://www.extremenetworks.com/support> for product or Hosted Service notices and articles, or to report a problem with your Extreme Networks product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Extreme Networks Support website: <http://www.extremenetworks.com/support/contact/> (or such successor site as designated by Extreme Networks), scroll to the bottom of the page, and select Contact Extreme Networks Support.

Contact Avaya Support

See the Avaya Support website: <https://support.avaya.com> for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: <https://support.avaya.com> (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Extreme Networks are the registered or unregistered Marks of Extreme Networks, Inc., its affiliates, its licensors, its suppliers, or other third parties. Users are not permitted to use such Marks without prior written consent from Extreme Networks or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Extreme Networks or the applicable third party.

Extreme Networks is a registered trademark of Extreme Networks, Inc.

All non-Extreme Networks trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

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

Contents

Chapter 1: Preface	9
Purpose.....	9
Training.....	9
Providing Feedback to Us.....	9
Getting Help.....	10
Extreme Networks Documentation.....	11
Chapter 2: New in this document	12
Notice about feature support.....	12
Chapter 3: VLAN configuration	13
VLAN fundamentals.....	13
Port-based VLANs.....	13
Private VLANs.....	14
Policy-based VLANs.....	17
SPBM B-VLAN.....	18
VLAN tagging and port types.....	18
VLAN router interfaces.....	21
IP routing and VLANs.....	21
VLAN implementation.....	21
VLAN configuration rules.....	22
VLAN feature support.....	22
Network Load Balancing	23
NLB and Directed Broadcast resource limits.....	27
VLAN MAC-layer filtering database and MAC security.....	27
Prevention of IP spoofing within a VLAN.....	28
VLAN loop prevention.....	29
Spanning tree and protection against isolated VLANs.....	31
IGMP Layer 2 Querier.....	32
Switched UNI Layer 3.....	33
VLAN configuration using CLI.....	34
Creating a VLAN.....	34
Creating a private VLAN.....	36
Assigning an IP address to a VLAN.....	39
Performing a general VLAN action.....	41
Configuring static MAC addresses for a VLAN.....	42
Limiting MAC address learning.....	43
Configuring the forwarding database timeout globally.....	44
Adding or removing ports in a VLAN.....	45
Adding or removing source MAC addresses for a VLAN.....	46
Configuring NLB support.....	47

Configuring a tagged port to discard untagged frames.....	48
Configuring SLPP.....	49
Configuring SLPP packet-rx on a port.....	50
Configuring SLPP packet-tx on a VLAN.....	52
Viewing SLPP information.....	53
Viewing SLPP information for a port.....	54
Configuring spoof detection.....	55
Viewing VLAN information.....	56
Viewing private VLAN information.....	57
Viewing brouter port information.....	58
Viewing VLAN port member status.....	59
Viewing VLAN source MAC addresses.....	60
Viewing VLAN forwarding database information.....	61
Viewing manual edit MAC addresses.....	62
Viewing port-level MAC security.....	63
Viewing NLB-mode information.....	63
Displaying C-VLAN and Switched UNI I-SID information.....	64
Enabling DVR on a Layer 2 VSN (VLAN)	67
VLAN configuration using EDM.....	68
Configuring the VLAN feature on a port.....	69
Viewing existing VLANs.....	70
Creating a port-based VLAN.....	71
Creating a private VLAN.....	73
Viewing Private VLAN information.....	75
Configuring an IP address for a VLAN.....	76
Changing VLAN port membership.....	77
Creating a protocol-based VLAN.....	77
Configuring source MAC addresses for a source MAC-based VLAN.....	78
Creating a SPBM B-VLAN.....	79
Configuring advanced VLAN features.....	80
Configuring NLB support.....	82
Configuring a port to accept tagged or untagged frames.....	83
Configuring untagging default VLAN on a tagged port.....	84
Configuring SLPP globally.....	84
Configuring the SLPP by VLAN.....	85
Configuring the SLPP by port.....	86
Configuring directed broadcast on a VLAN.....	88
Configuring the forwarding database timeout globally.....	89
Viewing VLAN forwarding database information.....	89
Viewing the forwarding database for a specific VLAN.....	90
Clearing learned MAC addresses by VLAN.....	91
Clearing learned MAC addresses for all VLANs by port.....	91
Viewing blocked MAC address information.....	92

Configuring static forwarding.....	92
Configuring limit learning.....	93
Chapter 4: Spanning Tree configuration.....	95
Spanning Tree fundamentals.....	95
BPDU Guard.....	97
Rapid Spanning Tree Protocol and Multiple Spanning Tree Protocol.....	98
Spanning Tree configuration using CLI.....	101
Configuring Spanning Tree.....	102
Configuring BPDU Guard.....	102
Configuring Rapid Spanning Tree Protocol.....	104
Configuring Rapid Spanning Tree Protocol for a port.....	105
Configuring the Rapid Spanning Tree Protocol version.....	107
Viewing the global RSTP configuration information.....	107
Viewing RSTP statistics.....	108
Viewing the RSTP status.....	108
Viewing the RSTP configuration information.....	109
Viewing the RSTP status for a port.....	110
Viewing RSTP information for a selected port.....	111
Viewing the RSTP role.....	112
Viewing spanning tree configuration.....	113
Configuring Multiple Spanning Tree Protocol.....	114
Configuring MSTP MSTI options.....	115
Configuring Ethernet MSTP.....	116
Configuring Ethernet MSTP MSTI.....	118
Viewing MSTP configurations.....	119
Viewing MSTP status.....	119
Viewing MSTP port information.....	120
Viewing MSTP MSTI information.....	121
Viewing MSTP statistics.....	122
Spanning Tree configuration using EDM.....	122
Configuring the Spanning Tree mode.....	123
Restarting the switch.....	123
Configuring BPDU Guard.....	124
Configuring RSTP global parameters.....	128
Configuring RSTP ports.....	130
Viewing RSTP port status.....	131
Configuring MSTP global parameters.....	132
Configuring CIST ports for MSTP.....	135
Configuring MSTI bridges for MSTP.....	138
Configuring MSTI ports for MSTP.....	139
Chapter 5: VXLAN Gateway configuration.....	141
VXLAN Gateway fundamentals.....	141
VXLAN Gateway licensing.....	143

- VXLAN modes 143
- Types of VXLAN Gateway deployments..... 144**
- VXLAN Gateway considerations and limitations..... 147
- VXLAN Gateway configuration using the CLI..... 150
 - Configuring the VXLAN Gateway boot flag..... 150
 - Configuring VXLAN Gateway..... 150
 - Flushing the MAC forwarding table..... 153
 - VXLAN Gateway show commands..... 154
- VXLAN Gateway configuration using EDM..... 157
 - Configuring the VXLAN Gateway boot flag..... 157
 - Configuring a VTEP source IP address..... 157
 - Configuring a remote VTEP..... 158
 - Configuring a VNID..... 159
 - Configuring VNID endpoints..... 160
 - Configuring ELAN endpoints..... 160
 - Displaying the VTEP next hop..... 161
 - Displaying the VNID forwarding database..... 162
- VXLAN Gateway configuration example..... 162
- Glossary..... 165**

Chapter 1: Preface

Purpose

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

- Extreme Networks Virtual Services Platform 4000 Series
- Extreme Networks Virtual Services Platform 7200 Series
- Extreme Networks Virtual Services Platform 8000 Series (includes VSP 8200 and VSP 8400 Series)
- Extreme Networks Virtual Services Platform 8600

This document contains procedural and conceptual information to help you configure and manage Virtual Local Area Networks (VLAN), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) on the VOSS platforms. This document also provides instructions to use Command Line Interface (CLI) and Enterprise Device Manager (EDM).

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.

Training

Ongoing product training is available. For more information or to register, you can access the Web site at www.extremenetworks.com/education/.

Providing Feedback to Us

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 especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short [online feedback form](#). You can also email us directly at internalinfodev@extremenetworks.com

Getting Help

Product purchased from Extreme Networks

If you purchased your product from Extreme Networks, use the following support contact information to get help.

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

- [GTAC \(Global Technical Assistance Center\) for Immediate Support](#)
 - Phone: 1-800-998-2408 (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
 - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) – Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) – A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) – Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) 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

Product purchased from Avaya

If you purchased your product from Avaya, use the following support contact information to get help.

Go to the Avaya Support website at <http://support.avaya.com> for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service

request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Extreme Networks Documentation

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for previous versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/software-licensing.

Chapter 2: New in this document

The following sections detail what is new in *Configuring VLANs, Spanning Tree, and NLB* since issue 08.xx.

There are no feature changes in this document.

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 appear on your hardware, it is not supported.

For information about feature support, see *Release Notes*.

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

Chapter 3: VLAN configuration

This chapter provides concepts and procedures to configure the virtual local area network (VLAN) features supported on the switch.

For information about the user-configuration interfaces, see *Using CLI and EDM*.

VLAN fundamentals

A VLAN is a switched network that is logically segmented by functions, project teams, or applications without regard to the physical location of users. By using a VLAN, you can divide the Local Area Network into smaller groups without interfering with the physical network.

The practical applications of VLAN include the following:

- You can create VLANs, or workgroups, for common interest groups.
- You can create VLANs, or workgroups, for specific types of network traffic.
- You can add, move, or delete members from these workgroups without making physical changes to the network.

By dividing the network into separate VLANs, you can create separate broadcast domains. This arrangement conserves bandwidth, especially in networks supporting broadcast and multicast applications that flood the network with traffic. A VLAN workgroup can include members from a number of dispersed physical segments on the network, improving traffic flow between them.

The switch performs the Layer 2 switching functions necessary to transmit information within VLANs, as well as the Layer 3 routing functions necessary for VLANs to communicate with one another. You can define a VLAN for a single switch or spanning multiple switches. A port can be a member of multiple VLANs. A VLAN is associated with a spanning tree group.

A VLAN packet is classified before it is forwarded. If the packet matches a classification rule, the port membership is checked. If the port is not an allowed member (potential, static, or active), the system drops the packet.

Port-based VLANs

A port-based VLAN is a VLAN in which you explicitly configure the ports to be in the VLAN. When you create a port-based VLAN on a device, you assign a VLAN identification number (VLAN ID) and specify the ports that belong to the VLAN. These port members are always active port members.

The VLAN ID is used to coordinate VLANs across multiple switches. Any type of frame can be classified to a port-based VLAN.

The example in the following figure shows two port-based VLANs: one for the marketing department, and one for the sales department. Ports are assigned to each port-based VLAN. A change in the sales area can move the sales representative at port 1/1 to the marketing department without moving cables. With a port-based VLAN, you only need to indicate in the Command Line Interface (CLI) that port 1/1 in the sales VLAN now is a member of the marketing VLAN.




Figure 1: Port-based VLAN

Private VLANs

Private VLANs provide isolation between ports within a Layer-2 service.

The primary and secondary VLAN make the private VLAN. Standard VLAN configuration takes place on the primary VLAN. The secondary VLAN is virtual and inherits configuration from the primary VLAN.

Ports in the private VLAN are configured as isolated, promiscuous, or trunk. The default value is None.

Port types

Table 1: Port types for private VLANs

Port type	Description
Promiscuous (tagged or untagged ports)	Promiscuous ports communicate with all other ports within the private VLAN. Uses the primary VLAN.
Isolated (tagged or untagged ports)	Isolated ports communicate with the promiscuous ports, but not with any other isolated port. Uses the secondary VLAN.
Trunk (tagged ports)	Trunk ports carry traffic between other port members within the private VLANs. Accepts either primary or secondary VLAN.

Trunk ports must have VLAN encapsulation enabled. A port may be a single port or may belong to an MLT.

The following figure shows a basic private VLAN topology with private VLAN configured on five switches. All ports connecting to other switches are trunk type ports and all other ports are either promiscuous or isolated ports. On the secondary VLAN, spokes can communicate with hubs, hubs can communicate with all spokes in the same private VLAN using the primary VLAN, but spokes cannot communicate with other spokes.




Figure 2: Private VLAN topology

E-Tree

The E-Tree allows private VLANs to traverse the Shortest Path Bridging MAC (SPBM) network.

For more information about E-Tree and SPBM configuration, see *Configuring Fabric Basics and Layer 2 Services*.

Private VLAN configuration rules

The following are private VLAN rules for the switch:

- Use private VLANs for Layer 2 services only
- Forwarding is based on MAC address based lookups
- IP routing and creation of IP interfaces are not supported on private VLANs
- Configuration of IP Source Guard (IPSG) is not supported on ports that are members of private VLANs.
- Do not use the untag-port default vlan parameter on private VLAN interfaces that are operating as trunk ports, because it impacts the private VLAN functionality.

Policy-based VLANs

Received frames are classified into a policy-based VLAN based on certain fields of the frame that matches the associated VLAN policy.

Port membership types

In a policy-based VLAN, a port can be designated as a potential member, a static member, or one not allowed to be a member of the VLAN.

If a port is designated as a potential member of the VLAN, and the incoming traffic matches the policy, the system dynamically adds the port to the active port list of the VLAN, making the port an active member of the VLAN. After the system adds a port to the active list, it can remove the port from the active list due to time-out. Potential member ports that join the VLAN are removed (timed out) from the active port list of the VLAN after the timeout (aging time) period expires.

All members of the Spanning Tree Group associated with a protocol-based VLAN are automatically considered potential members of the VLAN. In addition, all tagged ports (trunk ports) become static ports. If you do not want all the tagged ports to be static members of a protocol-based VLAN, put the port in the disallowed list.

Static port members are always members of the VLAN. Static port members are not aged out due to inactivity and they are not removed from the active list. If a server or router connects to a port, designate that port as a static member of a VLAN. If a server connects to a port that is only a potential member and the server sends very little traffic, a client fails to reach the server if the server port is timed out of the VLAN. It is recommended that you make these ports static members of the VLAN.

A disallowed port can never become a member of the VLAN until you add it as a port-member. After you remove a port from the VLAN, the system adds the port to the disallowed list.

On any single spanning-tree instance, an access (untagged) port can belong to one port-based VLAN and many policy-based VLANs. A trunk (tagged) port can belong to many port-based and policy-based VLANs.

The following table describes port membership types for policy-based VLANs.

Table 2: Port membership types for policy-based VLANs

Membership type	Description
Potential	Potential members of a VLAN become active members upon receiving data matching the policy defined for the VLAN (a packet tagged with that VLAN, or an untagged packet matching the policy).
Static (always a member)	Static members are always active members of the VLAN after you configure them as belonging to that VLAN.
Not allowed to join (never a member)	Ports of this type cannot join the VLAN.

The following table lists supported policy-based VLANs.

Table 3: Supported policy-based VLAN types

VLAN type	Support
Protocol-based	supported

Protocol-based VLANs

Protocol-based VLANs are an effective way to segment your network into broadcast domains according to the network protocols in use.

A port member of a port-based VLAN can belong to multiple protocol-based VLANs. Port tagging is not required for a port to be a member of multiple protocol-based VLANs.

The switch supports IPv6 protocol-based VLAN only.

SPBM B-VLAN

Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network.

This VLAN is used for both control plane traffic and dataplane traffic.

* Note:

Always configure two B-VLANs in the core to allow load distribution over both B-VLANs.

SPBM alters the behavior of the VLAN. When a B-VLAN is associated with an SPBM network the following VLAN attributes and behaviors are modified for the B-VLAN:

- Flooding is disabled
- Broadcasting is disabled
- Source MAC address learning is disabled
- Unknown MAC discard is enabled

Ports cannot be added to a B-VLAN manually, IS-IS takes care of adding ports to the B-VLAN.

Essentially the B-MAC addresses are programmed into the B-VLAN Forwarding Information Bases (FIBs) by IS-IS instead of the traditional VLANs flooding and learning approach.

Modification of the VLAN behavior is necessary to ensure proper control over the SPBM traffic.

VLAN tagging and port types

The switch supports the IEEE 802.1Q specification for tagging frames and coordinating VLANs across multiple switches.

[Figure 3: VLAN tag insertion](#) on page 19 shows how an additional four octet (tag) header is inserted in a frame after the source address and before the frame type. The tag contains the VLAN ID associated with the frame.




Figure 3: VLAN tag insertion

802.1Q tagged ports

Tagging a frame adds four octets to a frame, possibly making it bigger than the traditional maximum frame size. If a device does not support IEEE 802.1Q tagging, it can have problems interpreting tagged frames that it receives.

Whether tagged frames are sent depends on what you configure at the port level. Tagging is configured as true or false for the port, and is applied to all VLANs on that port.

A port with tagging enabled applies the VLAN ID tag to all packets sent on the port. Tagged ports are typically used to multiplex traffic belonging to multiple VLANs to other IEEE 802.1Q-compliant devices.

If you disable tagging on a port, it does not send tagged frames. A nontagged port connects a switch to devices that do not support IEEE 802.1Q tagging. If a tagged frame is forwarded to a port with tagging configured to false, the switch removes the tag from the frame before sending it to the port.

Treatment of tagged and untagged frames

The switch associates a frame with a VLAN based on the data content of the frame and the configuration of the receiving port. The treatment of the frame depends on whether the frame is tagged or untagged.

If a tagged frame is received on a port, if the port is a static or potential member of the VLAN ID specified in the tag, the switch directs it to that VLAN. If the port is not a member of the VLAN that is identified by the tag in the packet, the switch discards the packet. If a port is untagged, you can configure it to discard tagged frames received on the port. In this case the tagged frame is discarded.

For untagged frames, VLAN membership is implied from the content of the frame itself. You can configure a tagged port to accept or discard untagged frames received on the port.

The default VLAN of a port is the VLAN to which untagged frames are classified if they do not match the criteria of any policy-based VLAN of which the port is a member. The default VLAN of the port can be any port-based VLAN a port belongs to, or the unassigned VLAN (1). Frames classified to the unassigned VLAN are discarded.

The frame is forwarded based on the VLAN on which the frame is received, and on the forwarding options available for that VLAN. The switch tries to associate untagged frames with a VLAN in the following order:

- Does the frame belong to a protocol-based VLAN?
- What is the default VLAN for the receiving port?
- Is the default VLAN for the port not the unassigned VLAN?

If the frame meets none of these criteria, it is discarded.

Untagging default VLAN on a tagged port feature

This feature provides the ability to connect two devices such as an IP phone and a PC to a single port of the switch. Most IP phones ship with an embedded three port switch, and traffic coming from the phone is generally tagged (VLAN ID configured statically or remotely). However, the traffic originating from a PC is usually untagged traffic and must be separated from the IP phone traffic. This separation ensures that broadcast traffic from the PC does not impact voice quality.

After an IP phone is attached to an untagged port, it can fail to register with a remote Internet Telephony Gateway (or equivalent device) dependent on the netmask of the destination IP address (Call Server subnet).

For more information about the Network with IP phone and PC, see the following figure.




Figure 4: Network with IP phone and PC

IP phones and PCs coexist on the same port due to the use of an embedded IP Phone Layer 2 switch. In this scenario if you configure the port as untagged, the egress traffic on this port is untagged and no separation exists between the traffic to the IP phone and the PC. To avoid this condition, the port that connects to the IP phone must be tagged. If the port is tagged, the traffic for the PC is tagged with the default VLAN ID for the port. This configuration creates a problem because the PC does not expect tagged packets. Untag the default VLAN on a tagged port (in this example, port 1/1 that connects to the IP phone) to ensure that the traffic to the PC is sent untagged.

VLAN router interfaces

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. This IP address is not associated with a physical port. You can reach the VLAN IP address through any of the VLAN port members. Frames are routed to another VLAN IP address within the device. A port can belong to multiple VLANs; some, all, or none can perform routing.

IP routing and VLANs

The switch supports IP routing on the following types of VLANs:

- Port-based VLANs
- IP protocol-based VLANs

VLAN implementation

This section describes how to implement VLANs and describes default VLANs, the unassigned (NULL) VLAN, and brouter ports. This section also summarizes the defaults and rules regarding VLAN creation on the switch.

- [Default VLAN](#) on page 21
- [NULL VLAN](#) on page 21
- [Brouter ports](#) on page 21

Default VLAN

Devices are factory-configured so that all ports are in a port-based VLAN called the default VLAN. Because all ports are in the default VLAN, the device behaves like a Layer 2 device. The VLAN ID of this default VLAN is always 1, and it is always a port-based VLAN. You cannot delete the default VLAN.

NULL VLAN

Internally, the switch creates a special port-based VLAN called NULL VLAN or unassigned VLAN. This is a place holder VLAN for ports that are not members of any port-based VLAN. When a port is removed from all port-based VLANs, it is added to the NULL VLAN as a port member. Ports can belong to policy-based VLANs as well as to the NULL VLAN. If a frame does not meet the policy criteria and no underlying port-based VLAN exists, the port belongs to the NULL VLAN and the frame is dropped.

Because it is an internal construct, the NULL VLAN cannot be deleted.

Brouter ports

A brouter port is actually a one-port VLAN with an IP interface. The difference between a brouter port and a standard IP protocol-based VLAN configured to perform routing 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. Because a brouter port is a single-port VLAN, it uses one VLAN ID. Each brouter port decreases the number of available VLANs by one.

VLAN configuration rules

The following are VLAN rules for the switch:

- The switch supports configurable VLANs in the range of 1 to 4059. VLAN 0 is invalid. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. VLAN IDs on the switch range from 2 to 4084 but, 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.
- A tagged port can belong to multiple VLANs in multiple Spanning Tree Groups.
- Under the default configuration, the default Spanning Tree Group is number 1 if the chassis configuration permits multiple STGs.
- An untagged port can belong to only one port-based VLAN.
- You can configure only one protocol-based VLAN for a given protocol.
- The VLAN membership of a frame is determined by the following order of precedence, if applicable:
 1. IEEE 802.1Q tagged VLAN ID
 2. protocol-based VLAN
 3. port-based VLAN default VLAN of the receiving port

VLAN feature support

The following table summarizes supported features.

For the latest scalability information, see *Release Notes*.

Table 4: VLAN support

Feature	Description
Number of VLANs	4059
Port-based VLANs	Supported
Policy-based VLANs <ul style="list-style-type: none"> • Protocol-based • SPBM-based 	Supported
IEEE 802.1Q tagging	Supported

Table continues...

Feature	Description
IP routing and VLANs	Supported
Special VLANs <ul style="list-style-type: none"> • Default VLAN • Null VLAN • Brouter ports • Private VLAN 	Supported

Network Load Balancing

Microsoft Network Load Balancing (NLB) is a clustering technology available with the Microsoft Windows 2000, Microsoft Windows 2003, Microsoft Windows 2008, and Microsoft Windows 2012 Server family of operating systems. You can use NLB to share the workload among multiple clustering servers. NLB uses a distributed algorithm to load balance TCP/IP network traffic across a number of hosts, enhancing the scalability and availability of mission critical, IP based services, such as Web, VPN, streaming media, and firewalls. Network Load Balancing also provides high availability by detecting host failures and automatically redistributing traffic to remaining operational hosts.

NLB considerations and restrictions

Although the switch interoperates with NLB clusters operating in Unicast mode and Multicast mode, the following restrictions apply:

- VSP 7200 and VSP 8000 do not support true egress mirroring because packets are mirrored prior to the completion of packet processing, so egress mirrored packets can differ from the packets egressing the port. Therefore, NLB Unicast traffic and NLB Multicast traffic destined to the NLB Virtual Mac Address are not mirrored correctly.

Note:

To mirror the egress traffic of VSP 7200 and VSP 8000 platforms, you can use the NEXT-hop device ingress mirroring to capture the egress packets of the switch.

Of the VOSS platforms, only VSP 4000 supports true egress mirroring.

- Inter-VRF routing is not supported between an NLB client and an NLB cluster VLAN in Unicast mode or Multicast mode.
- You must configure NLB to use the same mode as the NLB Server.
- Static ARP entries are not supported for NLB Unicast or NLB Multicast.
- For interoperability with NLB, the switch provides configuration options at the VLAN level.
- ARP entries for NLB server IP addresses do not age out when there is still client traffic coming to the NLB servers, even after the NLB servers are no longer reachable.

NLB clustering in unicast mode

When the cluster is running in NLB unicast mode, all servers in the cluster share a common virtual MAC address, which is 02-bf-x-x-x-x (where x-x-x-x is the cluster IP address in hexadecimal form).

All traffic destined to this MAC address is sent to all the servers in the cluster. The virtual MAC address is specified in the Sender MAC Address field of the Address Resolution Protocol (ARP) reply from the cluster to the switch. ARP responses from the switch are sent to the virtual MAC address (rather than to the hardware MAC address).

You can configure the switch for NLB unicast mode support. After you enable the NLB unicast option, the switch floods traffic destined to the cluster IP address to all ports on the VLAN. Unicast mode supports connectivity to a secondary virtual IP address. For information about software scaling capabilities in unicast mode, see *Release Notes*.

NLB clustering in multicast mode

When the cluster is running in NLB multicast mode, a multicast virtual MAC address with the format 03-bf-x-x-x-x (where x-x-x-x is the cluster IP address in hexadecimal form) is bound to all cluster hosts but the real MAC address of the network adapter is retained. The multicast MAC address is used for client-to-cluster traffic and the real MAC address of the adapter is used for network traffic specific to the host server.

You can configure the switch for NLB multicast mode support. When you enable NLB multicast mode on a VLAN, the routed traffic destined to the NLB cluster is flooded by default on all ports of the VLAN. All VLANs support multiple cluster IPs by default. You can connect up to 200 NLB clusters to a single VLAN. For information about software scaling capabilities, see *Release Notes*.

Note:

SPBM supports NLB Unicast and Multicast modes. For more information on SPBM, see *Configuring Fabric Basics and Layer 2 Services*.

Supported NLB topologies

The switch supports Network Load Balancing (NLB) in the following topologies.

Supported NLB topology—example 1

The switch supports NLB when the NLB Cluster connections use a different physical port on the switch than the NLB clients.

The following figure illustrates this configuration where the NLB Server and the NLB Client workstations connect to different aggregation switches, which connect to the switch using different VLANs and different ports.

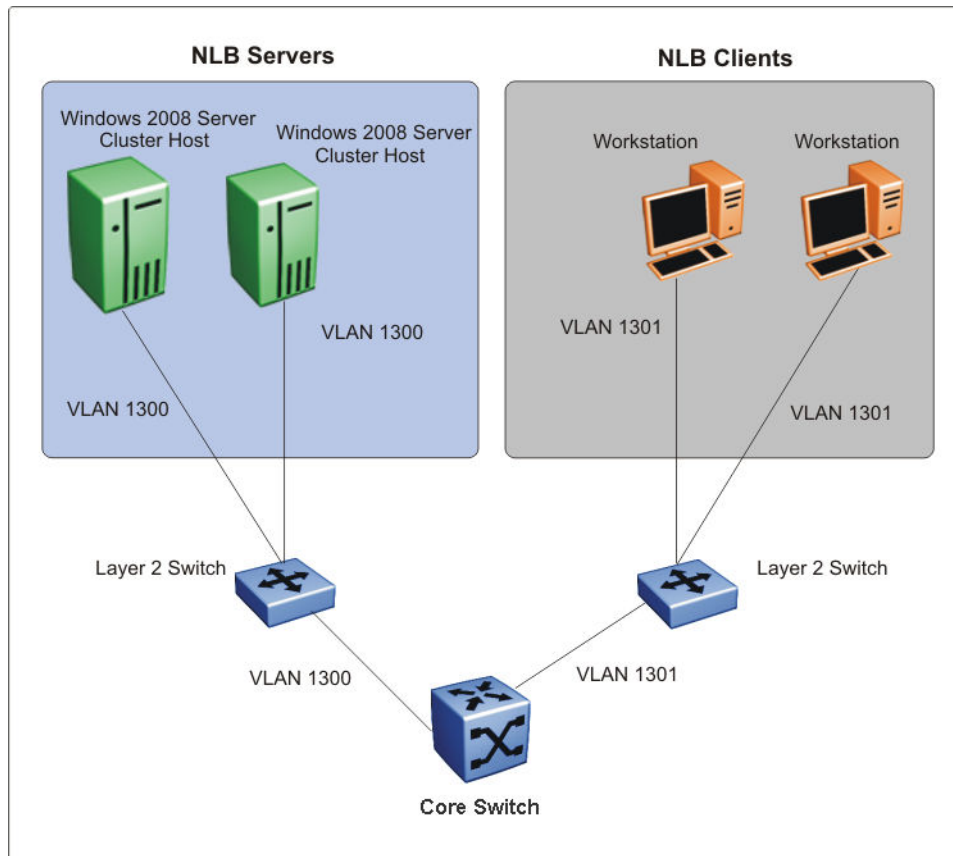


Figure 5: Supported NLB topology—example 1

Supported NLB topology—example 2

The switch also supports the following topology where the NLB Server and the NLB Client workstations connect to the same aggregation switch and then connect to the switch using the same port.

! Important:

The switch supports Layer 3 routing between an NLB-enabled VLAN and another VLAN on the same port.

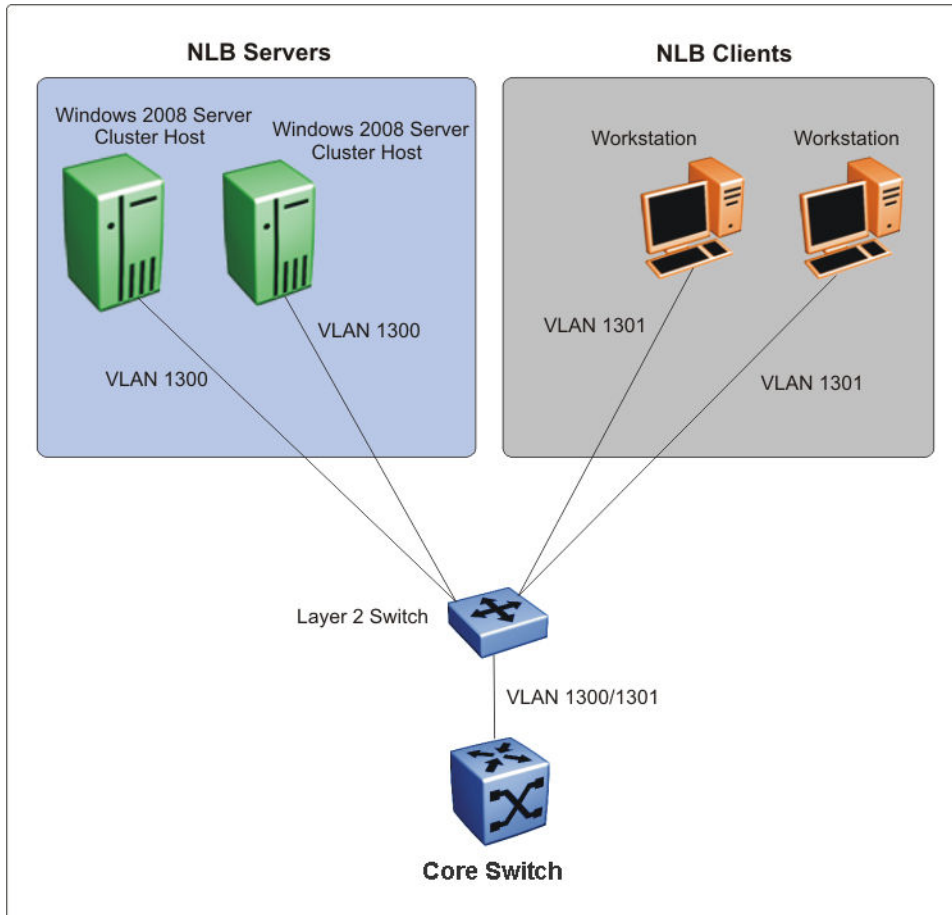


Figure 6: Supported NLB topology—example 2

Other supported NLB topologies

The switch supports NLB in the following other topologies:

- NLB cluster hosts and clients are connected to Layer 2 Ethernet switches that are SMLT connected to the SMLT cluster.
- NLB cluster hosts are directly connected and distributed between the switches and the clients are connected to Layer 2 Ethernet switch that is SMLT connected to the SMLT cluster.
- NLB cluster hosts and clients are directly connected and distributed between the switches in the SMLT cluster.
- NLB cluster hosts and clients are connected to Layer 2 Ethernet switches that are SMLT connected to the SMLT cluster core.

*** Note:**

For more information on the above topologies, see *Technical Configuration Guide for Microsoft Network Load Balancing*, NN48500-593.

NLB and Directed Broadcast resource limits

NLB and Directed Broadcast consume resources from the same pool of 200 resources. When you configure either NLB or Directed Broadcast, the switch uses one resource. If you configure both NLB and Directed Broadcast, the switch uses two resources.

To avoid a situation where there is a lack of resources, adhere to the following limits:

- The number of NLB cluster IP interfaces multiplied by the number of configured clusters must be equal to, or less than, 200. The number of NLB cluster IP interfaces is the *key*, not the number of VLANs. You can configure 1 VLAN with up to 200 NLB cluster IP interfaces or configure up to 200 VLANs with 1 NLB cluster IP interface per VLAN.

For example: 1 NLB cluster IP interface x 200 clusters = 200 or 2 NLB cluster IP interfaces x 100 clusters = 200

- If you configure VLANs with Directed Broadcast only, you can scale up to 200 VLANs.
- If you configure VLANs with both **NLB** and **Directed Broadcast**, you can only scale up to 100 VLANs assuming there is only 1 NLB cluster IP interface per VLAN.

For information on Directed Broadcast, see *Configuring Security*.

VLAN MAC-layer filtering database and MAC security

To perform MAC-layer bridging, the device must know the destination MAC-layer address of each device on each attached network, so it can forward packets to the appropriate destination. MAC-layer addresses are stored in the bridge forwarding database (FDB) table, and you can forward packet traffic based on the destination MAC-layer address information.

Note:

This feature is not supported on all hardware platforms. For more information about feature support, see *Release Notes*.

MAC security

Use MAC security to control traffic from specific MAC addresses. You can also limit the number of allowed MAC addresses. You can enable this feature at the port level.

Port—level security applies to traffic for all VLANs received on that port.

Port-level MAC security provides limit—learning option:

- **limit-learning:** This option protects the FDB from traffic from too many MAC addresses, which fill the FDB table.

This option limits the number of MAC addresses a port learns. You can specify a maximum and minimum number of addresses. After the number of addresses exceeds the maximum, learning stops. MAC address learning resumes after enough existing addresses age out and there is room to learn new MAC addresses. This option does not affect packet forwarding; it limits only MAC learning.

! Important:

Do not enable limit-learning and auto-learning for a port simultaneously.

Prevention of IP spoofing within a VLAN

VLAN IP as the default gateway

You can prevent VLAN logical IP spoofing by blocking the external use of the device IP address. A configurable option is provided, for each port, which detects a duplicate IP address (that is, an address that is the same as the device VLAN IP address) and blocks all packets with a source or destination address equal to that address.

*** Note:**

This feature is not supported on all hardware platforms. For more information about feature support, see *Release Notes*.

If an ARP packet is received that has the same source IP address as the logical VLAN IP address of the receiving port, all traffic coming to that port (with this MAC address as source/destination address) is discarded by the hardware. After detecting a duplicate IP address, the device sends a gratuitous ARP packet to inform devices on the VLAN about the correct MAC address for that IP address. You can specify a time on a configurable global timer after which the MAC discard record is deleted, and the device resumes accepting packets from that MAC address.

VRRP IP as the default gateway

Similarly, you can prevent VRRP IP spoofing by blocking the external use of the virtual IP address. A configurable option is provided, for each port, which detects a duplicate IP address (that is, an address that is the same as the device virtual IP address) and blocks all packets with a source or destination address equal to that address.

If an ARP packet is received that has the same source IP address as the virtual IP address of the receiving port, all traffic coming to that port (with this MAC address as source/destination address) is discarded by the hardware. After detecting a duplicate IP address, the device sends a gratuitous ARP packet to inform devices on the VRRP subnet about the correct virtual router MAC address for that IP address. You can specify a time on a configurable global timer after which the MAC discard record is deleted, and the device resumes accepting packets from that MAC address.

Packet spoofing

You can stop spoofed IP packets by configuring the switch to forward only IP packets that contain the correct source IP address of your network. By denying all invalid source IP addresses, you minimize the chance that your network is the source of a spoofed DoS attack.

A spoofed packet is one that comes from the Internet into your network with a source address equal to one of the subnet addresses on your network. The source address belongs to one of the address blocks or subnets on your network. To provide spoofing protection, you can use a filter that examines the source address of all outside packets. If that address belongs to an internal network or a firewall, the packet is dropped.

To prevent DoS attack packets that come from your network with valid source addresses, you need to know the IP network blocks in use. You can create a generic filter that:

- Permits valid source addresses

- Denies all other source addresses

To do so, configure an ingress filter that drops all traffic based on the source address that belongs to your network.

If you do not know the address space completely, it is important that you at least deny private (see RFC1918) and reserved source IP addresses. The following table lists the source addresses to filter.

Table 5: Source addresses to filter

Address	Description
0.0.0.0/8	Historical broadcast. High Secure mode blocks addresses 0.0.0.0/8 and 255.255.255.255/16. If you enable this mode, you do not need to filter these addresses.
10.0.0.0/8	RFC1918 private network
127.0.0.0/8	Loopback
169.254.0.0/16	Link-local networks
172.16.0.0/12	RFC1918 private network
192.0.2.0/24	TEST-NET
192.168.0.0/16	RFC1918 private network
224.0.0.0/4	Class D multicast
240.0.0.0/5	Class E reserved
248.0.0.0/5	Unallocated
255.255.255.255/32	Broadcast1

You can also enable the spoof-detect feature on a port.

VLAN loop prevention

Note:

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Loop prevention

Under certain conditions, such as incorrect configurations or cabling, loops can form. This is true mainly for layer 2 bridged domains, such as VLANs.

Simple Loop Prevention Protocol (SLPP) provides active protection against Layer 2 network loops on a per-VLAN basis. SLPP uses a lightweight hello packet mechanism to detect network loops. Sending hello packets on a per VLAN basis allows SLPP to detect VLAN based network loops for untagged as well as tagged IEEE 802.1Q VLAN link configurations. After SLPP detects a loop, the port is shutdown.

*** Note:**

If SLPP is used in a vIST environment, it must be enabled on both the vIST peers. Because, when an SLPP packet of a vIST peer is looped through UNI ports to the other device, that device will shut down its UNI port due to receiving SLPP packets from its peer. A device's own SLPP packets will go over a vIST connection but will not be forwarded by its vIST peer back onto its UNI ports.

Configure the SLPP functionality with the following criteria:

- **SLPP TX Process** – You decide on which VLANs a switch can send SLPP hello packets. The packets are then replicated out all ports which are members of the SLPP-enabled VLAN. It is recommended that you enable SLPP on all VLANs.
- **SLPP RX Process** – You decide on which ports the switch can act when receiving an SLPP packet that is sent by the same switch or by its SMLT peer. You must enable this process only on Access SMLT ports. You can enable this process only when the design permits on SMLT CORE ports in the case of a square/full mesh core design.
- **SLPP Action** – The action operationally disables the ports receiving the SLPP packet. You can also tune the network failure behavior. You can choose how many SLPP packets a port needs to receive before a switch takes an action. You need to stagger these values to avoid edge switch isolation – see the recommendations at the end of this section.

Loops can be introduced into the network in many ways. One way is through the loss of an MLT/link aggregation configuration caused by user error or malfunctioning equipment. This scenario does not always introduce a broadcast storm, but because all MAC addresses are learned through the looping ports, does significantly impact Layer 2 MAC learning. Spanning Tree cannot in all cases detect such a configuration issue, whereas SLPP reacts and disables the malfunctioning links and limits network impact to a minimum.

The desire is to prevent a loop from causing network problems, while also attempting not to isolate totally the edge where the loop was detected. Total edge closet isolation is the last resort to protect the rest of the network from the loop. With this in mind, some administrators adopt the concept of an SLPP primary switch and SLPP secondary switch. These are strictly design terms and are not configuration parameters. The Rx thresholds are staggered between the primary and secondary switch. Therefore, the primary switch disables an uplink immediately upon a loop occurring. If this resolves the loop issue, then the edge closet still has connectivity back through the SLPP secondary switch. If the loop is not resolved, then the SLPP secondary switch disables the uplink and isolates the closet to protect the rest of the network from the loop.

As the number of VLANs running SLPP scale off of a specific uplink port, the Rx-threshold value may need to be increased to prevent complete isolation of the offending edge. The primary goal of SLPP is to protect the core at all costs. In certain loop conditions, what can occur is the secondary switch also detects the loop and SLPP Rx-threshold of the secondary switch is reached before the primary can stop the loop by taking its port down. Therefore, both switches eventually take their ports down and the edge is isolated. The larger the number of VLANs associated with the port, the more likely this can occur, especially for loop conditions that affect all VLANs.

The loop detection functionality of the device must not be used under normal operating conditions. Only use it if directed by the technical support personnel.

You cannot configure the EtherType for SLPP. The switch uses an EtherType of 0x8102 .

Spanning tree and protection against isolated VLANs

Virtual Local Area Network (VLAN) isolation disrupts packet forwarding. The following figure illustrates the problem. Two VLANs (V1 and V2) connect four devices, and both VLANs are in the same spanning tree group. V2 includes three of the four devices, whereas V1 includes all four devices. After a spanning tree protocol detects a loop, it blocks the link with the highest link cost. In this case, the 100 Mbps link is blocked, which isolates a device in V2. To avoid this problem, either configure V2 on all four devices or use MSTP with a different Multiple Spanning Tree Instance (MSTI) for each VLAN.

*** Note:**

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

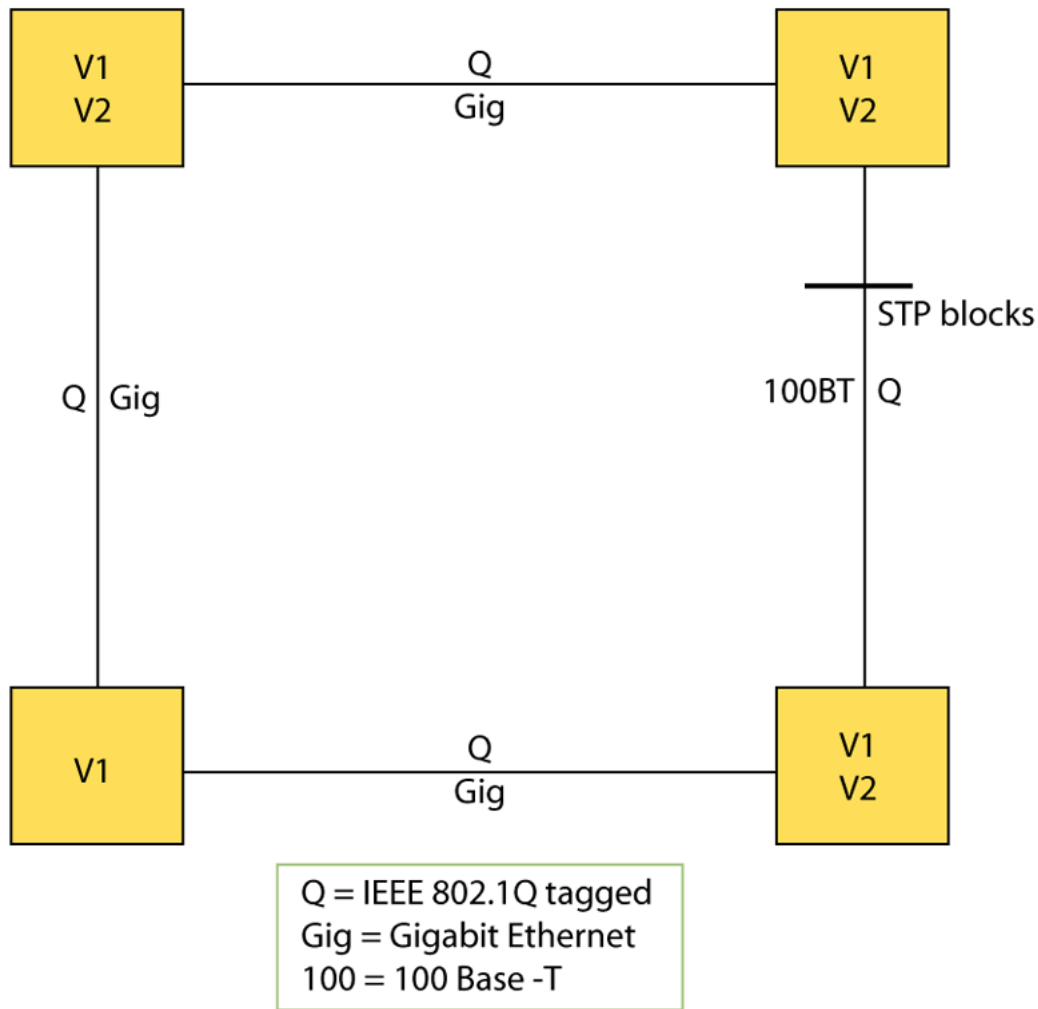


Figure 7: VLAN isolation

IGMP Layer 2 Querier

In a Layer 2 multicast network, you can enable Layer 2 querier on one of the switches in the VLAN. IGMP Layer 2 querier provides the IGMP querier function so that the switch can provide the recurring queries that maintain IGMP groups when you do not use multicast routing for multicast traffic.

Overview

In a multicast network, if you only need to use Layer 2 switching for the multicast traffic, you do not need multicast routing. However, you must have an IGMP querier on the network for multicast traffic to flow from sources to receivers. A multicast router provides the IGMP querier function. You can also use the IGMP Layer 2 Querier feature to provide a querier on a Layer 2 network without a multicast router.

The Layer 2 querier function originates queries for multicast receivers, and processes the responses accordingly. On the connected Layer 2 VLANs, IGMP snoop continues to provide services as normal. IGMP snoop responds to queries and identifies receivers for the multicast traffic.

You must enable Layer 2 querier and configure an IP address for the querier before it can originate IGMP query messages. If a multicast router exists on the network, the switch automatically disables the Layer 2 querier.

In a Layer 2 multicast network, enable Layer 2 querier on only one of the switches in the VLAN. A Layer 2 multicast domain supports only one Layer 2 querier. No querier election exists.

IGMP Snooping

IGMP Snooping enables Layer 2 switches in the network to examine IGMP control protocol packets exchanged between downstream hosts and upstream routers.

When Layer 2 switches examine the IGMP control protocol packets, they:

- Generate the Layer 2 MAC forwarding tables used for further switching sessions
- Regulate the multicast traffic to prevent it from flooding the Layer 2 segment of the network

IGMP Layer 2 Querier and IGMP interaction

IGMP Layer 2 Querier uses IGMP to learn which groups have members on each of the attached physical networks, and it maintains a list of multicast group memberships for each attached network and a timer for each membership. In this case, multicast group memberships means the presence of at least one member of a multicast group on a given attached network, not a list of all of the members.

IGMP Layer 2 Querier can assume one of two roles for each of the attached networks:

- Querier
- Non-Querier

After you enable IGMP Layer 2 Querier, the system assumes it is a multicast router, so it sends the General Query, Group Specific/Group, and Source Specific Query when Leave/BLOCK messages are received. IGMP queries are required to maintain an IGMP group.

* Note:

Group Specific When Leave does not apply to IGMPv1.

For more information about how to configure IGMP Layer 2 Querier, see *Configuring IP Multicast Routing Protocols*.

Switched UNI Layer 3

Create a platform VLAN using the command `vlan create <vlan-id> type port-mstprsp <msti-instance>`. Enable Layer 3 services on the platform VLAN and is associated with the Switched UNI (S-UNI) Service Instance Identifier (I-SID). All S-UNI ports are added to the platform VLAN.

You must associate the S-UNI I-SID to the platform VLAN. After you associate the platform VLAN with the I-SID, it becomes a CVLAN.

The switch performs MAC and ARP learning on the platform VLAN.

*** Note:**

You cannot add S-UNI ports or MLT to the S-UNI platform VLANs directly. Add the ports to the I-SID and assign the I-SID to the platform VLAN.

*** Note:**

You can associate only port based VLAN with S-UNI I-SID.

VLAN configuration using CLI

This chapter describes how to configure and manage a virtual local area network (VLAN) by using Command Line Interface (CLI).

Configure and manage a VLAN to create VLANs, including private VLANs, add or remove ports in the VLAN, configure priority, change a VLAN name, or perform other operations.

Creating a VLAN

Create a VLAN by port, protocol, or SPBM. Optionally, you can choose to assign the VLAN a name and color.

Assign an IP address to the VLAN. You can also assign a MAC-offset value that allows you to manually change the default MAC address.

About this task

Create a VLAN and assign an IP address.

If you configure the SLA Mon agent address under an IP interface or VLAN, you must remove the SLA Mon address before you can remove the IP address or VLAN.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a VLAN using CLI:

```
vlan create <2-4059>
```

3. Specify a name for the VLAN:

```
vlan create <2-4059> name WORD<0-64>
```

4. Create a VLAN by port:

```
vlan create <2-4059> type port-mstprstp <0-63>
```

5. Associate CVLAN I-SID to the platform VLAN.

```
vlan i-sid <1-4059> <1-16777215>
```

6. Create a VLAN using a user-defined protocol and specify the frame encapsulation header type:

```
vlan create <2-4059> type protocol-mstprstp <0-63> ipv6
```

7. Assign a color to the VLAN:

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

8. Log on to the VLAN Interface Configuration mode for the VLAN ID in CLI:

```
interface VLAN <1-4059>
```

9. Assign an IP address to a VLAN:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D>
```

10. Specify the MAC-offset value:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D> [<0-511>]
```

Example

```
Switch:1> enable
```

```
Switch:1# configure terminal
```

```
Switch:1(config)# vlan create 2 type port-mstprstp 0 color 4
```

```
Switch:1(config)#vlan i-sid 2 100
```

```
Switch:1(config)# interface vlan 2
```

```
Switch:1(config-if)# ip address 192.0.2.0/24
```

Variable Definitions

Use the data in the following table to use the **vlan create** command.

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 <i>WORD</i> <0-64>	Specifies the VLAN name. The name attribute is optional.

Table continues...

Variable	Value
type port-mstprstp <0-63> [color <0-32>]	Creates a VLAN by port: <ul style="list-style-type: none"> • <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. <p>* Note:</p> <p>MSTI instance 62 is reserved for SPBM if SPBM is enabled on the switch.</p>
type pvlan-mstprstp <0-63> [color <0-32>]	Creates a private VLAN by port: <ul style="list-style-type: none"> • <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.
type protocol-mstprstp <0-63> ipv6	Creates a VLAN by protocol: <ul style="list-style-type: none"> • <0-63> is the STP instance ID. • color <0-32> is the color of the VLAN in the range of 0 to 32.
type spbm-bvlan	Creates a SPBM B-VLAN.

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

Variable	Value
<A.B.C.D/X> <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.
<0-511>	Specifies the MAC-offset value.

Use the data in the following table to use the **vlan i-sid** 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.
<0-16777215>	Specifies the i-sid number. The value is in the range of <0-16777215>.

Creating a private VLAN

About this task

You can create a private VLAN and set the port type. The primary and secondary VLAN IDs are associated with the same MTSI, the secondary VLAN inherits the primary VLAN configuration. You cannot create another VLAN with the same VLAN ID as the secondary VLAN. The secondary VLAN cannot be any other type of VLAN other than a secondary VLAN.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Create a private VLAN:

```
vlan create <2-4059> type pvlan-mstprstp secondary <2-4059>
```

3. Specify a name for the VLAN:

```
vlan create <2-4059> name
```

4. Enter GigabitEthernet Interface Configuration mode:

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

*** Note:**

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

5. Set the port type:

```
private-vlan <isolated|promiscuous|trunk>
```

*** Note:**

If the port is a member of an MLT, the port inherits the private VLAN port type of the MLT. For more information about creating MLTs and LACPs, see, *Configuring Link Aggregation, MLT, SMLT, and vIST*.

6. Exit to Global Configuration mode:

```
exit
```

7. Enter VLAN Interface Configuration mode:

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

8. Add ports to the primary VLAN:

VLAN configuration


```
vlan members add <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# vlan create 2 type pvlan-mstprstp 6 secondary 5
Switch:1(config)# interface gigabitethernet 1/36
Switch:1(config-if)# private-vlan isolated
Switch:1(config-if)# exit
Switch:1(config)# interface vlan 2
Switch:1(config-if)# vlan members add 2 1/36
```

Variable Definitions

Use the data in the following table to use the `vlan create` command.

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 <i>WORD</i> <0-64>	Specifies the VLAN name. The name attribute is optional.
type pvlan-mstprstp <0-63>	Creates a private VLAN by port. The variable <0-63> is the STP instance ID from 0 to 63.  Note: MSTI instance 62 is reserved for SPBM if SPBM is enabled on the switch.
secondary<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.

Use the data in the following table to use the `private vlan port type` command.

Variable	Value
<code><isolated promiscuous trunk></code>	Specifies the port type. If not specified, the port type defaults to None. <ul style="list-style-type: none"> • Isolated: An Isolated port can belong only to one private VLAN • Promiscuous: A Promiscuous port can belong to many private VLANs • Trunk: A Trunk port can belong to many private VLANs, is tagged, and can also belong to non-private VLANs
<code>no private-vlan</code>	Port defaults to type None.
<code>default private-vlan</code>	Port defaults to type None.

*** Note:**

If there are other non-private VLANs using the defined port, the following message is displayed: All non private VLANs using this interface will be removed once this port becomes a member of a private VLAN. Ports with private-vlan type of isolated or promiscuous may only contain private VLANs. Do you wish to continue (y/n) ?

Use the data in the following table to use the `interface vlan` and `vlan members add` commands.

Variable	Value
<code><1-4059></code>	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.

Assigning 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.

*** Note:**

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:

- For VSP 4000 series:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D> [<0-127>]
```

- For VSP 8000 and 7200 series:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D> [<0-511>]
```

- For VSP 8600 series:

```
ip address <A.B.C.D/X>|<A.B.C.D> <A.B.C.D> [<0-1535>]
```

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
```

Variable definitions

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

Variable	Value
<A.B.C.D/X> <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.
<0-127>	Specifies the MAC-offset value for VSP 4000 series.
<0-511>	Specifies the MAC-offset value for VSP 8000 and 7200 series.
<0-1535>	Specifies the MAC-offset value for VSP 8600.

Use the data in the following table to use the **vrf** command.

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

Performing a general VLAN action

Perform a general VLAN action to initiate a specific function on a VLAN, such as clearing learned MAC addresses or ARP entries from the forwarding database by performing this procedure.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Perform a general VLAN action:

```
vlan action <1-4059> {none|flushMacFdb|flushArp|flushIp|
flushDynMemb|triggerRipUpdate|all}
```

Example

Perform a general VLAN action:

```
Switch(config)# vlan action 1 none
Switch(config)# vlan action 1 flushMacFdb
Switch(config)# vlan action 1 flushIp
Switch(config)# vlan action 1 flushDynMemb
```

Variable definitions

Use the data in the following table to use the `vlan action` command.

Table 6: Variable definitions

Variable	Value
none	Configures action to none. This action performs no updates.
flushMacFdb	Configures action to flushMacFdb. This action removes the learned MAC addresses from the forwarding database for the selected VLAN.
flushArp	Configures action to flushArp. This action removes the ARP entries from the address table for the selected VLAN.
flushIp	Configures action to flushIp. This action removes the learned IP addresses from the forwarding table for the selected VLAN.
flushDynMemb	Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based

Table continues...

Variable	Value
	VLAN, and removes MAC addresses learned on those ports for this VLAN.
flushDynMemb	Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN, and removes MAC addresses learned on those ports for this VLAN.
triggerRipUpdate	Configures action to triggerRipUpdate.
all	Configures action to all and performs all preceding actions.

Configuring static MAC addresses for a VLAN

Configure the static MAC address parameters.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure a static MAC address of a VLAN:

```
vlan mac-address-static <1-4059> <0x00:0x00:0x00:0x00:0x00:0x00>
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

Example

Configure a static MAC address of a VLAN:

```
Switch(config)# vlan mac-address-static 1 0x00:0x00:0x00:0x00:0x00:0x01
1/1
```

Variable definitions

Use the data in the following table to use the `vlan mac-address-static` command.

Table 7: Variable definitions

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

Table continues...

Variable	Value
	reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
<0x00:0x00:0x00:0x00:0x00:0x00>	Indicates the MAC address.
{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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Limiting MAC address learning

Configure the MAC security feature to control traffic from specific MAC addresses. The total number of MAC addresses that you can configure are fixed. The switch supports a maximum of 32k MAC entries for non-SPBM configurations. For SPBM configurations, the switch supports a maximum of 16k MAC entries. You can enable this feature at port level.

* Note:

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

About this task

Limit MAC address learning to limit the number of forwarding database (FDB) entries learned on a particular port to a user-specified value. After the number of learned forwarding database entries reaches the maximum limit, MAC learning stops on that port.

Procedure

1. Enter Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]} or interface mlt <1-512>
```

* Note:

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

2. Protect the FDB from hits by too many MAC addresses:

```
mac-security port {slot/port [-slot/port] [,...]} limit-learning
enable [max-addr <1-32000>]
```

Example

Protect the FDB from hits by too many MAC addresses:

```
Switch(config)# interface gigabitethernet 1/1
Switch(config-if)# mac-security limit-learning enable
Switch(config-if)# mac-security limit-learning max-addr 5000
```

Variable definitions

Use the data in the following table to use the `mac-security limit-learning` command.

Table 8: Variable definitions

Variable	Value
enable	Limits the MAC learning for the port. This feature does not affect the forwarding of the packets. If you enable limit-learning, the FDB entry for each port is limited to the number you specify in max-addr. If you enable the auto-learn parameter, after the maximum addresses are learned, all the new SA MAC packets are dropped. This feature provides no value if you enable unknown-mac-discard and disable auto-learn because all unknown packets are dropped. Do not enable auto-learning and limit-learning simultaneously.
max-addr <1–32000>	Specifies the maximum number of MAC addresses to learn. After the maximum value is reached, no further MAC learning occurs. The system does not drop packets; it forwards packets. The default is 1024.
port {slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following format: a single slot and port (1/1).

Configuring the forwarding database timeout globally

Use the following procedure to configure the aging time globally for the forwarding database.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enter the following command:

```
mac-address-table aging-time <10-1000000>
```

Variable definitions

Use the data in the following table to use the `mac-address-table` command.

Table 9: Variable definitions

Variable	Value
<code>aging-time</code>	Specifies the timeout period for dynamically learned mac addresses on the vlan. The default value is 300.
<code><10-1000000></code>	Specifies the range for the aging time.

Adding or removing ports in a VLAN

Add or remove the ports in a VLAN to configure the ports in the VLAN.

Procedure

1. Enter Interface Configuration mode:

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

*** Note:**

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

2. Add ports in a VLAN:

```
vlan members add <1-4059> {slot/port[/sub-port]} [-slot/port[/sub-port]][, ...] [{portmember|static|notallowed}]
```

3. Remove ports in a VLAN:

```
vlan members remove <1-4059> {slot/port[/sub-port]} [-slot/port[/sub-port]][, ...] [{portmember|static|notallowed}]
```

Example

Add ports in a VLAN:

```
Switch(config-if)# vlan members add 1 1/2 static
```

Remove ports in a VLAN:

```
Switch(config-if)# vlan members remove 1 1/2 notallowed
```

Variable definitions

Use the data in the following table to use the `vlan members add` and `vlan members remove` commands.

Table 10: Variable definitions

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 <code>vrf-scaling</code> and <code>spbm-config-mode boot</code> 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.
{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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
portmember	Configures the port type as port member.

Adding or removing source MAC addresses for a VLAN

Add or remove a VLAN source MAC addresses to configure the source MAC address for a source MAC-based VLAN.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Add a VLAN source MAC address:

```
vlan srcmac <1-4059> <0x00:0x00:0x00:0x00:0x00:0x00>
```

3. Remove a VLAN source MAC address:

```
no vlan srcmac <1-4059> <0x00:0x00:0x00:0x00:0x00:0x00>
```

Example

Add a VLAN source MAC address:

```
Switch(config)# vlan create 10 type srcmac-mstprstp 0
Switch(config)# vlan srcmac 10 0x00:0x00:0x00:0x00:0x00:0x11
```

Configuring NLB support

Use Microsoft Network Load Balancing (NLB) to share the workload among multiple clustering servers. For information about software scaling capabilities, see *Release Notes*.

*** Note:**

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Before you begin

- For all modes, configure an IP address on the VLAN enabled with NLB.
- To switch between Unicast NLB and Multicast NLB, you must first disable the NLB support.

About this task

Use the following procedure to configure NLB support on an IP interface to enable or disable NLB support.

The default value is NLB support disabled.

*** Note:**

SPBM supports Network Load Balancing (NLB) Unicast and Multicast modes.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable NLB support on an interface:

```
nlb-mode unicast
```

Or,

```
nlb-mode multicast
```

To switch from one nlb-mode to another, you must first disable the NLB support, and then enter the new nlb-mode.

3. **(Optional)** Disable NLB support on an interface:

```
no nlb-mode
```

Example

Configure unicast mode for VLAN 2, and then switch to multicast mode.

```
Switch:1(config)#interface vlan 2
Switch:1(config-if)#nlb-mode unicast
Switch:1(config-if)#no nlb-mode
```

```
Switch:1(config-if)#nlb-mode multicast
```

Configuring a tagged port to discard untagged frames

Configure a tagged port to discard all untagged packets so that the frame is not classified into the default VLAN for the port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

* Note:

If your 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 tagged port to discard untagged frames:

```
untagged-frames-discard [port {slot/port[/sub-port] [-slot/port[/sub-
port]][,...]]
```

3. Discard a tagged frame on an untagged port:

```
tagged-frames-discard [port {slot/port[/sub-port] [-slot/port[/sub-
port]][,...]] enable
```

4. Untag the default VLAN on a tagged port:

```
untag-port-default-vlan [port {slot/port[/sub-port] [-slot/port[/sub-
port]][,...]] enable
```

Example

Configure a tagged port to discard untagged frames:

```
Switch(config-if)#untagged-frames-discard port 1/1
```

Discard a tagged frame on an untagged port:

```
Switch(config-if)#tagged-frames-discard port 1/1 enable
```

Untag the default VLAN on a tagged port:

```
Switch(config-if)#untag-port-default-vlan port 1/2 enable
```

Variable definitions

Use the data in the following table to use optional parameters with the `untagged-frames-discard` command.

Table 11: Variable definitions

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

Configuring SLPP

Enable the Simple Loop Prevention Protocol (SLPP) globally and for a VLAN to detect a loop and automatically stop it. The VLAN configuration controls the boundary of SLPP-PDU transmission.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Enable SLPP:


```
slpp enable
```
3. Configure the transmission interval:


```
slpp tx-interval <500-5000>
```
4. Add a VLAN to the transmission list:


```
slpp vid <1-4059>
```

Example

Enable SLPP:

```
Switch(config)# slpp enable
```

Configure the transmission interval to 5000 milliseconds:

```
Switch(config)# slpp tx-interval 5000
```

Add a VLAN, with the VLAN ID 2, to the transmission list:

```
Switch(config)# slpp vid 1
```

Variable definitions

Use the data in the following table to use the `slpp` command.

Table 12: Variable definitions

Variable	Value
enable	Enables or disables the SLPP operation. You must enable the SLPP operation to enable the SLPP packet transmit and receive process. If you disable the SLPP operation, the system sends no SLPP packets and discards received SLPP packets. To set this option to the default value, use the default operator with the command. The default is disabled.
500–5000	Configures the SLPP packet transmit interval, expressed in milliseconds in a range from 500–5000. The default value is 500. To set this option to the default value, use the default operator with the command.
<1-4059>	Adds a VLAN, by VLAN ID, to a SLPP transmission list. Use the no operator to remove this configuration.

Job aid

The following table provides the recommended SLPP values.

Table 13: SLPP recommended values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configuring SLPP packet-rx on a port

Enable SLPP by port to detect a loop and automatically stop it.

Important:

To provide protection against broadcast and multicast storms, it is recommended that you enable Rate Limiting for broadcast traffic and multicast traffic.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

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

 **Note:**

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

2. Configure SLPP on a port:

```
slpp port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}
packet-rx [packet-rx-threshold <1-500>]
```


Example

```
Switch(config-if)# slpp port 1/1 packet-rx-threshold 5
```

Variable definitions

Use the data in the following table to use the `slpp port` command.

Table 14: Variable definitions

Variable	Value
<1-500>	<p>Specifies the SLPP reception threshold on the ports, expressed as an integer. The packet reception threshold specifies how many SLPP packets the port receives before it is administratively disabled. To set this option to the default value, use the default operator with the command. The default value is 1.</p> <p> Important:</p> <p>It is recommended that you configure the rx-threshold above 50 slpp packets only on lightly loaded switches. If you configure the rx-threshold to a value greater than 50 on a heavily loaded switch and a loop occurs, the system can experience high CPU utilization.</p>
{slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]}	<p>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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.</p>

Job aid

The following table provides the recommended SLPP values.

Table 15: SLPP recommended values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configuring SLPP packet-tx on a VLAN

Enable SLPP by VLAN to detect a loop and automatically stop it. This configuration controls the boundary of SLPP-PDU transmission.

Important:

To provide protection against broadcast and multicast storms, it is recommended that you enable Rate Limiting for broadcast traffic and multicast traffic.

Procedure

1. Enter Interface Configuration mode:

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

Note:

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

2. Enable SLPP:

```
slpp enable
```

3. Configure the transmission interval:

```
slpp tx-interval <500-5000>
```

4. Add a VLAN to the transmission list:

```
slpp vid <1-4059>
```

Example

Log on to the VLAN Interface Configuration mode:

```
Switch(config)# interface vlan 2
```

Enable SLPP:

```
Switch(config-if)# slpp enable
```

Configure the transmission interval to 500 milliseconds:

```
Switch(config-if)# slpp tx-interval 500
```

Add a VLAN, with the VLAN ID of 2, to the transmission list:

```
Switch(config-if)# slpp vid 2
```

Variable definitions

Use the data in the following table to use the `slpp` command.

Table 16: Variable definitions

Variable	Value
enable	<p>Activates or disables the SLPP operation.</p> <p>You must enable the SLPP operation to enable the SLPP packet transmit and receive process.</p> <p>If you disable the SLPP operation, the system sends no SLPP packets and discards received SLPP packets.</p> <p>To set this option to the default value, use the default operator with the command. The default is disabled.</p>
500–5000	<p>Configures the SLPP packet transmit interval, expressed in milliseconds in a range from 500–5000. The default value is 500. To set this option to the default value, use the default operator with the command.</p>
<1-4059>	<p>Adds a VLAN, by VLAN ID, to a SLPP transmission list. Use the no operator to remove this configuration.</p>

Viewing SLPP information

Use SLPP information to view loop information.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View SLPP information:

```
show slpp
```

Example

```
Switch# show slpp
```

```
=====
                                SLPP Info
=====
operation : enabled
tx-interval : 500
vlan : 2
```

Viewing SLPP information for a port

Show SLPP information for a port so that you can view the loop information for a port.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View SLPP information for a port:

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

3. Clear SLPP packet RX counters:

```
clear slpp stats port [{slot/port[/sub-port]}[-slot/port[/sub-port]]
[,...]]
```

Example

```
Switch# show slpp interface GigabitEthernet 1/7
```

```
=====
                                Port Interface
=====
PORT      PKT-RX      PKT-RX      INCOMING      SLPP PDU
NUM       COUNT      THRESHOLD   VLAN ID       ORIGINATOR
-----
1/7       enabled     5
PORT      PKT-RX      TIME LEFT
NUM       COUNT      TO CLEAR RX COUNT
-----
1/7       29         21600
```

Variable definitions

Use the data in the following table to use the `show slpp interface GigabitEthernet` command.

Table 17: Variable definitions

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

Configuring spoof detection

* Note:

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Configure spoof detection to prevent IP spoofing.

For more information about this feature, see [Prevention of IP spoofing within a VLAN](#) on page 28.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

* Note:

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

2. Enable or disable spoof detection:

```
spoof-detect [port {slot/port[-slot/port][,...]}] [enable]
no spoof-detect [port {slot/port[-slot/port][,...]}] [enable]
```

3. Enable or disable auto-recovery on a port:

```
auto-recover-port [port {slot/port[-slot/port][,...]}] [enable]
no auto-recover-port [port {slot/port[-slot/port][,...]}] [enable]
```

Example

Enable spoof detection:

VLAN configuration

```
Switch(config-if)# spoof-detect port 1/1 enable
```

Enable autorecovery on a port:

```
Switch(config-if)# auto-recover-port port 1/1 enable
```

Viewing VLAN information

View the VLAN information to display the basic configuration for all VLANs or a specified VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View VLAN information:

```
show vlan basic <1-4059>
```

3. View advanced parameters:

```
show vlan advance <1-4059>
```

Example

View VLAN information for VLAN 2:

```
Switch:1> show vlan basic 2
```

```
=====
                                Vlan Basic
=====
```

VLAN ID	NAME	TYPE	INST ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRFLD
2	VLAN-2	byPort	0	none	N/A	N/A	0

```
=====
```

View VLAN information:

```
Switch:1> show vlan basic
```

```
=====
                                Vlan Basic
=====
```

VLAN ID	NAME	TYPE	INST ID	PROTOCOLID	SUBNETADDR	SUBNETMASK	VRFLD
1	Default	byPort	0	none	N/A	N/A	0
2	abc	byPort	0	none	N/A	N/A	0
3	VLAN-VRRP	byPort	0	none	N/A	N/A	0
4	VLAN-6	byPort	0	none	N/A	N/A	1
5	VLAN-7	byPort	0	none	N/A	N/A	1
6	VLAN-8	byPort	0	none	N/A	N/A	1
19	VLAN-9	byPort	0	none	N/A	N/A	0

```
=====
```



```

10  VLAN-10      byPort      0  none      N/A        N/A        0
11  VLAN-11      byPort      0  none      N/A        N/A        0
12  VLAN-12      byPort      0  none      N/A        N/A        0
13  VLAN-13      spbm-bvlan  62 none      N/A        N/A        0
14  VLAN-14      spbm-bvlan  62 none      N/A        N/A        0
15  VLAN-15      byPort      1  none      N/A        N/A        0
--More-- (q = quit)

```

View advanced parameters:

```
Switch:1> show vlan advance
```

```

=====
                        Vlan Advance
=====
VLAN      IF      AGING  MAC      USER
ID        NAME    INDEX  TIME     ADDRESS  DEFINEPID ENCAP  DSAP/SSAP
-----
2         Default 2050   0        00:24:7f:9f:6a:03  0x0000

```

Variable definitions

Use the data in the following table to use optional parameters with the `show vlan basic` and `show vlan advance` commands.

Table 18: Variable definitions

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 <code>vrf-scaling</code> and <code>spbm-config-mode boot</code> 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.

Viewing private VLAN information

You can view the private VLAN information to display the primary and secondary VLANs and I-SIDs, and also view the private VLAN port types.

Procedure

1. Enter Privileged EXEC mode:
`enable`
2. View private VLAN information:

VLAN configuration

```
show vlan private-vlan <1-4059>
```

3. View private vlan port information:

```
show interfaces gigabitethernet private-vlan
```

Example

View VLAN information for private VLAN :

```
Switch:1(config)# show vlan private-vlan
```

```
=====
                        PRIVATE VLAN
=====
```

Primary VLAN	Primary ISID	Secondary VLAN	Secondary ISID
3	75	5	75
10	22	15	22

```
-----
All 2 out of 2 Total Num of Private Vlans displayed
```

View port information for private VLAN:

```
Switch:1(config)# show interfaces gigabitethernet private-vlan
```

```
=====
                        Port Private Vlans
=====
```

PORT NUM	TAGGING	PVLAN	PVLAN TYPE	VID TYPE	VID
1/1	enable	enable	isolated	secondary	5
1/2	enable	enable	promiscuous	primary	3
1/3	enable	enable	trunk	both	3/5

Viewing brouter port information

View the brouter port information to display the brouter port VLAN information for all VLANs on the device or for the specified VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View brouter port information:

```
show vlan brouter-port
```

Example

View brouter port information:

```
Switch:1> show vlan brouter-port
```

```
=====
```

Vlan Id	Port	VrfId
2202	1/11	0

```
-----
All 1 out of 1 Total Num of Vlan Brouter Port Entries displayed
```

Viewing VLAN port member status

View the VLAN port member status to display the port member status for all VLANs on the device or for the specified VLAN.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View VLAN port member status:

```
show vlan members [<1-4059>][null-vlan][port {slot/port[/sub-port]}[-
slot/port[/sub-port]][, ...]}
```

Example

View VLAN port member status:

```
Switch:1> show vlan members port 1/2
```

Vlan Port				
VLAN ID	PORT MEMBER	ACTIVE MEMBER	STATIC MEMBER	NOT_ALLOW MEMBER
2	1/2,1/5-1/8,1/11, 1/14,1/26,1/38	1/2,1/5-1/8,1/11, 1/14,1/26,1/38		
3	1/2,1/5-1/8,1/14, 1/26,1/38	1/2,1/5-1/8,1/14, 1/26,1/38		
4	1/1-1/2,1/5-1/8, 1/13-1/14,1/25- 1/26,1/37-1/38	1/1-1/2,1/5-1/8, 1/13-1/14,1/25- 1/26,1/37-1/38		
100	1/2,1/14,1/23- 1/24,1/26-1/28, 1/38	1/2,1/14,1/23- 1/24,1/26-1/28, 1/38		
300	1/2,1/5-1/8,1/14, 1/26,1/38	1/2,1/5-1/8,1/14, 1/26,1/38		

Variable definitions

Use the data in the following table to use optional parameters with the `show vlan members` command.

Table 19: Variable definitions

Variable	Value
null-vlan	Displays port members of the NULL VLAN. This is a place holder VLAN for ports that are not members of any port-based VLAN. When a port is removed from all port-based VLANs, it is added to the NULL VLAN as a port member. The NULL VLAN is an internal construct and cannot be deleted.
{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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. If you do not specify a port, the command shows information for all the ports.
<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. ! Important: Entering a VLAN ID is optional. If you enter a VLAN ID the command shows information for the specified VLAN or port. Without the VLAN ID the command shows information for all the configured VLANs.

Viewing VLAN source MAC addresses

View the VLAN source MAC addresses to display the source MAC address for a source MAC-based VLAN on the device or for the specified VLAN.

Procedure

View VLAN source MAC addresses:

```
show vlan src-mac [<1-4059>]
```

Example

View VLAN source MAC addresses:

```
Switch(config)# show vlan src-mac
```

```
=====
                        Vlan Srcmac
=====
VLAN_ID   MAC_ADDRESS
-----
10        00:00:00:00:00:11
All 1 out of 1 Total Num of Vlan Srcmac Entries displayed
```

Viewing VLAN forwarding database information

Use this procedure to display the MAC addresses that are learned or statically configured for a VLAN. In order to learn you have to be connected to another switch or host and receive some traffic.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View VLAN forwarding database information:

```
show vlan mac-address-entry [<1-4059>]
```

Example

View VLAN forwarding database information:

```
Switch:1> show vlan mac-address-entry
```

```
=====
                        Vlan Fdb
=====
VLAN      MAC              SMLT
ID  STATUS  ADDRESS          INTERFACE  REMOTE  TUNNEL
-----
1    learned  f8:15:47:e1:80:0c  Port-1/2  false  -
2    learned  32:20:d3:81:00:77  Port-1/9  false  -
4    learned  b4:a9:5a:2b:78:31  Port-2/1  false  -
3 out of 3 entries in all fdb(s) displayed.
```

View where entries are learned. The TUNNEL column indicates where in the SPBM network an entry is learned.

```
Switch:1> show vlan mac-address-entry spbm-tunnel-as-mac
```

```
=====
                        Vlan Fdb
=====
VLAN      MAC              SMLT
ID  STATUS  ADDRESS          INTERFACE  REMOTE  TUNNEL
-----
1    learned  f8:15:47:e1:80:0c  Port-1/2  false  -
2    learned  32:20:d3:81:00:77  Port-1/9  false  -
4    learned  b4:a9:5a:2b:78:31  Port-2/1  false  -
3 out of 3 entries in all fdb(s) displayed.
```

Variable definitions

Use the data in the following table to use optional parameters with the `show vlan mac-address-entry` command.

Table 20: Variable definitions

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.
mac <0x00:0x00:0x00:0x00:0x00:0x00>	Specifies the MAC address.
port {slot/port[-slot/port][,...]}	Specifies the port or port list.
spbm-tunnel-as-mac	Displays where entries are learned. The TUNNEL column indicates where in the SPBM network an entry is learned.

Viewing manual edit MAC addresses

Use the procedure to view the list of manual edit MAC addresses and the associated ports configured as allow-mac for MAC security.

Procedure

View manual edit MAC addresses:

```
show vlan manual-edit-mac
```

Example

View manual edit MAC addresses:

```
Switch(config)# show vlan manual-edit-mac
```

```
=====
Manual Edit Mac
=====
MAC ADDRESS          PORTS
-----
00:00:00:00:00:55   1/3
00:00:00:00:00:66   1/3
All 2 out of 2 Total Num of Manual Edit Mac Entries displayed
```

Viewing port-level MAC security

* Note:

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

View port-level MAC security to review the configuration.

Before you begin Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View port-level MAC security for limit-learning:

```
show interface gigabitethernet limit-fdb-learning [{slot/port[-slot/
port][,...]]]
```

Example

View port-level MAC security for limit-learning:

```
Switch(config)# show interface gigabitethernet limit-fdb-learning 1/4-1/5
```

```
=====
                        Port limit-fdb-learning
=====
```

PORT NUM	FDB PROTECT	MAXMAC COUNT	MINMAC COUNT	LOG TRAP	PORT DOWN	CURMAC COUNT	MAC LEARN
1/4	dis	1024	512	dis	dis	0	true
1/5	ena	5000	3000	dis	dis	0	true

```
=====
```

Viewing NLB-mode information

View Network Load Balancing-mode (NLB-mode) information.

* Note:

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View NLB port information:

```
show interface vlan nlb-mode [<1-4059>]
```

Example

View NLB-mode information.

```
Switch:1#show interface vlan nlb-mode
=====
                        Vlan Nlb
=====
VLAN_ID  NLB_ADMIN_MODE  NLB_OPER_MODE  PORT_LIST  MLT_GROUPS
-----
2         unicast          disable
22        multicast       multicast      1/19-1/21  2 3
Total Entries: 2
```

Displaying C-VLAN and Switched UNI I-SID information

Use the following procedure to display C-VLAN I-SID information.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Display the C-VLAN to I-SID associations:


```
show vlan i-sid <1-4059>
```
3. Display I-SID information and Switched UNI to I-SID associations:


```
show i-sid <1-16777215>
```
4. Display the IS-IS SPBM multicast-FIB calculation results by I-SID:


```
show isis spbm i-sid {all|config|discover} [vlan <1-4059>] [id <1-16777215>] [nick-name <x.xx.xx>]
```
5. Display all elan I-SID:
 - show i-sid elan
6. Display I-SID configured on MLT:
 - show mlt i-sid
7. Display I-SID configured on port:
 - show interfaces gigabitethernet i-sid

Example

```
Switch:1#show vlan i-sid
=====
                        Vlan I-SID
=====
VLAN_ID  I-SID
-----
1
2
5         5
```


10
20

Switch:1#show isis spbm i-sid all

```
=====
                        SPBM ISID INFO
=====
ISID   SOURCE NAME   VLAN   SYSID           TYPE           HOST_NAME
-----
200    1.11.16       1000   0014.c7e1.33df  config         Switch1
300    1.11.16       1000   0014.c7e1.33df  config         Switch1
400    1.11.16       1000   0014.c7e1.33df  config         Switch1
200    1.11.16       2000   0014.c7e1.33df  config         Switch1
300    1.11.16       2000   0014.c7e1.33df  config         Switch1
400    1.11.16       2000   0014.c7e1.33df  config         Switch1
200    1.12.45       1000   0016.ca23.73df  discover       Switch2
300    1.12.45       1000   0016.ca23.73df  discover       Switch2
=====

Total number of SPBM ISID entries configed: 6
-----
Total number of SPBM ISID entries discovered: 2
-----
Total number of SPBM ISID entries: 8
=====
```

switch:1#show i-sid

```
=====
                        Isid Info
=====
ISID   ISID          VLANID  PORT            MLT             ORIGIN
ID     TYPE          ID      INTERFACES      INTERFACES
-----
999    ELAN          99      -               c110:100       CONFIG
      99           99      1/21            -
c: customer vid    u: untagged-traffic
All 1 out of 1 Total Num of i-sids displayed
```

switch:1#show mlt i-sid

```
=====
                        MLT Isid Info
=====
MLTID  IFINDEX  ISID          VLANID  C-VID  ISID          ORIGIN  BPDU
      ID      ID           ID      ID     TYPE         TYPE
-----
10     6153    100          N/A     20     ELAN         CONFIG
-----

1 out of 1 Total Num of i-sid endpoints displayed
```

switch:1#show interfaces gigabitEthernet i-sid

```
=====
                        PORT Isid Info
=====
PORTNUM IFINDEX  ISID          VLANID  C-VID  ISID          ORIGIN  BPDU
      ID     ID      ID           ID      ID     TYPE         TYPE
-----
1/1     192     100          N/A     10     ELAN         CONFIG
1/2     193     100          N/A     10     ELAN         CONFIG
-----

2 out of 3 Total Num of i-sid endpoints displayed
```

Variable definitions

Use the data in the following table to use the **show vlan i-sid** commands.

Variable	Value
<1-4059>	Displays I-SID information for the specified C-VLAN. You can specify the VLAN ID.

Use the data in the following table to use the **show i-sid** commands

Variable	Value
<1-16777215>	Displays I-SID information. You can specify the I-SID ID.

Use the data in the following table to use the **show isis** commands.

Variable	Value
spbm i-sid {all config discover}	<ul style="list-style-type: none"> • all: displays all I-SID entries • config: displays configured I-SID entries • discover: displays discovered I-SID entries

Job aid

The following sections describe the fields in the outputs for the C-VLAN I-SID show commands.

show vlan i-sid

The following table describes the fields in the output for the **show vlan i-sid** command.

Parameter	Description
VLAN_ID	Indicates the VLAN IDs.
I-SID	Indicates the I-SIDs associated with the specified C-VLANs.

show i-sid

The following table describes the fields in the output for the **show i-sid** command.


Parameter	Description
I-SID	Indicates the I-SID IDs.
I-SID TYPE	Indicated the I-SID type. <ul style="list-style-type: none"> • T-UNI: Transparent UNI service. • ELAN: any to any service (switched service). • CVLAN: CVLAN based service.
VLANID	Indicates the VLAN IDs.
PORT INTERFACES	Indicated the port interface.
MLT INTERFACES	Indicates the MLT interface.

Table continues...

Parameter	Description
ORIGIN	Indicates if the I-SID is discovered by Fabric Attach or manually added.

show isis spbm i-sid

The following describes the fields in the output for the `show isis spbm i-sid` command.

Parameter	Description
ISID {all discover config}	Indicates the IS-IS SPBM I-SID identifier. <ul style="list-style-type: none"> all: display all SPBM I-SID discover: display discovered SPBM I-SID config: display configured SPBM I-SID
SOURCE NAME	Indicates the nickname of the node where this I-SID was configured or discovered. <p> Note: SOURCE NAME is equivalent to nickname.</p>
VLAN	Indicates the B-VLAN where this I-SID was configured or discovered.
SYSID	Indicates the system identifier.
TYPE	Indicates the SPBM I-SID type as either configured or discovered.
HOST_NAME	Indicates the host name of the multicast FIB entry.

Enabling DvR on a Layer 2 VSN (VLAN)

Before you begin

- Ensure that the VLAN on which to enable DvR exists, and is associated with an I-SID.

About this task

On a Controller, DvR must be manually enabled at the Layer 2 VSN (VLAN) level.

Use this procedure to configure a gateway IPv4 address for a Layer 2 VSN (VLAN) subnet, and then enable DvR on it. This address is pushed, along with other L3 configuration information, from the Controllers to the Leaf nodes within the domain, so that the Leaf nodes can create a gateway IP service for each DvR enabled Layer 2 VSN.

Procedure

- Enter VLAN Interface Configuration mode:

```
enable
configure terminal
```

VLAN configuration

```
interface vlan <1-4059>
```

2. Configure a gateway IPv4 address for the VLAN.

```
dvr gw-ipv4 {A.B.C.D}
```

! **Important:**

Ensure that you configure the same gateway IPv4 address on all Controllers in the domain that belong to a Layer 2 VSN (VLAN).

3. Enable DvR.

```
dvr enable
```

By default, DvR is disabled.

Example

Enable DvR on the Global Router VLAN.

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config-if)#vlan create 200 type port-mstprstp 0
Switch:1(config-if)#vlan i-sid 200 20200
Switch:1(config-if)#interface vlan 200
Switch:1(config-if)#dvr gw-ipv4 192.0.2.1
Switch:1(config-if)#dvr enable
Switch:1(config-if)#ip address 192.0.2.2 255.255.0.0
```

Enable DvR on a VLAN associated to a VRF:

```
Switch:1>en
Switch:1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config-if)#vlan create 400 type port-mstprstp 0
Switch:1(config-if)#vlan i-sid 400 20200
Switch:1(config-if)#interface vlan 400
Switch:1(config-if)#vrf vrf500
Switch:1(config-if)#dvr gw-ipv4 198.51.100.1
Switch:1(config-if)#dvr enable
Switch:1(config-if)#ip address 198.51.100.2 255.255.0.0
```

Variable definitions

Use the data in the following table to use the `dvr gw-ipv4` command.

Variable	Value
{A.B.C.D}	Specifies the gateway IPv4 address for the VLAN.

VLAN configuration using EDM

This chapter describes how to configure and manage Virtual Local Area Networks (VLAN) using Enterprise Device Manager (EDM).

Configuring the VLAN feature on a port

Configure the VLAN feature on a port.

Procedure

1. In the Device Physical View tab, select a port or multiple ports.
2. In the Navigation tree, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **VLAN** tab.
5. To perform tagging, select **PerformTagging**.
6. To discard tagged frames, select **DiscardTaggedFrames**.
7. To discard untagged frames, select **DiscardUntaggedFrames**.
8. To use the Untag Default VLAN feature, select **UntagDefaultVlan**.

Important:

It is recommended that you enable tagging on the port before you configure UntagDefaultVlans.

9. Enter a default VLAN ID.
10. To enable spoof detect, select **SpoofDetect**.
11. In the Classification area, select the types of VLAN to enable.
12. In the **Classification** area, select the Private VLAN port type. See [Creating a private VLAN](#) on page 73 for more information.
13. Click **Apply**.
14. Click **Close**.

VLAN field descriptions

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

Name	Description
PerformTagging	If checked, this port is a tagged (Trunk) Port. It can belong to multiple port-based VLANs and a VLAN tag is inserted in every frame it transmits. If it is not checked, the port is an untagged (Access) port. The default is disabled.
VlanIdList	Identifies which VLANs this port is assigned.
DiscardTaggedFrames	If selected, and the port is untagged (an access port), tagged frames received on the port are

Table continues...

Name	Description
	discarded by the forwarding process. If clear, tagged frames are processed normally. The default is disabled.
DiscardUntaggedFrames	If selected and the port is tagged (a trunk port), untagged frames received on the port are discarded by the forwarding process. If clear, untagged frames are processed normally. The default is disabled.
UntagDefaultVLAN	If selected, even if the port is tagged (a trunk port), frames forwarded to the default VLAN for the port are not tagged. The default is disabled.
UntaggedVlanIds	Identifies which VLANs this port is associated with as untagged.
DefaultVlanId	Specifies the VLAN ID assigned to untagged frames received on this trunk port that match no policy-based VLAN to which the port belongs.
SpoofDetect	Enables or disables spoof detection on the specified port.
Protocol	Enables protocol-based VLAN on the port. This feature is always enabled.
PrivateVlanPortType	Specifies the port type. If not specified, the port type defaults to None. <ul style="list-style-type: none"> • Isolated: Only private VLANs are permitted on isolated ports. • Promiscuous: Only private VLANs are permitted on promiscuous ports. • Trunk: The port is tagged.

Viewing existing VLANs

Display existing VLANs to view all defined VLANs, their configurations, and the current status.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. View the configured VLANs in the **Basic** tab.
4. View the configured private VLANs in the **Private VLAN** tab.

Creating a port-based VLAN

Create a port-based VLAN to add a new VLAN. To create a different type of VLAN, see one of the following procedures:

- [Creating a protocol-based VLAN](#) on page 77
- [Creating an SPBM B-VLAN](#) on page 79

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. In the **Basic** tab, click **Insert**.
4. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
5. In the **Name** box, type the VLAN name, or use the name provided.
6. In the **Color Identifier** box, click the down arrow and choose a color from the list, or use the color provided.
7. In the **MstpInstance** box, click the down arrow and choose an msti instance from the list.
8. In the **Type** box, select **byPort**.
9. In the **PortMembers** box, click the (...) button.
10. Click on the ports to add as member ports.
The ports that are selected are recessed, while the nonselected ports are not recessed. Port numbers that appear dimmed cannot be selected as VLAN port members.
11. Click **OK**.
12. Click **Insert**.

Basic field descriptions

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

Name	Description
Id	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	Specifies the name of the VLAN.

Table continues...

Name	Description
Ifindex	Specifies the logical interface index assigned to the VLAN.
Color Identifier	Specifies a proprietary color scheme to associate a color with the VLAN. Color does not affect how frames are forwarded.
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-bvlan • private
MstpInstance	Identifies the MSTP instance.
VrfId	Indicates the Virtual Router to which the VLAN belongs.
VrfName	Indicates the name of the Virtual Router to which the VLAN belongs.
PortMembers	Specifies the slot/port of each VLAN member. The sub-port only appears for channelized ports.
ActiveMembers	Specifies the slot/port of each VLAN member. The sub-port only appears for channelized ports.
StaticMembers	Specifies the slot/port of each static member of a policy-based VLAN. The sub-port only appears for channelized ports.
NotAllowToJoin	Specifies the slot/ports that are never allowed to become a member of the policy-based VLAN. The sub-port only appears for channelized ports.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.

 **Note:**

If you or another user changes the name of an existing VLAN using the VLAN **Basic** tab (or using CLI), the new name does not initially appear in EDM. To display the updated name, perform one of the following actions:

- Refresh your browser to reload EDM.
- Log out of EDM and log in again to restart EDM.
- Click **Refresh** in the VLAN **Basic** tab toolbar. If the old VLAN name appears in other tabs, click **Refresh** on those tabs as well.

Creating a private VLAN

Before you begin

- To create a private VLAN, you must set the VLAN type to private and set the private VLAN port type.
- The ports you add to a private VLAN must have a port type of isolated, promiscuous, or trunk.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. In the **Basic** tab, click **Insert**.
4. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
5. In the **Name** box, type the VLAN name, or use the name provided.
6. In the **Color Identifier** box, click the down arrow and choose a color from the list, or use the color provided.
7. In the **MstpInstance** box, click the down arrow and choose an msti instance from the list.
8. In the **Type** box, select **private**.
9. In the **PortMembers** box, click the (...) button.
10. Click on the ports to add as member ports.

The ports that are selected are recessed, while the non-selected ports are not recessed. Port numbers that appear dimmed cannot be selected as VLAN port members.

11. Click **OK**.
12. In the **Secondary Vlan** box, enter an unused VLAN ID.
13. Click **Insert**.
14. Collapse the **VLANs** tab.

The VLAN is added to the **Basic** tab.

To set the port type for the private VLAN:

15. In the navigation pane, expand the following folders: **Configuration > VLAN**.
16. Click **VLANs**.
17. In the Classification area, select the **PrivateVlanPortType**.
18. Click **Apply**.
19. Click **Close**.

Basic field descriptions

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

Name	Description
Id	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	Specifies the name of the VLAN.
IfIndex	Specifies the logical interface index assigned to the VLAN.
Color Identifier	Specifies a proprietary color scheme to associate a color with the VLAN. Color does not affect how frames are forwarded.
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-bvlan • private
MstpInstance	Identifies the MSTP instance.
VrfId	Indicates the Virtual Router to which the VLAN belongs.
VrfName	Indicates the name of the Virtual Router to which the VLAN belongs.
PortMembers	Specifies the slot/port of each VLAN member. The sub-port only appears for channelized ports.
ActiveMembers	Specifies the slot/port of each VLAN member. The sub-port only appears for channelized ports.
StaticMembers	Specifies the slot/port of each static member of a policy-based VLAN. The sub-port only appears for channelized ports.
NotAllowToJoin	Specifies the slot/ports that are never allowed to become a member of the policy-based VLAN. The sub-port only appears for channelized ports.
ProtocolId	Specifies the network protocol for protocol-based VLANs. This value is taken from the Assigned Numbers of remote function call (RFC). If the VLAN type is port-based, none is displayed in the Basic tab ProtocolId field.

*** Note:**

If you or another user changes the name of an existing VLAN using the VLAN **Basic** tab (or using CLI), the new name does not initially appear in EDM. To display the updated name, perform one of the following actions:

- Refresh your browser to reload EDM.
- Log out of EDM and log in again to restart EDM.
- Click **Refresh** in the VLAN **Basic** tab toolbar. If the old VLAN name appears in other tabs, click **Refresh** on those tabs as well.

VLAN field descriptions

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

Name	Description
PrivateVlanPortType	<p>Specifies the port type. If not specified, the port type defaults to None.</p> <ul style="list-style-type: none"> • Isolated: An Isolated port can belong only to one private VLAN • Promiscuous: A Promiscuous port can belong to many private VLANs. • Trunk: A Trunk port can belong to many private VLANs, is tagged, and can also belong to non-private VLANs.

Viewing Private VLAN information

You can view the private VLAN information to display the primary and secondary VLANs and I-SIDs, and also view the private VLAN port types.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **VLANs**.
3. Click **Private VLAN**.

Private VLAN field descriptions

Use the data in the following table to use the Private VLAN tab.

Name	Description
Primary Vlan	Shows the VLAN ID for the primary VLAN.
Secondary Vlan	Shows the VLAN ID for the secondary VLAN.
Primary / Secondary I-sid	Shows the I-SID for the VLAN.

Configuring an IP address for a VLAN

Assign an IP address to a VLAN to enable routing on the VLAN.

About this task

If you configure the SLA Mon agent address under an IP interface or VLAN, you must remove the SLA Mon address before you can remove the IP address or VLAN.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **VLANs**.
3. In the **Basic** tab, select the VLAN for which you are configuring an IP address.
4. Click **IP**.
5. Click **Insert**.
6. Configure the required parameters.
7. Click **Insert**.

IP Address field descriptions

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

Name	Description
Interface	Shows the interface to which this entry applies.
Ip Address	Specifies the IP address to associate with the VLAN.
Net Mask	Specifies the subnet mask associated with the IP address of this entry. The value of the mask is an IP address with all the network bits configured to 1 and all the hosts bits configured to 0.
BcastAddrFormat	Shows the IP broadcast address format on this interface.
ReasmMaxSize	Shows the size of the largest IP datagram which this entity can reassemble from incoming IP fragmented datagrams received on this interface.
VlanId	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.

Table continues...

Name	Description
BrouterPort	Indicates whether this entry corresponds to a brouter port, as oppose to a routable VLAN.
MacOffset	Specifies the MAC offset value. Routable VLANS are assigned MAC addresses arbitrarily or by offset. Their MAC addresses are: <ul style="list-style-type: none"> • 24 bits: Vendor ID • 12 bits: Chassis ID • 12 bits: 0xA00-0xFFFF If you enter the MAC offset, the lowest 12 bits are 0xA00 plus the offset. If not, they are arbitrary.
Vrflid	Associates the VLAN or brouter port with a VRF. VRF ID 0 is reserved for the administrative VRF.

Changing VLAN port membership

Modify VLAN port members to control access to the VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANS**.
3. Double-click the **PortMembers** number for the VLAN for which you want to modify port membership.
4. Click the port members you wish to add or remove.
5. Click **Ok**.
6. Click **Apply**.

The VLAN port membership is changed.

Creating a protocol-based VLAN

Use a protocol-based VLAN so that the VLAN only carries certain traffic types.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANS**.
3. In the **Basic** tab, click **Insert**.
4. In the **Id** box, type the unique VLAN ID or use the ID provided.

5. In the **Name** box, type the VLAN name or use the name provided.
6. In the **Color Identifier** box, select the color or use the color provided.
This color is used to visually distinguish the VLANs in a network.
7. In the **MstpInstance** box, click the down arrow and choose an MSTI instance from the list.
8. In the **Type** box, select **byProtocolId**.
This activates additional fields needed to configure protocol-based VLANs.
9. To specify the VLAN port membership, click the button (...) for one of the following fields:

Port Members

OR

StaticMembers

OR

NotAllowToJoin

10. Click each port button to choose the desired membership color.
Yellow: Potential members—dynamic (potential members are treated as always members)
OR
Green: Always members—static
OR
Red: Never members—not allowed to join

! Important:

In a protocol-based VLAN, a potential member becomes an active member of the VLAN after a frame of the specified protocol is received.

11. Click **Insert**.

Configuring source MAC addresses for a source MAC-based VLAN

*** Note:**

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Create a source MAC address for an existing source MAC VLAN.

Before you begin

- Configure the VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. In the **Basic** tab, select a source MAC address-based VLAN.
4. Click **Mac**.
5. To manually insert a MAC address, click **Insert**, and then enter it in the form nn:nn:nn:nn:nn:nn.

OR

6. To add a MAC address from a file, select **File, Add From File**.
7. Use the selection box to browse for the file location.
8. To save a MAC address to a file, select it, select **File, Save to File**, and then use the selection box to browse for a save location.
9. To delete a MAC address, select it, and then select **Delete Members On Device**.
10. Click **Yes**.
11. Click **Close**.

The Edit MAC box closes.

VLAN MAC field descriptions

Use the data in the following table to use the **VLAN MAC** tab.

Name	Description
MacAddr	Specifies the MAC addresses associated with this VLAN.

Creating a SPBM B-VLAN

Create a Shortest Path Bridging MAC (SPBM) Backbone VLAN (B-VLAN). Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network. This VLAN is used for both control plane traffic and dataplane traffic.

* Note:

It is recommended that you always configure two B-VLANs in an SPBM dual-homing environment.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. In the **Basic** tab, click **Insert**.

4. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
5. In the **Name** box, type the VLAN name, or use the name provided.
6. In the **Color Identifier** box, click the down arrow and choose a color from the list, or use the color provided.
7. In the **Type** box, select **spbm-bvlan**.
8. Click **Insert**.
9. Collapse the **VLANS** tab.
The VLAN is added to the **Basic** tab.

Configuring advanced VLAN features

Use advanced VLAN features to configure the VLAN name, aging time, VLAN operation action, and QoS level. The VLAN Operation Action parameter can be useful for troubleshooting.

You can also configure a DvR Gateway IPv4 address on a VLAN, and enable DvR on it.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > VLAN**.
2. Click **VLANS**.
3. Click the **Advanced** tab.
4. Configure the parameters as required by double-clicking fields to make changes.
You cannot make changes to fields that appear dim.
5. Click **Apply**.

Advanced field descriptions



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

Name	Description
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 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.
Name	Specifies the name of the VLAN.
Ifindex	Specifies the logical interface index assigned to the VLAN.

Table continues...

Name	Description
Type	Specifies the type of VLAN: <ul style="list-style-type: none"> • byPort • byProtocolId • spbm-bvlan • private
I-sid	Specifies the I-SID number assigned to a customer VLAN (C-VLAN). The range is 0 – 16777215. The default value is 0, which indicates that no I-SID is assigned.
ProtocolId	Specifies the network protocol for protocol-based VLANs. If the VLAN type is not protocol-based, None is displayed in the Basic tab ProtocolId field.
AgingTime	Specifies the timeout period for dynamic VLAN membership. A potential VLAN port is made ACTIVE after it receives a packet that matches the VLAN; if no such packet is received for AgingTime seconds, the port is no longer active. The default is 600.
MacAddress	Specifies the MAC address assigned to the virtual router interface for this VLAN. This field is relevant only after the VLAN is configured for routing. This MAC address is used as the Source MAC in routed frames and ARP replies.
Vlan Operation Action	Performs an operation on the VLAN. The values are: <ul style="list-style-type: none"> • none • flushMacFdb: Configures action to flushMacFdb. This action removes the learned MAC addresses from the forwarding database for the selected VLAN. • flushArp: Configures action to flushArp. This action removes the ARP entries from the address table for the selected VLAN. • flushIp: Configures action to flushIp. This action removes the learned IP addresses from the forwarding table for the selected VLAN. • flushDynMemb: Configures action to flushDynMemb. This action removes port members not configured as static from the list of active port members of a policy-based VLAN and removes MAC addresses learned on those ports.

Table continues...

Name	Description
	<ul style="list-style-type: none"> all: Configures action to all. This action performs all the supported actions; it does not perform the Snoop-related actions. <p>The default is none.</p>
Result	Specifies the result code after you perform an action.
NlbMode	Enables or disables Microsoft Network Load Balancing (NLB) operations on the VLAN. The default is disabled.
SpbMulticast	Enables or disables Multicast over Fabric Connect. The default is disabled.
SpbPimGatewayMulticast	Enables or disables SPB-PIM Gateway Multicast on a VLAN. The default is disabled.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
Ipv6FhsSnoopDhcpEnable	Enables or disables IPv6 dhcp snooping on a VLAN. The default is disabled.
Ipv6FhsNDInspectionEnable	Enables or disables neighbor discovery (ND) inspection on a VLAN. The default is disabled.
DvrEnable	<p>Enables or disables DvR on a VLAN that is configured on the DvR Controller. The default is disabled.</p> <p> Note:</p> <p>You must enable DvR on every VLAN that is configured on a DvR Controller.</p>
DvrGwIpv4Addr	<p>Specifies the DvR gateway IPv4 address for a VLAN.</p> <p> Important:</p> <p>Ensure that you configure the same gateway IPv4 address on all Controllers in the DvR domain that belong to a VLAN.</p>

Configuring NLB support

Use Microsoft Network Load Balancing (NLB) to share the workload among multiple clustering servers. For more information about software scaling capabilities, see *Release Notes*.

Before you begin

Ensure that the VLAN exists and has an associated IP address.

About this task

Use the following procedure to configure NLB support on an IP interface to enable or disable NLB support. The default value is NLB support disabled.

 **Note:**

- SPBM supports Network Load Balancing (NLB) Unicast and Multicast modes.
- Static multicast ARP entries are not supported for NLB Unicast or NLB Multicast.
- Multicast MAC flooding is not supported for NLB.
- ARP entries for NLB server IP addresses do not age out when there is still client traffic coming to the NLB servers, even after the NLB servers are no longer reachable.

Procedure

1. In the navigation pane, expand the following folders: **Configuration --> VLAN**.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. In the row for the VLAN, double-click the value in the **NlbMode** column.
5. Select the appropriate value.
6. Click **Apply**.

Configuring a port to accept tagged or untagged frames

Configure a port to accept tagged or untagged frames.

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 **General**.
4. Click the **VLAN** tab.
5. To configure tagging on the port, select the **PerformTagging** check box.

This setting applies to all VLANs associated with the port.

 **Important:**

If the check box is selected, tagging is enabled. All frames sent from this port are tagged.

If the check box is cleared, tagging is disabled. The port does not send tagged frames. The switch removes the tag before sending the frame out of the port.

6. To discard tagged frames on a port for which tagging is disabled, select **DiscardTaggedFrames**.
7. To discard untagged frames on a port for which tagging is enabled, select **DiscardUntaggedFrames**.

8. To designate a default VLAN to associate with a packet that does not match a policy-based VLAN, enter a VLAN ID in the **DefaultVlanId** box or use the default VLAN 1.
9. Click **Apply**.
10. Click **Close**.

Configuring untagging default VLAN on a tagged port

Configure an untagged default VLAN on a tagged port to separate untagged packets originating from a PC from the tagged packets originating from an IP phone.

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 **General**.
4. Click the **VLAN** tab.
5. Select **UntagDefaultVlan**.
6. In the **DefaultVlanId**, enter a default VLAN ID.
7. Click **Apply**.
8. Click **Close**.

Configuring SLPP globally

Enable the Simple Loop Prevention Protocol (SLPP) to detect a loop and automatically stop it.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **SLPP**.
3. Click the **Global** tab.
4. Select **GlobalEnable**.
5. In the **TransmissionInterval** box, type a value for the time interval for loop detection.
6. Click **Apply**.

Global field descriptions

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

Name	Description
GlobalEnable	Activates or disables SLPP globally. The default is disabled.
TransmissionInterval	Configures the interval for which loop detection occurs. The interval is expressed in milliseconds in a range from 500–5000. The default value is 500.

Job aid

The following table provides the recommended SLPP values.

Table 21: SLPP recommended values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configuring the SLPP by VLAN

Activate SLPP on a VLAN to enable forwarding of the SLPP packet over the VLAN. This configuration controls the boundary of SLPP-PDU transmission.

Before you begin

- Enable SLPP globally before you configure it on a VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **SLPP**.
3. Click the **VLANS** tab.
4. Click **Insert**.
5. Click the **VlanId** ellipses (...).
6. Select the desired VLAN ID.
7. Click **Ok**.
8. Select **SlppEnable**.

- Click **Insert**.

Insert VLANs field descriptions

Use the data in the following table to use the **Insert VLANS** dialog box.

Name	Description
VlanId	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.
SlppEnable	<p>Activates SLPP on the selected VLAN.</p> <p>The SLPP packet transmission and reception process is active only if you enable the SLPP operation. If you disable the SLPP operation, the following occurs:</p> <ul style="list-style-type: none"> • the system sends no SLPP packets • the system discards received SLPP packets <p>The default is enabled.</p>

Configuring the SLPP by port

Use SLPP on a port to avoid traffic loops on the port.

Important:

To provide protection against broadcast and multicast storms, it is recommended that you enable Rate Limiting for broadcast traffic and multicast traffic.

Before you begin

- Enable SLPP globally before you configure it on a port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **SLPP**.
3. Click the **Ports** tab.
4. Double-click the **PktRxThreshold** box for the desired port to edit the threshold value for packet reception.
5. Double-click the **SlppEnable** box for the desired port.
6. Select **true** to enable SLPP.

7. Click **Apply**.

Ports field descriptions

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

Name	Description
IfIndex	Specifies the interface index number for a port.
PktRxThreshold	Specifies the threshold for packet reception. Configure the SLPP packet receive threshold to a value (1- 500) that represents the number of received SLPP-PDUs to shut down the port. This variable is a port-level parameter, therefore if the port is tagged, SLPP-PDUs from the various VLANs increment this single threshold counter. The default is 1.
SlppEnable	Activates SLPP on the selected interface. The default is disabled.
IncomingVlanId	Shows the VLAN ID of the classified packet on a port disabled by SLPP.
SrcNodeType	Specifies the source node type of the received SLPP packet.
PktRxCount	Shows the total number of SLPP packets the port received.
TimeToClrPktRxCount	Specifies the timer to clear the SLPP receive counter. After you enable SLPP and the port receives SLPP PDUs, the timer starts. After the timer exceeds the configured value, the system resets the count to zero. The default is 21,600 seconds.
RemainingTimeToClrPktRxCount	Shows the time remaining before the SLPP receive counter is reset to zero.

Job aid

The following table provides the recommended SLPP values.

Table 22: SLPP recommended values

Enable SLPP	Setting
Access SMLT	Yes
Core SMLT	No
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	

Table continues...

Enable SLPP	Setting
Packet Rx threshold	50
Transmission interval	500 ms (default)

Configuring directed broadcast on a VLAN

Configure directed broadcast on a VLAN to enable or disable directed broadcast traffic forwarding for an IP interface.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN** folders.
2. Click **VLANs**.
3. Select the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **Direct Broadcast** tab.
7. Select **DirectBroadcastEnable**.

 **Important:**

Configure multiple VLANs or IPs in the same subnet but in different systems simultaneously.

8. Click **Apply**.

Direct Broadcast field descriptions

Use the data in the following table to use the **Direct Broadcast** tab.

Name	Description
DirectBroadcastEnable	<p>Specifies that an Isolated Routing Port (IRP) can forward directed broadcast traffic. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcast on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling this function protects a host from possible denial of service (DoS) attacks.</p> <p>With the feature enabled, the Control Processor (CP) does not receive a copy of the directed broadcast. As a result, the system does not respond to a subnet broadcast ping sent from a remote subnet.</p> <p>The default is disabled.</p>

Configuring the forwarding database timeout globally

Configure the forwarding database global timeout to age out dynamically learned forwarding information.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**
2. Click **VLANS**.
3. Click the **FdbAging** tab.
4. Type an interval, in seconds, for aging out dynamically learned forwarding information, or keep the default.
5. Click **Apply**.

FDB Aging field descriptions

Use the data in the following table to use the **FDB Aging** tab.

Name	Description
FdbAging	Specifies the timeout period for dynamically learned mac addresses on the vlan. The default value is 300.

Viewing VLAN forwarding database information

Perform this procedure to view forwarding database entries for all VLANs on the device.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANS**.
3. In the VLANs tab, click the **Forwarding** tab.

Forwarding field descriptions

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

Name	Description
VlanId	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

Table continues...

Name	Description
	IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
Address	Specifies a unicast MAC address for which the VLAN has forwarding or filtering information.
Status	Specifies the status of the VLAN. The values are: <ul style="list-style-type: none"> • other • invalid • learned • self • mgmt
Port	Specifies either a value of zero (0) or the port number of the port on which a frame having the specified MAC address was seen. A value of cpp indicates a self-assigned MAC address.
BMac	Shows the backbone MAC address if the entry is learned from a Shortest Path Bridging MAC (SPBM) network.
Cvid	Specifies the customer VID.

Viewing the forwarding database for a specific VLAN

Use the forwarding database for VLANs to determine how the system forwards a received frame.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**
2. Click **VLANs**.
3. Select a VLAN.
4. Click **Bridge**.
5. Click the **Forwarding** tab and the VLAN forwarding database information is displayed.

Forwarding field descriptions

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

Name	Description
VlanId	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

Table continues...

Name	Description
	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.
Address	Specifies a unicast MAC address for which the bridge has forwarding or filtering information.
Status	Specifies the status. Values include the following: <ul style="list-style-type: none"> • self—one of the bridge addresses • learned—a learned entry that is being used • mgmt—a static entry
Port	Specifies either a value of zero (0) or the port number of the port on which a frame having the specified MAC address was seen. A value of cpp indicates a self-assigned MAC address.
Cvid	Identifies the customer VID for this interface.

Clearing learned MAC addresses by VLAN

Use the clear learned MAC addresses feature to flush the bridge forwarding database.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.
3. Click the **Advanced** tab.
4. Double-click in the **VLAN Operation Action** field.
5. Choose **FlushMacFdb** from the list.
6. Click **Apply**.

Clearing learned MAC addresses for all VLANs by port

Use the following procedure to clear all the forwarding database (FDB) for VLANs associated with this port.

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 **General**.

4. In the Interface tab **Action** box, select **FlushMacFdb**.
5. Click **Apply**.

All learned MAC addresses are cleared from the forwarding database (FDB) for VLANs associated with this port.

6. Click **Close**.

Viewing blocked MAC address information

Perform this procedure to view blocked MAC address information on Spoof-Detect-enabled ports and VLANs.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IP**.
3. Click the **SpoofMac** tab.

SpoofMac field descriptions

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

Name	Description
PortNum	Identifies the Spoof-Detect-enabled port number where the spoofed MAC is detected.
VlanId	Identifies the VLAN to which the port belongs.
IpAddress	Shows the local IP address on this interface.
MacAddress	Shows the spoofing MAC address detected on this interface.
ClearMac	Indicates that the switch will clear or delete the Spoof-Detect MAC entry when this field is true.

Configuring static forwarding

Configure static forwarding to specify the group of ports that are allowed to forward frames.

Important:

Entries are valid for unicast and for group/broadcast addresses.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN**.
2. Click **VLANs**.

3. Select the **Basic** tab.
4. Select a VLAN.
5. Click **Bridge**.
6. In the Bridge, VLAN tab, click the **Static** tab.
7. Click **Insert**.
8. In the **MacAddress** box, enter a forwarding destination MAC address.
9. In the **Port** box, click the ellipsis button (...).
10. Select the port on which the frame is received.
11. Click **Ok**.
12. Click **Insert**.

Static field descriptions

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

Name	Description
MacAddress	Specifies the destination MAC address in a frame to which the forwarding information for this entry applies. This object can take the value of a unicast address.
Port	Specifies the port number of the port on which the frame is received.
VlanId	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.
Status	Specifies the status of the VLAN.

Configuring limit learning

Limit MAC address learning to limit the number of forwarding database (FDB) entries learned on a particular port to a user-specified value. After the number of learned forwarding database entries reaches the maximum limit, MAC learning stops at that port.

 **Note:**

This feature is not supported on all hardware platforms. If you do not see this command in the command list or EDM, the feature is not supported on your hardware. For more information about feature support, see *Release Notes*.

Procedure

1. In the Device Physical View tab, select a port or multiple ports.
2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
3. Click **General**.
4. Click the **Limit-Learning** tab.
5. Configure the parameters as required.

Limit Learning field descriptions

Use the data in the following table to use the **Limit-Learning** tab.

Name	Description
PortNum	Shows the slot and port number to configure.
MaxMacCount	Configures the number of entries in the MAC table for the port that causes learning to stop. The default is 1024.
CurrentMacCount	Shows the number of entries currently in the MAC table for the port.
Enable	Enables or disables limit learning for the port.
MacLearning	Shows if MAC learning is enabled or disabled for the port.

Chapter 4: Spanning Tree configuration

This chapter provides concepts and procedures to configure the Spanning Tree features supported on the switch.

Spanning Tree fundamentals

The switch supports Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP).

Spanning Tree

Spanning Tree protocols detect and eliminate logical loops in a bridged or switched network. If multiple paths exist, the spanning Tree algorithm configures the network so that a bridge or device uses the root bridge path based on hop counts. Although link speed is taken into account, the path is based on the root bridge rather than on an optimized path. If that path fails, the protocol automatically reconfigures the network and makes another path active, thereby sustaining network operations. The switch supports RSTP and MSTP but can downgrade a port automatically if it receives an STP Bridge Protocol Data Unit (BPDU) from a switch that runs STP.

Note:

Spanning Tree is disabled on all Switched UNI (S-UNI) ports. The ports will move into forwarding state as soon as the physical port or VLACP or LACP comes up on the port. If the platform VLAN is associated to the S-UNI Service Instance Identifier (I-SID), then the S-UNI ports added to the platform VLAN will become the member of MSTP instances associated with the platform VLAN. To enable SLPP on the S-UNI ports, the platform VLAN must be associated with the S-UNI I-SID.

Spanning Tree Groups

Spanning Tree Groups (STGs) represent logical topologies. A topology is created based on bridge configuration values such as root bridge priority. In the case of multiple STGs, you can map a VLAN to the most appropriate logical topology in the physical network.

The switch supports Spanning Tree modes RSTP and MSTP. The default Spanning Tree mode is MSTP. The default STG is 0. In RSTP mode, all VLANs run in the default STG. In MSTP mode, you can create additional STGs by using the VLAN create command. The switch supports up to 64 STGs.

Although STP and MSTP are variations of the same Spanning Tree protocol, they communicate information differently. A switch in MSTP mode cannot recognize the Spanning Tree groups running on a chassis configured with STP. MSTP Spanning Tree groups are not the same as STP Spanning

Tree groups. Using a switch in MSTP mode with a chassis in STP mode can create a loop in the network.

The root bridge for Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP) is determined by comparing attributes of each bridge in the network.

The protocol considers bridge priority first. If more than one bridge has the same priority, then the protocol must consider the bridge ID. The bridge with the lowest ID becomes the root bridge. For MSTP, this bridge is called the Common and Internal Spanning Tree (CIST) Root because it is the root of the entire physical network.

In MSTP mode, you can create additional Spanning Tree instances, by using the VLAN command. These instances, known as Multiple Spanning Tree Instances (MSTIs), can assign different priorities to switches. The MSTIs have different link costs or port priorities and as a result create separate logical topologies.

MSTP also allows the creation of MSTP regions. A region is a collection of switches sharing the same view of physical and logical topologies. For switches to belong to the same region, the following attributes must match:

- MSTP configuration ID selector
- MSTP configuration name
- MSTP configuration revision number
- VLAN instance mapping

Links connecting sections are called boundary ports. In a region, the boundary switch that contains the boundary port providing the shortest external path cost to the CIST Root is the CIST Regional Root.

STGs and VLANs

When you map VLANs to STGs, be aware that all links on the bridge belong to all STGs. Because each Spanning Tree group can differ in its decision to make a link forwarding or blocking, you must ensure that the ports you add to a VLAN are in the expected state.

Untagged ports can only belong to one VLAN and therefore can only belong to one STG. Tagged ports can belong to multiple VLANs and therefore to multiple STGs.

BPDU handling on S-UNI port/MLT

The switch handles Bridge Protocol Data Units (BPDUs) according to whether or not you configure a platform VLAN.

- When you configure a platform VLAN:
 - BPDUs are forwarded to the CPU by default.
 - BPDUs are not flooded in the S-UNI I-SID associated with the platform VLAN.

Note:

If the platform VLAN is configured for the S-UNI port, you cannot enable BPDU forwarding.

- When you DO NOT configure a platform VLAN:
 - BPDUs received on untagged-traffic ports are dropped by default.
 - To flood BPDUs in its I-SID, enable BPDU forwarding under S-UNI I-SID using the command `untagged-traffic port <port no> bpdu enable`.

BPDU Guard

The switch supports Bridge Protocol Data Unit (BPDU) Guard for STGs, RSTP, and MSTP.

Overview

Spanning Tree eliminates loops in a network. A bridge that participates in spanning tree uses BPDUs to exchange information with other bridges. The bridges select a single bridge as the root bridge based on the BPDU information exchange. The bridge with the lowest priority becomes the root bridge. If all bridges share the same priority, the bridge with the lowest bridge ID becomes the root bridge. This process is the root selection process.

After you add a new bridge to the network, or remove an existing bridge, the bridges repeat the root selection process, and then select a new root bridge.

To ensure the correct operation of Spanning Tree in the network, BPDU Guard protects the stability of the Root Bridge by dropping stray, unexpected, or unwanted BPDU packets entering a port, and immediately shutting down those ports for a specified time period. BPDU Guard is normally enabled on access ports connecting to end user devices such as servers that are not expected to operate Spanning Tree.

Use BPDU Guard to achieve the following results:

- Block the root selection process after an edge device, such as a laptop that uses Linux with STP enabled, is added to the network. Blocking the root selection process prevents unknown devices from influencing the spanning tree topology.
- Block BPDU flooding of the switch from an unknown device.

Operation

You can enable or disable BPDU Guard on an individual port basis, regardless of the spanning tree state. Each port uses a timer to determine port-state recovery.

After you enable BPDU Guard on a port and the port receives a BPDU, the following actions occur:

1. The guard disables the port.
2. The switch generates an SNMP trap and alarm, and the following log message:

```
BPDU Guard - Port <slot/port> is being shutdown by BPDU Guard,  
timeout <time_seconds>
```

3. The port timer begins.
4. The port remains in the disabled state until the timer expires.

If you disable BPDU Guard before the timer expires, the timer stops and the port remains in the disabled state. You must manually enable the port.

BPDU Guard is enabled at the interface level. You can configure the BPDU Guard timer for each port, for 10 to 65535 seconds. If you set the port timer to zero, it will not expire.

Rapid Spanning Tree Protocol and Multiple Spanning Tree Protocol

The Rapid Spanning Tree Protocol (RSTP or IEEE 802.1w) reduces the recovery time after a network breakdown. It also maintains backward compatibility with IEEE 802.1d (the spanning tree implementation prior to RSTP). In certain configurations, the recovery time of RSTP can be reduced to less than 1 second. RSTP also reduces the amount of flooding in the network by enhancing the way Topology Change Notification (TCN) packets are generated.

With Multiple Spanning Tree Protocol (MSTP or IEEE 802.1s), you can configure multiple instances or Spanning Tree groups on the same device. Each instance or Spanning Tree group can include one or more VLANs.

By using RSTP and MSTP, the switch achieves the following:

- reduces convergence time after a topology change (from 30 seconds to less than 1 or 2 seconds)
- eliminates unnecessary flushing of the MAC database and the flooding of traffic to the network
- creates backward compatibility with classic 802.1d switches
- creates support for 64 instances of spanning tree in MSTP mode

The following sections relate to RSTP and MSTP:

- [RSTP interoperability with STP](#) on page 98
- [Differences in port roles for STP and RSTP](#) on page 99
- [Port roles: root forwarding role](#) on page 99
- [Port roles: designated forwarding role](#) on page 100
- [Port roles: alternate blocking role](#) on page 100
- [Edge port](#) on page 100
- [Path cost values](#) on page 100
- [RSTP negotiation process](#) on page 100

RSTP interoperability with STP

RSTP provides a parameter called ForceVersion to provide backward compatibility with standard STP. A user can configure a port in either STP-compatible mode or RSTP mode:

- An STP-compatible port transmits and receives only STP Bridge Protocol Data Units (BPDUs). An RSTP BPDU that the port receives in this mode is discarded.
- An RSTP-compatible port transmits and receives only RSTP BPDUs. If an RSTP port receives an STP BPDU, it becomes an STP port. User intervention is required to change this port back to RSTP mode. This process is called Port Protocol Migration.

Note:

You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.

You must be aware of the following recommendations before you implement MSTP or RSTP:

- The default mode is MSTP. A special boot configuration flag identifies the mode.
- You can lose your configuration if you change the spanning tree mode from MSTP to RSTP and the configuration file contains VLANs configured with MSTI greater than 0. RSTP only supports VLANs configured with the default instance 0.
- For best interoperability results, contact your vendor representative.

Differences in port roles for STP and RSTP

RSTP is an enhanced version of STP. These two protocols have almost the same parameters.

The following table lists the differences in port roles for STP and RSTP. STP supports two port roles, while RSTP supports four port roles.

Table 23: Differences in port roles for STP and RSTP

Port Role	STP	RSTP	Description
Root	Yes	Yes	This port receives a better BPDU than its own and has the best path to reach the Root. The root port is in Forwarding state. The root port and designated ports can be in the Discarding state before they go to root forwarding.
Designated	Yes	Yes	This port has the best BPDU on the segment. The designated port is in the Forwarding state.
Alternate	No	Yes	This port receives a better BPDU than its own BPDU, and a root port exists within the same device. The alternate port is in the Discarding state.
Backup	No	Yes	This port receives a better BPDU than its own BPDU, and this BPDU is from another port within the same device. The backup port is in the Discarding state.

Port roles: root forwarding role

MSTP and RSTP root forwarding roles are as follows:

- The port that receives the best path BPDU on a device is the root port, and is referred to as a Root Forwarding (RF) port. This is the port that is the closest to the root bridge in terms of path cost.
- The spanning tree algorithm elects a single root bridge in a bridged network. With MSTP, a root bridge is selected for the Common and Internal Spanning Tree (CIST). A root bridge is selected for the region, and a root bridge is selected for each spanning tree instance.
- The root bridge is the only bridge in a network that does not have root ports; all ports on a root bridge are Designated Forwarding (DF).
- Only one path towards a root bridge can exist on a given segment; otherwise, loops can occur.

Port roles: designated forwarding role

MSTP and RSTP designated forwarding roles are as follows:

- All bridges connected on a segment monitor the BPDUs of all other bridges. The bridge that sends the best BPDU is the root bridge for the segment.
- The corresponding port on the bridge is referred to as a Designated Forwarding Port.

Port roles: alternate blocking role

MSTP and RSTP alternate blocking roles are as follows:

- A blocked port is defined as not being the designated or root port. An alternate port provides an alternate path to the root and can replace the root port if it fails.
- An alternate blocked port is a port that is blocked because it received better path cost BPDUs from another bridge.

Port roles: backup blocking role

MSTP and RSTP backup blocking roles are as follows:

- A backup port receives the more useful BPDUs from the bridge on which the port exists.

Edge port

RSTP uses a parameter called the edge port. After a port connects to a nonswitch device, such as a PC or a workstation, it must be configured as an edge port. An active edge port enters the forwarding state without delay. An edge port becomes a nonedge port if it receives a BPDU.

Path cost values

RSTP and MSTP recommend path cost values that support a wide range of link speeds. The following table lists the recommended path cost values.

Table 24: Recommended path cost values

Link speed	Recommended value
Less than or equal to 100 Kbps	200 000 000
1 Mbps	20 000 000
10 Mbps	2 000 000
100 Mbps	200 000
1 Gbps	20 000
10 Gbps	2000
100 Gbps	200
1 Tbps	20
10 Tbps	2

RSTP negotiation process

The following section describes the negotiation process between switches that takes place before PCs can exchange data (see the following figure).




Figure 8: RSTP negotiation process

After turning on, all ports assume the role of designated ports. All ports are in the discarding state except edge ports. Edge ports go directly into the forwarding state without delay.

Unit A port 1 and Unit B port 1 exchange BPDUs. Unit A knows that it is the root and that Unit A port 1 is the designated port. Unit B learns that Unit A has higher priority. Unit B port 1 becomes the root port. Both Unit A port 1 and Unit B port 1 are still in the discarding state.

Unit A starts the negotiation process by sending a BPDU with the proposal bit set.

Unit B receives the proposal BPDU and configures its nonedge ports to discarding state. This operation occurs during the synchronization process.

Unit B sends a BPDU to Unit A with the agreement bit set.

Unit A configures port 1 to the forwarding state, and Unit B configures port 1 to the forwarding state. PC 1 and PC 2 can now communicate. The negotiation process now moves on to Unit B port 3 and its partner port. PC 3 cannot exchange data with either PC 1 or PC 2 until the negotiation process between Unit B and Unit C finishes.

The RSTP convergence time depends on how quickly the switches can exchange BPDUs during the negotiation process, and on the number of switches in the network.

Spanning Tree configuration using CLI

This chapter describes how to configure the Spanning Tree mode, MSTP, and RSTP using Command Line Interface (CLI) commands.

! Important:

The switch supports up to 64 STGs on a device, however, SPBM uses STG 63 and MSTI 62 for internal use. STG 63 or MTSI 62 cannot be used by other VLANs or MSTIs.

Configuring Spanning Tree

Configure the STP mode to configure the spanning tree mode on the device.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the STP mode:

```
boot config flags spanning-tree-mode {rstp|mstp}
```

Example

Configure the STP mode:

```
Switch:1(config)# boot config flags spanning-tree-mode mstp
```

```
Warning: Please save the configuration and reboot the switch
for this to take effect.
```

```
Warning: Please carefully save your configuration files before
starting configuring the switch in RSTP or MSTP mode.
```

Variable definitions

Use the data in the following table to use the `boot config flags spanning-tree-mode` command.

Table 25: Variable definitions

Variable	Value
rstp mstp	Specifies the Spanning Tree modes: Rapid Spanning Tree Protocol (RSTP) or Multiple Spanning Tree Protocol (MSTP).

Configuring BPDU Guard

Configure BPDU Guard to block the root selection process or to prevent BPDU flooding from unknown devices.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

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

*** Note:**

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

2. Enable BPDU Guard for the port:

```
spanning-tree bpduguard enable
```

3. **(Optional)** Configure the timer for port-state recovery:

```
spanning-tree bpduguard timeout <0, 10-65535>
```

4. **(Optional)** Enable BPDU Guard on an additional port or group of ports:

```
spanning-tree bpduguard port {slot/port[/sub-port] [-slot/port[/
subport]][, ...]} enable
```

5. **(Optional)** Configure the timer for port-state recovery for an additional port or group of ports:

```
spanning-tree bpduguard port {slot/port[/sub-port] [-slot/port[/
subport]][, ...]} timeout <0-65535>
```

6. Verify the configuration:

```
show spanning-tree bpduguard [GigabitEthernet {slot/port[/sub-port]
[-slot/port[/subport]][, ...]] [{slot/port[/sub-port] [-slot/port[/
subport]][, ...]]]
```

Example

Enable BPDU Guard on port 1/8, and specify a timer value of 200 seconds. Verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#interface gigabitEthernet 1/8
Switch:1(config-if)#spanning-tree bpduguard enable
Switch:1(config-if)#spanning-tree bpduguard timeout 200
Switch:1(config-if)#show spanning-tree bpduguard 1/8
```

```
=====
                        Bpdu Guard
=====
Port      PORT      PORT      TIMER    BPDUGUARD
NUM MLTID  ADMIN_STATE  OPER_STATE  TIMEOUT  COUNT    ADMIN_STATE
-----
1/8      Up        Up         200        0        Enabled
```

Variable definitions

Use the data in the following table to use the **spanning-tree bpduguard** commands.

Variable	Value
enable	Enables BPDU Guard on the port. The default is disabled.
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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
timeout <0, 10-65535>	Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires. You can configure a value from 10 to 65535. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.

Use the data in the following table to use the **show spanning-tree bpduguard** 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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Configuring Rapid Spanning Tree Protocol

Configure Rapid Spanning Tree Protocol (RSTP) to reduce the recovery time after a network breakdown.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure RSTP:

```
spanning-tree rstp [forward-time <400-3000>] [group-stp enable]
[hello-time <100-1000>] [max-age <600-4000>] [pathcost-type <bits16|
bits32>] [priority <0-61440>] [tx-holdcount <1-10>] [version <rstp|
stp-compatible>]
```


Example

Configure RSTP:

```
Switch:1(config)# spanning-tree rstp forward-time 1000 hello-time 200 max-age 4000 pathcost-type bits16 priority 4096 tx-holdcount 10 version rstp group-stp enable
```

Variable definitions

Use the data in the following table to use the `spanning-tree rstp` command.

Table 26: Variable definitions

Variable	Value
forward-time <400-3000>	Configures the RSTP forward delay for the bridge in hundredths of a second.
group-stp enable	Enables or disables RSTP for a specific STG. Enter the no form of the command to disable RSTP for the STG (no <code>spanning-tree rstp group-stp enable</code>).
hello-time <100-1000>	Assigns the RSTP hello time delay for the bridge in hundredths of a second.
max-age <600-4000>	Assigns the RSTP maximum age time for the bridge in hundredths of a second.
pathcost-type {bits16 bits32}	Assigns the RSTP default pathcost version. The default is 32 bits.
priority <0-61440>	Assigns the RSTP bridge priority.
tx-holdcount <1-10>	Assigns the RSTP transmit hold count from 1 to 10. The default value is 6.
version {rstp/stp-compatible}	Sets the version to RSTP or STP compatible.

Configuring Rapid Spanning Tree Protocol for a port

Configure RSTP to reduce the recovery time after a network breakdown.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

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

*** Note:**

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

2. Configure RSTP:

```
spanning-tree rstp cost <1-200000000> edge-port <false|true> p2p
<auto|force-false|force-true> priority <0-240> protocol-migration
<false|true> stp enable
```

Example

Configure RSTP:

```
Switch:1(config-if)# spanning-tree rstp cost 100 edge-port true p2p auto
priority 32 protocol-migration true stp enable
```

Variable definitions

Use the data in the following table to use the `spanning-tree rstp` command.

Table 27: Variable definitions

Variable	Value
cost <1-200000000>	Specifies the contribution of this port to the path cost.
edge-port <false true>	Configures the edge-port value for the port. A value of true indicates that this port is an edge-port, and a value of false indicates that this port is a nonedge-port.
p2p <auto force-false force-true>	Specifies the point-to-point status of the LAN segment attached to this port. A value of force-true indicates that this port is treated as if it connects to a point-to-point link. A value of force-false indicates that this port is treated as having a shared media connection. A value of auto indicates that this port is considered to have a point-to-point link if it is an aggregator and all of its members are aggregatable, or if the MAC entity is configured for full-duplex operation, either through autonegotiation or by management means.
priority <0-240>	Assigns the RSTP bridge priority in a range of 0–240. The value has to increment in steps of 16.
protocol-migration <false true>	If you chose true, the option initiates protocol migration for a port. If you chose false, the option terminates protocol migration for a port. An RSTP-compatible port transmits and receives only RSTP BPDUs. If an RSTP port receives an STP BPDU, it becomes an STP port. User intervention is

Table continues...

Variable	Value
	required to change this port back to RSTP mode. This process is called Port Protocol Migration.
stp enable	Configures STP for the port.

Configuring the Rapid Spanning Tree Protocol version

Perform this procedure to specify the RSTP mode.

Procedure

1. Enter Global Configuration mode:


```
enable
configure terminal
```
2. Configure Rapid Spanning Tree Protocol version:


```
spanning-tree rstp version {rstp|stp-compatible}
```

Example

Configure Rapid Spanning Tree Protocol version:

```
Switch:1(config)# spanning-tree rstp version rstp
```

Variable definitions

Use the data in the following table to use the `spanning-tree rstp version` command.

Table 28: Variable definitions

Variable	Value
<code>rstp version {rstp stp-compatible}</code>	Sets the version to RSTP or to STP compatible. The default is RSTP.

Viewing the global RSTP configuration information

View the global RSTP configuration information to display the Rapid Spanning Tree Protocol (RSTP) configuration details.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View global RSTP configuration information:


```
show spanning-tree rstp config
```

Example

View global RSTP configuration information:

```
Switch:1> show spanning-tree rstp config
```

```
=====
                        RSTP Configuration
=====
Rstp Module Status      : Enabled
Priority                 : 32768 (0x8000)
Stp Version              : rstp Mode
Bridge Max Age          : 20 seconds
Bridge Hello Time       : 2 seconds
Bridge Forward Delay Time : 15 seconds
Tx Hold Count           : 6
PathCost Default Type   : 32-bit
```

Viewing RSTP statistics

Perform this procedure to view RSTP statistics.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RSTP statistics:

```
show spanning-tree rstp statistics
```

Example

View RSTP statistics:

```
Switch:1> show spanning-tree rstp statistics
```

```
=====
                        RSTP Statistics
=====
Rstp UP Count           : 1
Rstp Down Count         : 0
Count of Root Bridge Changes : 0
Stp Time since Topology change: 0 day(s), 00H:00M:00S
Total No. of topology changes : 0
```

Viewing the RSTP status

View the RSTP status to display the RSTP related status information for the selected bridge.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View the RSTP status:

```
show spanning-tree rstp status
```

Example

View the RSTP status:

```
Switch:1> show spanning-tree rstp status
```

```

=====
                        RSTP Status Information
=====
Designated Root          : 80:00:00:24:7f:9f:60:00
Stp Root Cost            : 0
Stp Root Port            : cpp
Stp Max Age              : 20 seconds
Stp Hello Time           : 2 seconds
Stp Forward Delay Time   : 15 seconds

```

Viewing the RSTP configuration information

View the RSTP configuration information to display the RSTP-related port level configuration details.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View RSTP configuration information:

```
show spanning-tree rstp port config {slot/port[/sub-port]}[-slot/
port[/sub-port]][,...]
```

* Note:

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

View RSTP configuration information:

```
Switch:1> show spanning-tree rstp port config 1/1
```

```

=====
                        RSTP Port Configurations
=====
Port Number                : 1/1
Port Priority               : 128 (0x80)
Port PathCost              : 200000000
Port Protocol Migration    : False
Port Admin Edge Status    : False
Port Oper Edge Status     : False
Port Admin P2P Status     : Auto
Port Oper P2P Status      : False
Port Oper Protocol Version : Rstp

```

Variable definitions

Use the data in the following table to use optional parameters with the `show spanning-tree rstp port config` command.

Table 29: Variable definitions

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

Viewing the RSTP status for a port

View the RSTP status for a port to display the RSTP-related status information for a selected port.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View the RSTP status for a port:

```
show spanning-tree rstp port status {slot/port[/sub-port][/-slot/
port[/sub-port]][,...]}
```

*** Note:**

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

View the RSTP status for a port:

```
Switch:1> show spanning-tree rstp port status 1/2
```

```
=====
                        RSTP Port Status
                        (Port Priority Vector)
=====
Port Number           : 1/2
Port Designated Root  : 80:00:00:24:7f:9f:60:00
Port Designated Cost  : 0
Port Designated Bridge : 80:00:00:24:7f:9f:60:00
Port Designated Port  : 80:c1
```

Variable definitions

Use the data in the following table to use optional parameters with the `show spanning-tree rstp port status` command.

Table 30: Variable definitions

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

Viewing RSTP information for a selected port

View the RSTP information for a selected port to display the RSTP-related configuration information for the selected port.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. View the RSTP information for a selected port:

```
show spanning-tree rstp port statistics [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]]
```

Example

View the RSTP information for a selected port:

```
Switch:1# show spanning-tree rstp port statistics 1/4
```

```
=====
                        RSTP Port Statistics
=====
Port Number                : 1/4
Number of Fwd Transitions  : 0
Rx RST BPDUs Count        : 0
Rx Config BPDU Count      : 0
Rx TCN BPDU Count         : 0
Tx RST BPDUs Count        : 9
Tx Config BPDU Count      : 0
Tx TCN BPDU Count         : 0
Invalid RST BPDUs Rx Count : 0
Invalid Config BPDU Rx Count : 0
Invalid TCN BPDU Rx Count : 0
Protocol Migration Count   : 0
=====
```

Variable definitions

Use the data in the following table to use optional parameters with the `show spanning-tree rstp port statistics` command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}</code>	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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Viewing the RSTP role

View the RSTP role to display the RSTP information.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View the RSTP role:

```
show spanning-tree rstp port role [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]]
```

Example

View the RSTP role:

```
Switch:1> show spanning-tree rstp port role 1/3
```

```
=====
                        RSTP Port Roles and States
=====
Port-Index  Port-Role  Port-State  PortSTPStatus  PortOperStatus
-----
1/3         Designated  Forwarding  Enabled         Enabled
```

Variable definitions

Use the data in the following table to use optional parameters with the **show spanning-tree rstp port role** command.

Variable	Value
<code>{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}</code>	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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Viewing spanning tree configuration

Perform this procedure to view configuration and status information for spanning tree in your network.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View spanning tree configuration information:

```
show spanning-tree config
```

3. View spanning tree status information:

```
show spanning-tree status
```

Example

View spanning tree configuration information:

```
Switch:1> show spanning-tree config
```

```

=====
                        Spanning Tree Config
=====
ID      PRIORITY  BRIDGE  BRIDGE  FORWARD
      MAX_AGE  HELLO_TIME  DELAY  STATE
-----
0       32768     20      0       15     Enabled
1       32768     20      0       15     Enabled

ID      TAGGBPDU  PORT
      ADDRESS  MEMBER
-----
0       01:80:c2:00:00:00  mstp  1/1-1/9,1/11-1/48
1       01:80:c2:00:00:00  mstp  1/10

Total number of Spanning Tree IDs :  2

```

View spanning tree status information:

```
Switch:1> show spanning-tree status
```

```

=====
                        Spanning Tree Status
=====
STG  BRIDGE  NUM  PROTOCOL  TOP
ID   ADDRESS  PORTS SPECIFICATION  CHANGES
-----
0    00:24:7f:a1:70:00  47   ieee8021s  1
1    00:24:7f:a1:70:00  1    ieee8021s  1

STG  DESIGNATED  ROOT  ROOT  MAX  HELLO  HOLD  FORWARD
ID   ROOT        COST  PORT  AGE  TIME   TIME  DELAY
-----
0    80:00:00:24:7f:a1:70:00  0    cpp  20  0     1    15
1    80:00:00:24:7f:a1:70:00  0    cpp  20  0     1    15

Total number of Spanning Tree IDs :  2

```

Configuring Multiple Spanning Tree Protocol

Use the following procedure to configure the Multiple Spanning Tree Protocol.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure MSTP:

```
spanning-tree mstp
```

Example

Configure Multiple Spanning Tree Protocol to configure the MSTP configuration version.

```
Switch:1(config)# spanning-tree mstp forward-time 500 max-age 3000 max-hop
200 pathcost-type bits32 priority 8192 tx-holdcount 10 version mstp
```

Variable definitions

Use the data in the following table to use the `spanning-tree mstp` command.

Table 31: Variable definitions

Variable	Value
forward-time <400-3000>	Configures the MSTP forward delay for the bridge from 400 to 3000 hundredths of a second.
max-age <600-4000>	Assigns the MSTP maximum age time for the bridge from 600 to 4000 one hundredths of a second.
max-hop <100-4000>	Assigns the MSTP bridge maximum hop count. The range is 100 to 4000 one hundredths of a second. The original MIB erroneously designated the value in hundredths of a second, when it should have been in hops. The replacement MIB kept the range at 100-4000 to remain backwards compatible. To convert this value to hops, divide by 100 so 100-4000 equals 1-40 hops.
msti <1-63> priority <0-65535>	Assigns the MSTP MSTI instance parameter.
pathcost-type {bits16 bits32}	Assigns the MSTP default pathcost type to either 16 bits or 32 bits. The default is 32 bits.
priority <0-61440>	Assigns the MSTP bridge priority in a range of 0 to 61440 in steps of 4096.

Table continues...

Variable	Value
region [config-id-sel <0-255>] [region-name <WORD 1-32>] [region-version <0-65535>]	<p>Assigns the MSTP region commands:</p> <ul style="list-style-type: none"> • config-id-sel—Assigns the MSTP region configuration ID number. The range is 0 to 255. • region-name—Assigns the MSTP region name. The character string can be a range of 1 to 32 characters • region-version—Assigns the MSTP region version. The range is 0 to 65535.
tx-holdcount <1-10>	Assigns the MSTP transmit hold count. The range is 1 to 10. The default value is 3.
version {mstp rstp stp-compatible}	<p>Assigns the bridge version.</p> <p>Although STP and MSTP are variations of the same spanning tree protocol, they communicate information differently. A switch in MSTI mode cannot recognize the spanning tree groups running on a chassis configured with STP. MSTP spanning tree groups are not the same as STP spanning tree groups. Using a switch in MSTP mode with another chassis in STP mode can create a loop in the network.</p> <p>You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.</p>

Configuring MSTP MSTI options

Use the following procedure to configure MSTP multiple spanning tree instance (MSTI) options.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure MSTP MSTI:

```
spanning-tree mstp msti <1-63> priority <0-65535>
```

Example

Configure MSTP MSTI:

```
Switch:1(config)# spanning-tree mstp msti 62 priority 4096
```

Variable definitions

Use the data in the following table to use the `spanning-tree mstp msti <1-63> priority <0-65535>` command.

Table 32: Variable definitions

Variable	Value
<1-63>	Specifies the instance ID.
<0-65535>	Specifies the priority value. Enter values in increments of 4096: <ul style="list-style-type: none"> • 4096 • 8192 • 12288 • 16384 • 20480 • 24576 • 28672 • 32768 • 36864 • 40960 • 45056 • 49152 • 53248 • 57344 • 61440

Configuring Ethernet MSTP

Configure Ethernet MSTP on a port to enable this feature.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
```

```
configure terminal
```

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

*** Note:**

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

2. Configure Ethernet MSTP:

```
spanning-tree mstp [cost <1-200000000>] [edge-port <false|true>]
[force-port-state enable] [hello-time <100-1000>] [msti <1-63>] [p2p
{auto|force-false|force-true}] [port {slot/port[/sub-port]}]
[priority <0-240>] [protocol-migration <false|true>]
```

*** Note:**

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

Configure Ethernet MSTP:

```
Switch:1(config)# spanning-tree mstp cost 1 edge-port true force-port-
state enable hello-time 100 p2p auto priority 2 protocol-migration true
```

Variable definitions

Use the data in the following table to use the `spanning-tree mstp` command.

Table 33: Variable definitions

Variable	Value
cost <1-200000000>	Configures the path cost for a port. Valid values are 1 to 200000000
edge-port <false true>	Enables or disables the port as an edge port.
force-port-state enable	Enables STP.
hello-time <100-1000>	Configures the hello-time for a port.
msti <1-63>	Configures the port MSTP MSTI.
p2p {auto force-false force-true}	Enables or disables point-to-point for a port.
{slot/port[/sub-port]}	Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
priority <0-240>	Configures priority for the port.
protocol-migration {false true}	If you chose true, the option initiates protocol migration for a port. If you chose false, the option terminates protocol migration for a port. An MSTP-compatible port transmits and receives only RSTP BPDUs. If an MSTP port receives an STP

Table continues...

Variable	Value
	<p>BPDU, it becomes an STP port. User intervention is required to change this port back to MSTP mode. This process is called Port Protocol Migration.</p> <p>You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.</p>

Configuring Ethernet MSTP MSTI

Use the following procedure to configure the Ethernet MSTP MSTI parameters on a port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

*** Note:**

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

2. Configure Ethernet MSTP MSTI:

```
spanning-tree mstp msti <1-63> [cost <1-200000000>] [force-port-state enable] [port {slot/port[/sub-port] [-slot/port[/sub-port]] [, ...]] [priority <0-240>]
```

Example

Configure Ethernet MSTP MSTI:

```
Switch(config-if)# spanning-tree mstp msti 62 priority 32
```

Variable definitions

Use the data in the following table to use the `spanning-tree mstp msti <1-63>` command.

Table 34: Variable definitions

Variable	Value
<1-63>	Specifies the instance ID.

Table continues...

Variable	Value
cost <1–200000000>	Configures the path cost for the port
force-port-state enable	Enables MSTI learning for the 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 your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.
priority <0–240>	Configures the priority for the port. Enter the priority value (0–240) as increments of 16.

Viewing MSTP configurations

View the MSTP configurations to display the MSTP-related bridge-level VLAN and region information.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View the MSTP configurations:

```
show spanning-tree mstp config
```

Example

View the MSTP configurations:

```
Switch:1> show spanning-tree mstp config
```

```

=====
                        MSTP Configurations
=====
Mstp Module Status      : Enabled
Number of Msti Supported : 64
Cist Bridge Priority     : 32768 (0x8000)
Stp Version             : Mstp Mode
Cist Bridge Max Age     : 20 seconds
Cist Bridge Forward Delay : 15 seconds
Tx Hold Count           : 3
PathCost Default Type   : 32-bit
Max Hop Count           : 2000
Msti Config Id Selector : 0
Msti Region Name        : 00:15:e8:9e:10:01
Msti Region Version     : 0
Msti Config Digest      : b2:96:8d:23:9d:73:39:e4:4f:bd:94:c2:14:d4:8d:09
=====

```

Viewing MSTP status

View the MSTP status to display the MSTP-related status information known by the selected bridge.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View the MSTP status:

```
show spanning-tree mstp status
```

Example

View the MSTP status:

```
Switch:1> show spanning-tree mstp status
```

```

=====
                        MSTP Status
=====
Bridge Address           : 00:15:e8:9e:10:01
Cist Root                : 80:00:00:15:e8:9e:10:01
Cist Regional Root      : 80:00:00:15:e8:9e:10:01
Cist Root Port          : cpp
Cist Root Cost           : 0
Cist Regional Root Cost : 0
Cist Instance Vlan Mapped : 1-9,11-12,14-100,102-1024
Cist Instance Vlan Mapped2k : 1025-2048
Cist Instance Vlan Mapped3k : 2049-3072
Cist Instance Vlan Mapped4k : 3073-3999,4001-4094
Cist Max Age            : 20 seconds
Cist Forward Delay      : 15 seconds
=====

```

Viewing MSTP port information

View the MSTP port information to display the MSTP, CIST port, and MSTI port information maintained by every port of the common spanning tree.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. View the MSTP port information:

```
show spanning-tree mstp port role [slot/port[/sub-port]][-slot/port[/sub-port]][,...]
```

*** Note:**

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

Example

View the MSTP port information:

```
Switch:1> show spanning-tree mstp port role 1/3
```

```

=====
                        CIST Port Roles and States
=====

```


Port-Index	Port-Role	Port-State	PortSTPStatus	PortOperStatus
1/3	Disabled	Discarding	Enabled	Disabled

Viewing MSTP MSTI information

View MSTP MSTI information to ensure the feature is configured correctly for your network.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Show MSTI information:

```
show spanning-tree mstp msti [config <1-63>] [port <config {slot/
port[/sub-port] [-slot/port[/sub-port]] [, ...]}|role {slot/port[/sub-
port] [-slot/port[/sub-port]] [, ...]}|statistics {slot/port[/sub-port]
[-slot/port[/sub-port]] [, ...]}
```

* Note:

If your platform supports channelization for 40 Gbps ports and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Example

Show MSTI information:

```
Switch:1> show spanning-tree mstp msti config 62
```

```
=====
MSTP Instance Status
=====
Instance Id                : 62
Msti Bridge Regional Root  : 80:00:00:15:e8:9e:10:01
Msti Bridge Priority       : 32768 (0x8000)
Msti Root Cost             : 0
Msti Root Port             : cpp
Msti Instance Vlan Mapped  :
Msti Instance Vlan Mapped2k :
Msti Instance Vlan Mapped3k :
Msti Instance Vlan Mapped4k : 4000
```

```
Switch(config)# show spanning-tree mstp msti port statistics 1/1
```

```
=====
MSTP Instance-specific Per-Port Statistics
=====
Port Number                : 1/1
Instance Id                : 1
Msti Port Fwd Transitions  : 0
Msti Port Received BPDUs   : 0
Msti Port Transmitted BPDUs : 0
Msti Port Invalid BPDUs Rcvd : 0
```

Variable definitions

Use the data in the following table to use the `show spanning-tree mstp msti` command.

Table 35: Variable definitions

Variable	Value
config [<1-63>]	Shows the configuration for one or all MSTP instance IDs.
port	Shows the configuration, role, or statistics information of a MSTP port. <ul style="list-style-type: none"> • config {slot/port[/sub-port][/-slot/port[/sub-port]][,...]} • role {slot/port[/sub-port][/-slot/port[/sub-port]][,...]} • statistics {slot/port[/sub-port][/-slot/port[/sub-port]][,...]}

Viewing MSTP statistics

View MSTP MSTI information to ensure the feature is configured correctly for your network.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Show MSTP statistics:

```
show spanning-tree mstp statistics
```

Example

Show MSTP statistics:

```
Switch:1> show spanning-tree mstp statistics
```

```
=====
MSTP Bridge Statistics
=====
Mstp UP Count           : 1
Mstp Down Count         : 0
Region Config Change Count : 4
Time Since Topology Change : 0 seconds
Topology Change Count   : 0
New Root Bridge Count    : 1
```

Spanning Tree configuration using EDM

This chapter describes how to create, manage, and monitor spanning tree groups (STG). It also describes how to configure the Rapid Spanning Tree Protocol (RSTP) and the Multiple Spanning Tree Protocol (MSTP) using Enterprise Device Manager (EDM).

! Important:

The switch supports up to 64 STGs in a device, however, SPBM uses STG 63 and MSTI 62 for internal use. STG 63 or MTSI 62 cannot be used by other VLANs or MSTIs.

Configuring the Spanning Tree mode

Configure the Spanning Tree mode to change the mode to MSTP or RSTP mode.

! Important:

After you change the mode, restart the system for the changes to take effect.

Procedure

1. In the navigation pane, expand the **Configuration > VLAN > Spanning Tree** folders.
2. Click **Globals**.
3. Select the required spanning tree mode.
4. Click **Apply**.

The system notifies you that the setting takes effect after you save the configuration and restart the server.

Globals field descriptions

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

Name	Description
SpanningTreeAdminMode	Configures the spanning tree mode as either RSTP or MSTP. The default is MSTP.
SpanningTreeOperMode	Specifies the current mode of the spanning tree.

Restarting the switch

Restart the switch so that changes to the bootconfig parameters (or other parameters) take effect. For example, you must restart the device to enable a change to the Spanning Tree mode.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit**.
2. Click **Chassis**.
3. In the System tab, locate the **ActionGroup1** box.
4. Select **saveRuntimeConfig**.
5. Click **Apply**.

6. In the **ActionGroup4** box, select **softReset** .
7. Click **Apply**.

Configuring BPDU Guard

Configure BPDU Guard to block the root selection process or to prevent BPDU flooding from unknown devices.

About this task

To configure multiple ports simultaneously, select more than one port in the Device Physical View tab. The **BPDU Guard** tab appears as a table-based tab. For more information about how to use a table-based tab, see *Using CLI and EDM*.

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 **General**.
4. Click the **Interface** tab.
5. Select **BpduGuardAdminEnabled** to enable BPDU Guard for the port.
6. **(Optional)** Type a value in **BpduGuardTimeout** to configure the timer for port-state recovery
7. Click **Apply**.

Interface field descriptions

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

Name	Description
Index	Displays the index of the port, written in the slot/port[/sub-port] format.
Name	Configures the name of the port.
Descr	Displays the description of the port. A textual string containing information about the interface.
Type	Displays the type of connector plugged in the port.
Mtu	Displays the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram which can be sent or received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.
PhysAddress	Displays the physical address of the port. The address of the interface at the protocol layer

Table continues...


Name	Description
	immediately `below' the network layer in the protocol stack. For interfaces which do not have such an address (e.g., a serial line), this object should contain an octet string of zero length.
VendorDescr	Displays the vendor of the connector plugged in the port.
DisplayFormat	Identifies the slot and port numbers (slot/port). If the port is channelized, the format also includes the sub-port in the format slot/port/sub-port
AdminStatus	Configures the port as enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
OperStatus	Displays the current status of the port. The status includes enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
LicenseControlStatus	Shows the port license status. This field only applies to VSP 7200 Series.
ShutdownReason	Indicates the reason for a port state change.
LastChange	Displays the timestamp of the last change.
LinkTrap	Enable or disable link trapping.
AutoNegotiate	Enables or disables auto-negotiation for this port. The default is enabled for VSP 4000 Series, VSP 8000 Series, and VSP 8600 (for all ports except 10G SFP+ ports) but disabled for VSP 7200 Series.
AutoNegAd	Specifies the port speed and duplex abilities to be advertised during link negotiation.  Note: The 8424XT ESM does not support the following speeds: 10-full, 10-half, and 1000-half. The abilities specified in this object are only used when auto-negotiation is enabled on the port. If all bits in this object are disabled, and auto-negotiation is enabled on the port, then the physical link process on the port will be disabled (if hardware supports this ability). Any change in the value of this bit map will force the PHY to restart the auto-negotiation process. This will have the same effect as physically unplugging and reattaching the cable plant attached to the port.

Table continues...

Name	Description
	<p>The capabilities being advertised are either all the capabilities supported by the hardware or the user-configured capabilities, which is a subset of all the capability supported by hardware.</p> <p>The default for this object will be all of the capabilities supported by the hardware.</p>
AdminDuplex	<p>Configures the administrative duplex setting for the port.</p> <p>The switch does not support half duplex.</p>
OperDuplex	<p>Indicates the operational duplex setting for the port.</p> <p>The switch does not support half duplex.</p>
AdminSpeed	Configures the administrative speed for the port.
OperSpeed	Indicates the operational speed for the port.
QoSLevel	Selects the Quality of Service (QoS) level for this port. The default is level1.
DiffServ	Enables the Differentiated Service feature for this port. The default is disabled.
Layer3Trust	Configures if the system should trust Layer 3 packets coming from access links or core links only. The default is core.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled (selected) or disabled (cleared) on the port. The default is disabled (clear).
MltId	Shows the MLT ID associated with this port. The default is 0.
Locked	Shows if the port is locked. The default is disabled.
UnknownMacDiscard	Discards packets that have an unknown source MAC address, and prevents other ports from sending packets with that same MAC address as the destination MAC address. The default is disabled.
DirectBroadcastEnable	Specifies if this interface forwards direct broadcast traffic.
OperRouting	Shows the routing status of the port.
HighSecureEnable	Enables or disables the high secure feature for this port.
RmonEnable	Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.
FlexUniEnable	Enables Flex UNI on the port. The default is disabled.

Table continues...

Name	Description
IngressRateLimit	Limits the traffic rate accepted by the specified ingress port.
IngressRatePeak	Configures the peak rate in Kbps. The default is 0.
IngressRateSvc	Configures the service rate in Kbps. The default is 0.
EgressRateLimitState	Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the port. The default is disabled.
EgressRateLimit	Configures the egress rate limit in Kbps. Different hardware platforms provide different port speeds. The software supports the following ranges: <ul style="list-style-type: none"> • 10 Gbps ports — 1000 to 10000000 • 40 Gbps ports — 1000 to 40000000 • 100 Gbps ports — 1000 to 100000000 If you configure this value to 0, shaping is disabled on the port.
TxFlowControl	Configures if the port sends pause frames. By default, an interface does not send pause frames. You must also enable the flow control feature globally before an interface can send pause frames.
TxFlowControlOperState	Shows the operational state of flow control.
BpduGuardTimerCount	Shows the time, starting at 0, since the port became disabled. When the BpduGuardTimerCount reaches the BpduGuardTimeout value, the port is enabled. Displays in 1/100 seconds.
BpduGuardTimeout	Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires. You can configure a value of 0 or to 65535. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.
BpduGuardAdminEnabled	Enables BPDU Guard on the port. The default is disabled.
Action	Performs one of the following actions on the port <ul style="list-style-type: none"> • none - none of the following actions • flushMacFdb - flush the MAC forwarding table • flushArp - flush the ARP table • flushIp - flush the IP route table • flushAll - flush all tables

Table continues...

Name	Description
	<ul style="list-style-type: none"> triggerRipUpdate — manually triggers a RIP update <p>The default is none.</p>
Result	Displays result of the selected action. The default is none.
IsPortShared	<p>Indicates whether the port is combo or not.</p> <ul style="list-style-type: none"> portShared—Combo port. portNotShared—Not a combo port.
PortActiveComponent	<p>Specifies whether the copper port is active or fabric port is active if port is a combo port.</p> <ul style="list-style-type: none"> fixed port—Copper port is active. gbic port—Fabric port is active.

Configuring RSTP global parameters

Perform this procedure to configure the RSTP global parameters.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **RSTP**.
3. Configure the parameters as required.
4. Click **Apply**.

Globals field descriptions

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

Name	Description
PathCostDefault	Specifies the version of the spanning tree default path costs that are used by this bridge. A value of 8021d1998 indicates the use of the 16-bit default path costs from IEEE Std. 802.1d-1998. A value of stp8021t2001 indicates the use of the 32-bit default path costs from IEEE Std. 802.1t.
TxHoldCount	Specifies the value used by the port transmit state machine to limit the maximum transmission rate. The default is 3.
Version	Specifies the version of STP that the bridge currently runs. The value stpCompatible indicates that the

Table continues...

Name	Description
	Spanning Tree Protocol as specified in IEEE 802.1d is in use; rstp indicates that the Rapid Spanning Tree Protocol as specified in IEEE 802.1w is in use.
EnableStp	Indicates whether the spanning tree protocol is active in this STG. The default is enabled.
Priority	Specifies the RSTP bridge priority.
BridgeMaxAge	Specifies the value that all bridges use for MaxAge while this bridge acts as the root.
BridgeHelloTime	The value that all bridges use for HelloTime while this bridge acts as the root.
BridgeForwardDelay	Specifies the value that all bridges use for forward delay while this bridge acts as the root.
DesignatedRoot	Specifies the unique bridge identifier of the bridge recorded as the root in the configuration BPDUs transmitted by the designated bridge for the segment to which the port is attached.
RootCost	Specifies the cost of the path to the root from this bridge.
RootPort	Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge.
MaxAge	Specifies the maximum age of Spanning Tree Protocol information in hundredths of a second learned from the network on any port before the port is discarded.
HelloTime	Specifies the amount of time in hundredths of a second between the transmission of configuration bridge PDUs by this node on any port while it is the root of the spanning tree (or trying to become the root).
ForwardDelay	Specifies a time value, measured in hundredths of a second, controls how fast a port changes its spanning state after moving towards the forwarding state. The value determines how long the port stays in each of the listening and learning states, which precede the forwarding state. This value is also used after a topology change is detected, and is underway, to age all dynamic entries in the forwarding database.
RstpUpCount	Specifies the number of times the RSTP module is enabled. A trap is generated on the occurrence of this event.

Table continues...

Name	Description
RstpDownCount	Specifies the number of times the RSTP module is disabled. A trap is generated on the occurrence of this event.
NewRootIdCount	Specifies the number of times this bridge detects a root identifier change. A trap is generated on the occurrence of this event.
TimeSinceTopology Change	Specifies the time (in hundredths of a second) since the TcWhile Timer for any port in this bridge was nonzero for Common Spanning Tree.
TopChanges	Specifies the number of times that there was at least one nonzero TcWhile Timer on this bridge for Common Spanning Tree.

Configuring RSTP ports

Configure RSTP to reduce the recovery time after a network breakdown.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **RSTP**.
3. Click the **RSTP Ports** tab.
4. Use the fields in the **RSTP Ports** tab to configure the RSTP ports.
5. Click **Apply**.

RSTP Ports field descriptions

Use the data in the following table to use the **RSTP Ports** tab.

Name	Description
Port	Specifies a unique value, greater than zero, indicating the port number.
Priority	Specifies the value of the priority field.
PathCost	Specifies the contribution of this port to the path cost of paths towards the root that includes this port.
ProtocolMigration	Specifies a port to transmit RSTP BPDUs if operating in RSTP mode. Any other operation on this object has no effect, and RSTP mode returns false if read.
AdminEdgePort	Specifies the administrative value of the Edge Port parameter. A value of true indicates that this port is

Table continues...

Name	Description
	an edge-port, and a value of false indicates that this port is a nonedge-port.
OperEdgePort	Specifies the operational value of the Edge Port parameter. The object is initialized to the value of AdminEdgePort and is configured to false on reception of a BPDU.
AdminPointToPoint	Specifies the administrative point-to-point status of the LAN segment attached to this port. A value of forceTrue indicates that this port is treated as if it is connected to a point-to-point link. A value of forceFalse indicates that this port is treated as having a shared media connection. A value of auto indicates that this port is considered to have a point-to-point link if it is an aggregator and all of its members are aggregatable, or if the MAC entity is configured for full-duplex operation, either through autonegotiation or by management means.
OperPointToPoint	Specifies the operational point-to-point status of the LAN segment attached to this port. It indicates whether a port is considered to have a point-to-point connection or not. The value is determined by management or by autodetection as described in the AdminPointToPoint object.
OperVersion	Indicates if the port is in MSTP mode, RSTP mode or STP-compatible mode. MSTP mode transmits MST BDUs, RSTP mode transmits RST BPDUs and STP-compatible transmits Config/TCN BPDUs.

Viewing RSTP port status

View the RSTP port status to ensure proper functioning of RSTP.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **RSTP**.
3. In the RSTP tab, click the **RSTP Status** tab.

RSTP Status field descriptions

Use the data in the following table to use the **RSTP Status** tab.

Name	Description
Port	Specifies a unique value, greater than zero, indicating the port number.
State	Specifies the current state of the port as defined by application of the Spanning Tree Protocol. This state controls what action a port takes on reception of a frame.
Role	Indicates the current port role assumed by this port.
OperVersion	Indicates whether the port is operationally in the RSTP- or STP-compatible mode; that is, whether the port transmits RSTP BPDUs or Config/TCN BPDUs.
EffectivePortState	Specifies the effective operational state of the port. This object is configured to true if the port is operationally up in the Interface Manager, and if Force Port State for this port and the specified port state is enabled. Otherwise, this object is configured to false.

Configuring MSTP global parameters

Configure the global MSTP parameters to determine how MSTP operates for the system. Interface-level parameters override global settings.

Before you begin

- The system must be in MSTP mode.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **MSTP**.
3. Click the **Globals** tab.
4. Configure MSTP as required.
5. Click **Apply**.

Globals field descriptions

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

Name	Description
PathCostDefaultType	Specifies the version of the spanning tree default path costs to be used by this bridge. A value of 8021d1998 denotes the use of the 16-bit default path costs from IEEE 802.1d-1998. A value of

Table continues...

Name	Description
	stp8021t2001 denotes the use of the 32-bit default path costs from IEEE 802.1t.
TxHoldCount	Specifies the value used by the port transmit state to limit the maximum transmission rate. The default is 3.
MaxHopCount	<p>Assigns the MSTP bridge maximum hop count. The range is 100 to 4000 one hundredths of a second.</p> <p>The original MIB erroneously designated the value in hundredths of a second, when it should have been in hops. The replacement MIB kept the range at 100-4000 to remain backwards compatible. To convert this value to hops, divide by 100 so 100-4000 equals 1-40 hops.</p>
NoOfInstancesSupported	Indicates the maximum number of spanning tree instances supported.
MstpUpCount	The number of times the MSTP module is enabled. A trap is generated on the occurrence of this event.
MstpDownCount	The number of times the MSTP module is disabled. A trap is generated on the occurrence of this event.
ForceProtocolVersion	<p>Specifies the version of Spanning Tree Protocol that the bridge currently runs. stpCompatible indicates that the Spanning Tree Protocol as specified in IEEE 802.1d is in use; rstp indicates that the Rapid Spanning Tree Protocol as specified in IEEE 802.1w is in use; and mstp indicates that the multiple spanning tree protocol as specified in IEEE 802.1s is in use.</p> <p>Although STP and MSTP are variations of the same spanning tree protocol, they communicate information differently. A switch in MSTP mode cannot recognize the spanning tree groups running on a chassis configured with STP. MSTP spanning tree groups are not the same as STP spanning tree groups. Using a switch in MSTP mode with a chassis in STP mode can create a loop in the network.</p> <p>The default is MSTP.</p>
BrgAddress	Specifies the MAC address used by this bridge if it must be referred to in a unique fashion. It is recommended that this should be the numerically smallest MAC address of all ports that belong to this bridge. If concatenated with MstCistBridgePriority or MstBridgePriority, a unique bridge identifier is formed, which is used in the STP.

Table continues...

Name	Description
Root	Specifies the bridge identifier of the root of the common spanning tree as determined by the STP by this node. This value is used as the CIST root identifier parameter in all configuration bridge PDUs originated by this node.
RegionalRoot	Specifies the bridge identifier of the root of the multiple spanning tree region as determined by the STP as executed of this node. This value is used as the common and internal spanning tree (CIST) regional root identifier parameter in all configuration bridge PDUs originated by this node.
RootCost	Specifies the cost of the path to the CIST root from this bridge.
RegionalRootCost	Specifies the cost of the path to the CIST regional root from this bridge.
RootPort	Specifies the port number of the port which offers the lowest path cost from this bridge to the CIST root bridge.
BridgePriority	Specifies the value of the writable portion of the bridge identifier comprising the first two octets. The values you enter for bridge priority must be in steps of 4096. The default is 32768.
BridgeMaxAge	Specifies the value that all bridges use for MaxAge while this bridge acts as the root. The granularity of this timer is specified as 1 second. An agent can return a bad value error if you attempt to configure a value which is not a whole number of seconds. The default is 2000.
BridgeForwardDelay	Specifies the value that all bridges use for forward delay if this bridge acts as the root. Note that 802.1d specifies that the range for this parameter is related to the value of BridgeMaxAge. The granularity of this timer is specified as 1 second. An agent can return a bad value error if you attempt to configure a value which is not a whole number of seconds. The default is 1500.
HoldTime	Determines the interval length in hundredths of a second during which no more than two configuration bridge PDUs can be transmitted by this node.
MaxAge	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the value that this bridge currently uses.

Table continues...

Name	Description
ForwardDelay	Specifies the time value, measured in units of hundredths of a second, that controls how fast a port changes its spanning state after moving towards the forwarding state. This value determines how long the port stays in a particular state before moving to the next state.
TimeSinceTopology Change	Specifies the time (in hundredths of a second) since the TcWhile Timer for any port in this bridge was nonzero for Common Spanning Tree.
TopChanges	Specifies the number of times that there was at least one nonzero TcWhile Timer on this bridge for Common Spanning Tree.
NewRootBridgeCount	Specifies the number of times this bridge detects a root bridge change for Common Spanning Tree. A trap is generated on the occurrence of this event.
RegionName	Specifies the name for the region configuration. By default, the region name is equal to the bridge MAC Address.
RegionVersion	Specifies the version of the MST region.
ConfigIdSel	Specifies the configuration identifier format selector used by the bridge. This has a fixed value of 0 to indicate RegionName. RegionVersions are specified as in the standard.
ConfigDigest	Specifies the configured MD5 digest value for this region, which must be 16 octets long.
RegionConfigChange Count	Specifies the number of times a region configuration identifier change is detected. A trap is generated on the occurrence of this event.

Configuring CIST ports for MSTP

Configure Common and Internal Spanning Tree (CIST) ports to configure ports for MSTP.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **MSTP**.
3. Click the **CIST Port** tab.

Important:

The MSTP, CIST Port tab contains information for each port that is common to all bridge and spanning tree instances.

4. Use the fields in the **CIST Port** box to configure the MSTP CIST port.
5. Click **Apply**.

CIST Port field descriptions

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

Name	Description
Port	Specifies the port number of the port for which this entry contains spanning tree information.
PathCost	Specifies the contribution of this port to the path cost of paths towards the CIST root that includes this port.
Priority	<p>Specifies the four most significant bits of the port identifier of the spanning tree instance which are modified by setting the CistPortPriority value. The values that are configured for port priority must be in steps of 16.</p> <p>Although port priority values can range from 0 to 255, only the following values are used: 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, 240.</p> <p>The default is 128.</p>
DesignatedRoot	Specifies the unique bridge identifier of the bridge recorded as the CIST root in the configuration BPDUs transmitted.
DesignatedCost	Specifies the path cost of the designated port of the segment that connects to this port.
DesignatedBridge	Specifies the unique bridge identifier of the bridge which that port considers to be the designated bridge for the ports segment.
DesignatedPort	Specifies the port identifier of the port on the designated bridge for this port segment.
RegionalRoot	Specifies the unique bridge identifier of the bridge recorded as the CIST regional root identifier in the configuration BPDUs transmitted.
RegionalPathCost	Specifies the contribution of this port to the path cost of paths towards the CIST regional root that include this port.
ProtocolMigration	<p>Indicates the protocol migration state of this port. If you chose true, the option initiates protocol migration for a port. If you chose false, the option terminates protocol migration for a port.</p> <p>An MSTP-compatible port transmits and receives only RSTP BPDUs. If an MSTP port receives an STP BPDU, it becomes an STP port. User intervention is</p>

Table continues...

Name	Description
	<p>required to change this port back to MSTP mode. This process is called Port Protocol Migration.</p> <p>You must configure protocol migration to true on all spanning-tree enabled interfaces when you change the spanning tree version from STP-compatible to MSTP for those interfaces to work in the proper mode.</p>
AdminEdgeStatus	Specifies the administrative value of the Edge Port parameter. A value of true indicates that this port is an edge-port, and a value of false indicates that this port is a nonedge-port.
OperEdgeStatus	Specifies the operational value of the Edge Port parameter. The object is initialized to the value of AdminEdgeStatus and is configured to false on reception of a BPDU.
AdminP2P	Specifies the administrative point-to-point status of the LAN segment attached to this port. A value of forceTrue indicates that this port is treated as if it connects to a point-to-point link. A value of forceFalse indicates that this port is treated as having a shared media connection. A value of auto indicates that this port is considered to have a point-to-point link if it is an aggregator and all of its members are aggregatable, or if the MAC entity is configured for full-duplex operation, either through autonegotiation or by management means.
OperP2P	Specifies the operational point-to-point status of the LAN segment attached to this port. It indicates whether a port is considered to have a point-to-point connection or not. The value is determined by management or by autodetection as described in the AdminP2P object.
HelloTime	Specifies the amount of time in hundredths of a second between the transmission of configuration bridge PDUs by this node on this port.
OperVersion	<p>Indicates whether the port is operationally in the MSTP mode, the RSTP mode, or the STP-compatible mode; that is, whether the port transmits MST BPDUs, RST BPDUs, or Config/TCN BPDUs.</p> <p>Although STP and MSTP are variations of the same spanning tree protocol, they communicate information differently. A switch in MSTI mode cannot recognize the spanning tree groups running on a chassis configured with STP. MSTP spanning tree groups are not the same as STP spanning tree</p>

Table continues...

Name	Description
	groups. Using a switch in MSTP mode with another chassis in STP mode can create a loop in the network.
EffectivePortState	Specifies the effective operational state of the port for CIST. This is true only if the port is operationally up at the interface and protocol levels for CIST. This is configured to false for all other conditions.
State	Specifies the current state of the port as defined by the common spanning tree protocol. It can be disabled, discarding, learning, or forwarding.
ForcePortState	Specifies the current state of the port. You can change the port to either Disabled or Enabled for the base spanning tree instance.
SelectedPortRole	Specifies the selected port role of the port for this spanning tree instance.
CurrentPortRole	Specifies the current port role of the port for this spanning tree instance.

Configuring MSTI bridges for MSTP

Perform this procedure to configure multiple spanning tree instance (MSTI) bridges for MSTP.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **MSTP**.
3. Click the **MSTI Bridges** tab.

! **Important:**

The systems generates MSTI bridge instances after you create a VLAN in MSTP mode.

4. Use the fields in the **MSTI Bridges** box to configure the MSTP bridge.
5. Click **Apply**.

MSTI Bridges field descriptions

Use the data in the following table to use the **MSTI Bridges** tab.

Name	Description
Instance	Specifies the spanning tree instance to which this information belongs.

Table continues...

Name	Description
RegionalRoot	Specifies the MSTI regional root identifier value for the instance. This value is used as the MSTI regional root identifier parameter in all configuration bridge PDUs originated by this node.
Priority	Specifies the writable portion of the MSTI bridge identifier comprising the first two octets. The values that are configured for bridge priority must be in steps of 4096. The default is 32768.
RootCost	Specifies the cost of the path to the MSTI regional root as seen by this bridge.
RootPort	Specifies the port number of the port that offers the lowest path cost from this bridge to the MSTI region root bridge.
TimeSinceTopologyChange	Specifies the time (in hundredths of a second) since the TcWhile Timer for any port in this bridge was nonzero for this spanning tree instance.
TopChanges	Specifies the number of times that there was at least one nonzero TcWhile Timer on this bridge for this spanning tree instance.
NewRootCount	Specifies the number of times this bridge detects a root bridge change for this spanning tree instance. A trap is generated on the occurrence of this event.
InstanceUpCount	Specifies the number of times a new spanning tree instance is created. A trap is generated on the occurrence of this event.
InstanceDownCount	Specifies the number of times a spanning tree instance is deleted. A trap is generated on the occurrence of this event.

Configuring MSTI ports for MSTP

Perform the following procedure to configure MSTI ports for MSTP.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > VLAN > Spanning Tree**.
2. Click **MSTP**.
3. Click the **MSTI Port** tab.

 **Important:**

Port members you select on the VLAN, **Basic** tab appear in the **MSTI Port** tab.

4. Use the fields in the **MSTI Port** box to configure the MSTP.

5. Click **Apply**.

MSTI Port field descriptions

Use the data in the following procedure to use the **MSTI Port** tab.

Name	Description
Port	Specifies the port number of the port for which this entry contains spanning tree information.
Instance	Specifies the spanning tree instance to which the information belongs.
PathCost	Specifies the contribution of this port to the path cost of paths towards the MSTI root that includes this port.
Priority	Specifies the four most significant bits of the port identifier for a given spanning tree instance can be modified independently for each spanning tree instance supported by the bridge. The values configured for port priority must be in steps of 16. The default is 128.
DesignatedRoot	Specifies the unique bridge identifier of the bridge recorded as the MSTI regional root in the configuration BPDUs transmitted.
DesignatedBridge	Specifies the unique bridge identifier of the bridge that this port considers to be the designated bridge for the port segment.
DesignatedPort	Specifies the port identifier of the port on the designated bridge for this port segment.
State	Specifies the current state of the port, as defined by the MSTP. A port which is in forwarding state in one instance can be in discarding (blocking) state in another instance.
ForcePortState	Specifies the current state of the port, that is changed to either disabled or enabled for the specific spanning tree instance.
DesignatedCost	Specifies the path cost of the designated port of the segment connected to this port.
CurrentPortRole	Specifies the current port role of the port for this spanning tree instance.
EffectivePortState	Specifies the effective operational state of the port for a specific instance. This is configured to true if the port is operationally up at the interface and protocol levels for the specific instance. This is configured to false at all other times.

Chapter 5: VXLAN Gateway configuration

This chapter provides concepts and procedures to configure the VXLAN Gateway features supported on the switch.

VXLAN Gateway fundamentals

VXLAN Gateway terminates virtual extensible LAN (VXLAN) tunnels that “stretch” emulated Layer 2 segments over an IP network.

VXLAN is an encapsulation protocol for running an overlay network over an existing Layer 3 infrastructure. This tunneling scheme uses a VLAN-like encapsulation mechanism to encapsulate MAC-based traffic. The VXLAN overlays separate workloads from physical networks, which facilitates the movement of VMs in a multitenant environment.

For more information about VXLAN, see the Internet Engineering Task Force (IETF) standard RFC 7348.


*** Note:**

VXLAN Gateway is supported on the VSP 7200 and VSP 8000 Series switches.

VXLAN Gateway as a VTEP

The VXLAN Gateway feature terminates VXLAN tunnels and is known as a VXLAN tunnel endpoint (VTEP). The VXLAN Gateway is a *hardware-based* VTEP that allows VXLAN to seamlessly communicate with VLANs, other VXLANs, as well as Fabric Connect I-SIDs.

The following figure shows a VXLAN Gateway (labeled as VTEP Source) connecting VXLAN segments to an SPB network. This type of deployment provides a solution for customers using VXLAN in their data centers to interoperate with Fabric Connect.



VXLAN segments

VXLAN tunnels are called VXLAN segments. They provide the ability to separate, abstract, and decouple the physical topology from a *logical* or *virtual* topology by using encapsulated tunneling. This means that only VMs within the same VXLAN segment can communicate with each other.

VNID

Each VXLAN segment has a 24-bit segment ID called a VXLAN Network Identifier (VNID). This VNID mechanism allows up to *16 million* VXLAN segments to coexist within the same administrative domain, which enables VMs to migrate easily between servers in the same data center or between distributed data centers. Each VTEP can support multiple VNIDs. For information about maximum scaling numbers, see *Release Notes*.

The VNID uses the inner MAC frame originated by the individual VM. Thus, you could have overlapping MAC addresses across segments but never have traffic "crossover" because the traffic is isolated using the VNID qualifier. Due to this encapsulation, only the VTEP knows the VNID and its associated VXLAN tunnel. VMs are unaware of VXLAN.

* Note:

VXLAN Gateway uses Head End replication to flood broadcast, unknown destination address, and multicast traffic packets to all VTEPs in a VNID.

VM traffic

Only VMs within the same VXLAN segment can communicate with each other. The following steps show how a VM communicates with a VM on a different host within the same VXLAN segment:

1. The VM sends a MAC frame destined to the target.
2. The VTEP looks up the VNID that this VM is associated with to verify that the destination MAC is on the same segment. If the destination MAC is on the same segment, the process continues.

3. The VTEP inserts an outer header comprising an outer MAC, outer IP address, and VXLAN header in front of the original MAC frame.
4. The VTEP transmits the final packet out to the destination, which is the IP address of the remote VTEP connecting the destination VM addressed by the inner MAC destination address.
5. The remote VTEP verifies that the VNID is a valid one and is used by the destination VM.
6. If the VNID is valid, the remote VTEP strips the packet of its outer header and passes it on to the destination VM. The destination VM never knows about the VNID or that the frame was transported with a VXLAN encapsulation.

VXLAN Gateway licensing

VXLAN Gateway requires a Premier or Premier with MACsec license. If a Premier License is not present, you cannot configure VXLAN Gateway.

Important:

The VSP 7200 and the VSP 8000 Series support PLDS licenses **only** and share the same order codes.

For more information about licensing, see *Administering*.

VXLAN modes

The VXLAN Gateway implementation is available in two modes: Base Interworking Mode and Full Interworking Mode. By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes through a boot configuration flag.

Important:

Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.

Base Interworking Mode

Base Interworking Mode is the default mode. In this mode, VXLAN Gateway supports Layer 2 gateway communication at line rate between VXLAN and traditional VLAN environments.

Base Interworking Mode has the following limitations:

- No support for VXLAN-to-VXLAN communication.
- No support for VXLAN-to-SPB communication.
- No support for SMLT, vIST, or Simplified vIST.

Each VNID represents a service instance for the VXLAN Gateway. To enable a VNID to extend into the SPB network, you must use Full Interworking Mode.

*** Note:**

If your configuration does not have a VXLAN Gateway but has SMLT or vIST, you can enable VXLAN Gateway in the Full Interworking Mode *only*. SMLT and vIST are incompatible and mutually exclusive with the Base Interworking Mode.

Full Interworking Mode

Full Interworking Mode is **not** the default mode. To configure Full Interworking Mode, you must enable the `vxlan-gw-full-interworking-mode` boot flag, save the configuration, and then reboot the switch.

Full Interworking Mode supports the Base Interworking Mode communication between VXLAN and traditional VLAN environments as well as VXLAN-to-VXLAN communication and all SPB functionality including vIST and SMLT. Layer 2 and Layer 3 services extend across VXLAN, VLAN, and SPB domains using an internal loopback. In Full Interworking Mode, VXLAN traffic flows at less than full line rate and the total amount of available internal loopback bandwidth depends on the platform used.

Full Interworking Mode has the following limitations:

- No support for Simplified vIST.
- A VNID cannot have the same value as an I-SID. A consistency check ensures that this does not happen if you change the boot configuration flag to turn off Full Interworking Mode and go into Base Interworking Mode.

Types of VXLAN Gateway deployments

VXLAN Gateway provides the following deployment solutions:

- [VXLAN-to-VLAN Layer 2 gateway deployment](#) on page 144
- [VXLAN-to-VXLAN Layer 3 gateway deployment](#) on page 145
- [VXLAN-to-SPB gateway deployment](#) on page 146

These gateways provide VXLAN tunnel termination functions for switches that could be Top of Rack (ToR) switches, access switches, or switches higher up in the data center network topology such as a core switch or even WAN edge devices. The last case (WAN edge) could involve a Provider Edge (PE) router that terminates VXLAN tunnels in a hybrid cloud environment.

VXLAN-to-VLAN Layer 2 gateway deployment

The VXLAN-to-VLAN deployment is where the VXLAN Gateway communicates between VXLAN segments and VLAN segments. Use this scenario when nodes on a VXLAN overlay network need to communicate with nodes on legacy networks that are VLAN based as shown in the following figure.




Figure 9: VXLAN-to-VLAN deployment

In this type of deployment, the VXLAN Gateway maps the VXLAN VNID to a VLAN in the following way:

- When packets traverse a VXLAN segment into the VLAN segment, the VXLAN VTEP removes the VXLAN header and forwards regular VLAN packets to a physical port based on the destination MAC address of the inner Ethernet frame.

The VTEP discards decapsulated frames with the inner VLAN ID unless you explicitly configure them to be passed on to the non-VXLAN interface. The VXLAN VNID to VLAN mapping is statically configured.

- In the reverse direction when packets traverse a VLAN segment into a VXLAN segment, the VXLAN VTEP maps the incoming frames from the non-VXLAN interfaces to a specific VXLAN overlay network based on the VLAN ID in the frame. This complies with the standard VXLAN RFC.

The VTEP also accepts both untagged and tagged inner frames and adds or replaces the tag of the inner frame with the local VLAN mapped to the VNID before sending out the frame.

The VTEP removes the VLAN ID of the original frame before it encapsulates the frame for VXLAN.

VXLAN-to-VXLAN Layer 3 gateway deployment

The VXLAN-to-VXLAN deployment is where the VXLAN Gateway communicates between VXLAN segments. Use this scenario when one VXLAN segment needs to communicate with another VXLAN segment.

When a packet traverses a VXLAN segment and enters the VXLAN Gateway, the VXLAN VTEP maps one VXLAN segment to another in the following way:

- Decapsulates the VXLAN header.
- Performs a route lookup of the customer packet.
- Identifies that the next hop VLAN is in another VXLAN segment.

- Encapsulates the packet with a new VXLAN header with a new VNID.
- Forwards the packet out.

The following figure shows two virtualized servers attached to a Layer 3 infrastructure. The servers could be on the same rack, on different racks, or across data centers within the same administrative domain.

There are four VXLAN overlay networks identified by the VNIDs: 22, 34, 74, and 98. Notice how the VMs are associated with a VNID. For example, VM1-1 in Server 1 and VM2-4 on Server 2 are on the same VXLAN overlay network identified by VNID 22.

The VMs do not know about the overlay networks and transport method because the encapsulation and decapsulation happen transparently at the VTEPs on Servers 1 and 2. The other overlay networks and the corresponding VMs are: VM1-2 on Server 1 and VM2-1 on Server 2 both on VNID 34, VM1-3 on Server 1 and VM2-2 on Server 2 on VNID 74, and finally VM1-4 on Server 1 and VM2-3 on Server 2 on VNID 98.




Figure 10: VXLAN-to-VXLAN deployment

VXLAN-to-SPB gateway deployment

The VXLAN-to-SPB deployment is where the VXLAN Gateway communicates between a VXLAN segment and an SPB domain. This type of deployment provides a solution for customers using VXLAN in their data-centers to interoperate with Fabric Connect as shown in the following figure.




Figure 11: VXLAN-to-SPB deployment

The VXLAN Gateway maps the VXLAN VNID to I-SID in the following way:

- When packets traverse a VXLAN segment into an SPB domain, the VXLAN VTEP decapsulates the VXLAN packet, encapsulates it with MAC-in-MAC encapsulation, and then forwards it into the SPB domain.
- In the reverse direction, the VXLAN VTEP decapsulates the MAC-in-MAC packet from an SPB domain, encapsulates it with a VXLAN header, and then forwards it into a VXLAN segment.

VXLAN Gateway considerations and limitations

Review the following considerations, limitations, and behavioral characteristics associated with VXLAN Gateway.

Unsupported features

VXLAN Gateway does not support the following features:

- IP Multicast over Fabric Connect traffic cannot be carried over into a VXLAN domain through a VXLAN Gateway.
- SPB-PIM Gateway and VXLAN Gateway interoperability is not supported.
- VXLAN Gateway does not support Simplified vIST in any mode.

QoS and VXLAN Gateway

- If present in the packet, VXLAN Gateway honors user-configured ingress dot1p and DSCP mappings to derive the internal CoS.
- Customer packets are not remarked.
- Control packets are handled by high priority CoS queues.
- DSCP bits in the outer IP header of VXLAN-encapsulated packets are always derived from the internal QoS, irrespective of the ingress port DiffServ configuration. Customer packet IP DSCP bits are not modified as part of VXLAN encapsulation. For more information about QoS, see *Configuring QoS and ACL-Based Traffic Filtering*.

- DSCP bits in the inner IP header of VXLAN-encapsulated packets change when traffic comes in on a Layer 2 Trusted or a Layer 3 Untrusted port with DiffServ enabled.

ECMP support for VXLAN Gateway and Fabric Extend

VXLAN Gateway requires ECMP support to communicate with remote VTEPs. The software extended this ECMP support to Fabric Extend Layer 3 core tunnels. Therefore, if your switch supports VXLAN Gateway, it also supports ECMP for both VXLAN Gateway and Fabric Extend.

Base Interworking Mode considerations and limitations

- After associating a VNID with an I-SID (`vnid <vnid value> i-sid <isid value>`), you cannot create an I-SID with the same value as the `<vnid value>` and vice versa.
- There is no support for SMLT, vIST, or Simplified vIST in Base Interworking Mode.
- Neither Base Interworking Mode nor Full Interworking Mode support Fabric Attach endpoints.
- Base Interworking Mode does not allow SMLT and vIST configurations. If you change from Full Interworking Mode to Base Interworking Mode, make sure there are none of these configurations.
- In Base Interworking Mode, the `untagged-traffic <port >< mlt>` support is limited to untagged traffic forwarding only. No control packet forwarding is supported on this port or MLT. LACP MLT is not supported as part of `untagged-traffic <mlt>`.

Full Interworking Mode considerations and limitations

- Whenever you change from Base Interworking Mode to Full Interworking Mode, you have to reboot the switch. The software automatically changes the configuration to be in line with Full Interworking Mode.
- Full Interworking Mode does not have VXLAN endpoints. Configure ELAN endpoints under the ELAN I-SID only.
- Neither Base Interworking Mode nor Full Interworking Mode support Fabric Attach endpoints.

General configuration guidelines

- In VXLAN environments that have an underlying IP network with SMLT deployment, you must enable RSMLT in the underlying network on vIST VTEP devices.
- If you use more than one VXLAN Gateway node for the same I-SID, you must ensure that loops are prevented.
- The VTEP remote destination IP table in the datapath hardware is shared with the Fabric Extend IP core remote tunnel destination IP table.
 - The maximum number of FE tunnel destination IP addresses is 256.
 - The maximum number of VTEP IP addresses is 500.

For every FE tunnel destination IP you configure, you must reduce the number of VTEP IPs by one. For example, if you configure the maximum number of FE tunnel destination IPs then the maximum of VTEP remote destination that you can configure is $500-256=244$.

- The total number of MACs that the VTEP learns in the switch is constant and includes the MACs learned on VNID, ISID, and on VLANs (normal and CVLANs).
 - MAC addresses that the VTEP learns from the VNID table are also learned when an I-SID is associated with the VNID. Therefore, the maximum number of MAC addresses that the VTEP can support is reduced accordingly.
 - If a VLAN is associated with an I-SID, and that I-SID is associated with a VNID, then the VTEP learns each customer MAC address from three different tables (VLAN, I-SID, and VNID). This consumes three MAC entry records. In a vIST scenario where all I-SIDs have associated VNIDs, then the max number of C-MACs that the VTEP learns is reduced by $224K \div 3 = 74K$ MAC addresses.
 - If a Switched-UNI endpoint is associated with an I-SID, then the VTEP learns the C-MAC only once in the I-SID table. If this I-SID is associated with a VNID, then the VTEP learns the C-MAC twice: once in the I-SID table and once in the VXLAN VNID table.
- IP filter rules that you create do not work for VXLAN tunnel terminated packets. This restriction also applies to the Fabric Extend Layer 3 core.
- All ports in an MLT configured as a Switched-UNI endpoint or a VXLAN endpoint have the same properties as configured under the MLT interface.
- A VXLAN tunnel VTEP source IP cannot be a broadcast or multicast IP address. Configure it as a loopback IP address and not a router IP address.
- The VXLAN source address (**vtep-source-ip**) should be different from the SPBM IP Shortcut (**ip-source-address**) and Fabric Extend (**ip-tunnel-source-address**) source addresses.
- If you configure the loopback IP address under a VRF, the **vtep-source-ip** address must be set up under the same VRF.
- The remote VTEP IP address (**vtep <id> ip <ip address>**) cannot be local to the system.
- You can configure a maximum of 500 remote VTEPs per VNID.
- You can change the VTEP source IP and remote VTEP IP dynamically.
- VTEPs have to be reachable for the tunnel to be up.
- The tunnel ID for a VTEP has to be the same on both vIST peers.
- If the VTEP source IP is configured under a VRF, then the VRF cannot have an I-SID associated with it or vice versa.
- Two VNIDs cannot be mapped to the same I-SID.
- An I-SID associated with a VNID cannot be a T-UNI or E-Tree I-SID.
- To prevent routing loops, the platform VLANs used for VXLAN interworking should NOT be in the same VRF or GRT as the IP interfaces in the underlay network.
- You cannot delete a VTEP source IP address (**vtep-source-ip**) when VTEPs are configured.
- If an ELAN I-SID attached to a VNID has a platform VLAN associated with it, then you cannot delete the platform VLAN or change the I-SID associated with that VLAN.
- You cannot delete the loopback IP associated with a VTEP source IP.

- If the VTEP source IP is configured under a VRF, then you cannot delete the VRF.
- You cannot delete or modify the VLAN with I-SID that is associated with a VNID.
- If an ELAN I-SID attached to a VNID also has a platform VLAN associated with it, then you cannot delete the platform VLAN or change the I-SID associated with that VLAN.

VXLAN Gateway configuration using the CLI

This section provides procedures to configure VXLAN Gateway using the CLI.

Configuring the VXLAN Gateway boot flag

There are two configuration modes: Base Interworking Mode and Full Interworking Mode. The Base Interworking Mode is the default mode. To switch to the Full Interworking Mode, use this procedure.

! **Important:**

After you change this boot flag, you must save the change to the configuration file and reboot the switch.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable Full Interworking Mode:

```
boot config flags vxlan-gw-full-interworking-mode
```

3. Save the changed configuration.

4. Restart the switch.

5. Verify that VXLAN Gateway is in Full Interworking Mode:

```
show boot config flags
```

The `true` parameter in the output shows that Full Interworking Mode is enabled.

```
flags vxlan-gw-full-interworking-mode true
```

Configuring VXLAN Gateway

Use the following procedure to configure VXLAN Gateway parameters to allow SPBM to operate on the switch.

Before you begin

Use the `show boot config flags` to verify that you are in the mode you want. The `vxlan-gw-full-interworking-mode` should be set to `false` for the Base Interworking Mode (default) and `true` for the Full Interworking Mode.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the VTEP source IP address:

```
vtep source-ip <A.B.C.D> [vrf WORD<1-16>]
```

3. Configure the remote VTEP IP address:

```
vtep <1-500> ip <A.B.C.D> [name WORD<1-64>]
```

4. Create a VNID instance:

```
vnid <1-16777215> i-sid <1-16777215>
```

*** Note:**

The command prompt changes to **#vxlan** to indicate that you are now in VXLAN Configuration mode for the VNID specified in `vnid <1-16777215>`.

5. Associate VLANs in a port or MLT list to this VNID instance:

```
c-vid <1-4094> port <{slot/port[/sub-port] [-slot/port[/sub-port]]
[,...]}>
```

or

```
c-vid <1-4094> mlt <1-4094>
```

6. Specify the ports or MLT in this VNID instance that will support untagged traffic:

```
untagged-traffic port <{slot/port[/sub-port]}>
```

or

```
untagged-traffic mlt <1-4094>
```

7. Associate VTEPs to this VNID instance:

```
vtep <1-255>
```

Examples

For an example of this procedure, see [VXLAN Gateway configuration example in Base Interworking mode](#) on page 162 or [VXLAN Gateway configuration example in Full Interworking mode](#) on page 162.

Variable definitions

Use the data in the following table to use the `vteb` command.

Variable	Value
source-ip <A.B.C.D> [vrf WORD<1–16>]	Specifies the VXLAN tunnel end point (VTEP) source IP address, which can be on the GRT or a VRF. * Note: The VTEP source IP address must be on a loopback interface.
<1–500> ip <A.B.C.D> [name WORD<1–64>]	Specifies an index value and an IP address that uniquely identifies this remote VTEP. Optionally, you can assign a specific name to this tunnel. By default, the switch assigns a name in this format: VTEP-<#ID>-<IP address> * Note: The remote VTEP IP address cannot be a local, broadcast, or multicast IP address.

Use the data in the following table to use the `vnid` command.

Variable	Value
<1–16777215> i-sid <1–16777215>	Uses this VNID and I-SID information to create a VNID instance and enter VXLAN Configuration Mode in the CLI. * Note: A VNID must not have the same value as an I-SID.

Use the data in the following table to configure commands in the VXLAN Configuration Mode (`#vxlan`). These commands apply to the VNID instance that you specified in the `vnid <1–16777215> i-sid <1–16777215>` command.

Variable	Value
c-vid <1-4094> port <{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}>	Specifies a value that uniquely identifies the customer VLAN ID and ports of this ELAN end point. The switch reserves the following VLAN IDs: <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode.

Table continues...

Variable	Value
	<ul style="list-style-type: none"> • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
c-vid <1-4094> mlt <1-4094>	<p>Specifies a value that uniquely identifies the customer VLAN ID and MLTs of this ELAN end point.</p> <p>The switch reserves the following VLAN IDs:</p> <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
untagged-traffic port <{slot/port[/sub-port] [-slot/port[/sub-port]] [...]}>	Specifies the ports that support untagged traffic.
untagged-traffic mlt <1-4094>	Specifies the MLTs that support untagged traffic.
vtep <1-255>	Lists the remote VTEP destinations to associate with the specified VNID.

Flushing the MAC forwarding table

Use this procedure to flush all the learned MAC addresses from the forwarding database of the selected VNID.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Flush the MAC FDB:

```
vnid mac-address-entry <1-16777215> flush
```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# vnid mac-address-entry 100 flush
```

VXLAN Gateway show commands

Use the procedures in this section to display specific information about VXLAN Gateway on the switch.

Displaying VTEP source information

Use the following procedure to display information about the VTEP source.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the VTEP's source IP address and the name of the VRF:
show vtep local

Example

```
Switch:1# show vtep local
```

```
=====
                          VTEP General Info
=====
vtep source-ip-address: 198.51.100.1
vtep vrf : GlobalRouter
```

Displaying remote VTEP information

Use the following procedure to display information about the remote VTEP.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about all of the remote VTEPs or use the option to specify a particular VTEP:
show vtep remote [<1-500>]

Example

```
Switch:1# show vtep remote
```

```
=====
                          Remote VTEP Info
=====
ID    NAME          VTEP           L3_VTEP_NEXT_HOP_INFO
  DEST-IP      PORT/MLT      VLAN          VRF
-----
1     vtep-1        198.51.100.6   M1t78        3851         GlobalRouter
56    vtep-56       198.51.100.22  M1t130       3855         underlay-vxlan
-----
2 out of 2 Total Num of remote vteps
=====
```

Displaying the name of the remote VTEP

Use the following procedure to display the name of the remote VTEP.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display the names of all remote VTEPs or use the option to specify a particular VTEP:
show vtep remote name [<vtepId>]

Example

```
Switch:1# show vtep remote name
=====
Remote VTEP Name
=====
ID      NAME
-----
1       vxlan_tunnel_1
2       --
3       --
-----
3 out of 3 Total Num of remote vteps
=====
```

Displaying VNID information

Use the following procedure to display information about the VNID.

Procedure

1. Enter Privileged EXEC mode:
enable
2. Display information about all the VNIDs or use the option to specify a particular VNID:
show vnid i-sid [<1-500>]

Example

```
Switch:1# show vnid i-sid
=====
vnid Info
=====
VNID      ISID      PORT      MLT      VTEP
ID        ID        INTERFACES INTERFACES IDS
-----
110       120       c3:1/10   -        2,56
11000060  1060     c60:1/10  -        56
11000065  1065     c65:1/10  -        56
11000066  1066     c66:1/48  -        2
.
.
.
```

Displaying the MAC addresses in the VNID FDB

Use the following procedure to display the MAC addresses in the VNID forwarding database.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display all the VNIDs in the FDB table or use one of the options to specify a particular VNID:

```
show vnid mac-address-entry [<1-16777215> | port <{slot/port [-slot/
port] [,...]}> | mac <0x00:0x00:0x00:0x00:0x00:0x00> | remote]
```

Example

```
Switch:1# show vnid mac-address-entry
```

```
=====
                          VNID Fdb Table
=====
VNID          STATUS      MAC-ADDRESS          INTERFACE          TYPE
-----
11000065      learned    00:00:5e:00:01:01    vtep_56            REMOTE
14000100      learned    a4:25:1b:51:a5:04    vtep_67            REMOTE
14000100      learned    a4:25:1b:52:a5:04    vtep_67            REMOTE
14000100      learned    a4:25:1b:52:bd:03    vtep_2             REMOTE
```

```
c: customer vid    u: untagged-traffic
```

```
All 4 out of 4 Total Num of vnid FDB Entries displayed
=====
```

Displaying the VXLAN running configuration

Use the following procedure to display the currently running VXLAN configuration.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display the running configuration:

```
show running-config module vxlan
```

Example

This example shows the VXLAN information only. It does not show the hardware-specific information that also appears in the output for this command.

The example shows the VXLAN Gateway in Base Interworking Mode. It also shows an external loopback (see `c-vid` commands), which enables communication between the VNID and an I-SID in an SPB network.

```
Switch:1# show running-config module vxlan
```

```
#
# VTEP CONFIGURATION
#
```

```

vtep source-ip 1.1.1.1 vrf underlay-vxlan

#
# REMOTE VTEP CONFIGURATIONS
#
vtep 2 ip 1.1.1.1 name "vtep-2"
vtep 67 ip 1.1.1.67 name "vtep-67"

#
# VNID CONFIGURATION
#
vni 12002105 i-sid 102105
c-vid 105 port 1/11
c-vid 2105 port 1/10
vtep 67
exit

```

VXLAN Gateway configuration using EDM

This section provides procedures to configure VXLAN Gateway using the Enterprise Device Manager (EDM).

Configuring the VXLAN Gateway boot flag

There are two configuration modes: Base Interworking Mode and Full Interworking Mode. The Base Interworking Mode is the default mode. To switch to the Full Interworking Mode, use this procedure.

Important:

After you change this boot flag, you must save the change to the configuration file and reboot the switch.

Procedure

1. In the navigation pane, expand the **Configuration > Edit > Chassis** folders.
2. Click the **Boot Config** tab.
3. Select **EnableVxlanGwFullInterworkingMode** to enable Full Interworking Mode.
4. Click **Apply**.

Configuring a VTEP source IP address

Use the following procedure to configure the VXLAN Gateway source IP address.


Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.

3. Click the **Globals** tab.
4. In the **Sourcelp** field, enter the VTEP source IP address.
5. In the **Vrf** field, select `GlobalRouter` or a VRF ID.
6. Click **Apply**.

Globals field descriptions

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

Name	Description
Sourcelp	Specifies the VXLAN tunnel end point (VTEP) source IP address in an A.B.C.D format.
Vrf	<p>Specifies the VRF name used for the source IP address, which you select from the drop down menu. This IP address can be on the GRT or a VRF.</p> <p> Note: The VTEP source IP address must be on a loopback interface.</p>

Configuring a remote VTEP

Use the following procedure to configure the remote VTEP destination. The total number of VTEP remote destination IPs that you can configure is 500.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vtep** tab.
4. In the **Id** field, assign a unique index value to this remote VTEP.
5. In the **IpAddr** field, enter the IP address of this remote VTEP.
6. In the **Name** field, assign a name to this remote VTEP.
7. Click **Apply**.

Vtep field descriptions

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

Name	Description
Id	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.

Table continues...

Name	Description
IpAddr	Specifies an IP address that uniquely identifies this remote VTEP. * Note: The remote VTEP IP address cannot be a local, broadcast, or multicast IP address.
Name	Assigns a name to this remote VTEP tunnel.
NextHopVrfName	Displays the VRF name of the next hop to reach the remote VTEP.

Configuring a VNID

Use the following procedure to configure a VNID, which identifies one of the potential *16 million* VXLAN segments that can coexist within the same administrative domain.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vnid** tab.
4. In the **Identifier** field, assign a unique value to this VNID.
5. In the **Isid** field, specify the I-SID associated with this VNID.
6. In the **Action** field, select one of the available actions: none or flush.
7. Click **Apply**.

Vnid field descriptions

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

Name	Description
Identifier	Specifies a value that uniquely identifies the VNID of this VXLAN tunnel end point.
Isid	Specifies a value that uniquely identifies the I-SID associated with this VXLAN tunnel end point. * Note: A VNID must not have the same value as an I-SID.
Action	Choose one of the following actions to take with the MAC FDB: • none: Don't take any action. This is the default.

Table continues...

Name	Description
	<ul style="list-style-type: none"> • flushMacFdb: Flush all the learned MAC addresses from the forwarding database of the selected VNID.

Configuring VNID endpoints

Use the following procedure to configure a VNID and the VXLAN tunnel end point.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vnid End Points** tab.
4. In the **Vnid** field, enter the VNID ID.
5. In the **VtepId** field, enter the ID of the remote VTEP.
6. Click **Apply**.

Vnid End Points field descriptions

Use the data in the following table to use the **Vnid End Points** tab.

Name	Description
Vnid	Specifies a value that uniquely identifies the VNID of this VXLAN tunnel end point.
VtepId	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.
Isid	Displays a value that uniquely identifies the I-SID associated with this VXLAN tunnel end point.

Configuring ELAN endpoints

Use the following procedure to configure information associated with ELAN end points.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Elan End Points** tab.
4. In the **Vnid** field, specify the VNID associated with this ELAN.
5. In the **Cvid** field, specify the customer VLAN ID.

6. In the **IfIndex** field, specify the interface index for this ELAN endpoint.
7. Click **Apply**.

Elan End Points field descriptions

Use the data in the following table to use the **Elan End Points** tab.

Name	Description
Vnid	Specifies a value that uniquely identifies the VNID of this ELAN tunnel end point.
Cvid	Specifies a value that uniquely identifies the customer VLAN ID of this ELAN endpoint. The switch reserves the following VLAN IDs: <ul style="list-style-type: none"> • VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. • VLAN IDs 4060 to 4094 are reserved for internal use. • VLAN IDs 3500 to 3999 are reserved if you enable VRF scaling and SPBM mode. • VLAN ID 4095 is not used. • VLAN ID 4096 is reserved for untagged traffic.
Ifindex	Specifies the interface index for the MLT or port for this ELAN endpoint.
Isid	Specifies the I-SID associated with this VNID.

Displaying the VTEP next hop

Use the following procedure to display information about the remote VTEP.

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Next Hop** tab.

Next Hop field descriptions

Use the data in the following table to use the **Next Hop** tab.

Name	Description
Vtepld	Specifies an index value in the range from 1 to 500 that uniquely identifies this remote VTEP.

Table continues...

Name	Description
Ip	Specifies an IP address that uniquely identifies this remote VTEP tunnel.
IfIndex	Specifies the interface index of the next hop to reach the remote VTEP.
Vid	Specifies the VLAN ID of the next hop to reach remote VTEP.

Displaying the VNID forwarding database

Use the following procedure to display information about the VNID forwarding database (FDB).

Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **VTEP**.
3. Click the **Vnid FDB** tab.

Vnid FDB field descriptions

Use the data in the following table to use the **Vnid FDB** tab.

Name	Description
Vnid	Displays the VNID IDs of the VXLAN tunnel endpoints.
Address	Displays all the learned MAC addresses in the forwarding database of the selected VNID.
Status	Displays whether or not the MAC address was learned.
InterfaceIndex	Displays the interface index name for the VNID.
Type	Displays whether the VNID is local or remote.

VXLAN Gateway configuration example

This example shows how to configure the VXLAN Gateway. Remember that there are two modes:

- *Base Interworking Mode* is the default mode so there is no need to set the mode boot flag. However, you should use the `show boot config flags` command to verify that the `vxlan-gw-full-interworking-mode` is set to `false`.
- *Full Interworking Mode* supports all SPB functionality including vIST and SMLT using an internal loopback. Full Interworking Mode is **not** the default mode. To configure Full Interworking mode, you must enable the `vxlan-gw-full-interworking-mode` boot flag.

! Important:

Changing the mode requires a reboot for the change to take effect, which can cause a loss of traffic.




Figure 12: VXLAN Gateway configuration example

```
#
#SPBM CONFIGURATION
#
spbm
spbm ethertype 0x8100

#
#VLAN CONFIGURATION
#
vlan members remove 1 1/1,1/30,2/1/1,2/1/4
vlan create 3000 type port-mstprstp 0
vlan members 3000 1/30 portmember
interface Vlan 3000
ip address 192.168.30.1 255.255.255.0 0
ip ospf enable
exit

#
#PORT CONFIGURATION - PHASE II
#
interface GigabitEthernet 1/1
default-vlan-id 0
flex-uni enable

#
#CIRCUITLESS IP INTERFACE CONFIGURATION - GlobalRouter
#
interface loopback 1
ip address 1 203.0.113.3/255.255.255.255
ip ospf 1
exit

#
#VTEP CONFIGURATION
#
vtep source-ip 198.51.100.1
```

VXLAN Gateway configuration

```
#
#REMOTE VTEP CONFIGURATIONS
#
vtep 1 ip 203.0.113.3

#
#OSPF CONFIGURATION - GlobalRouter
#
router ospf enable
router ospf
exit

#
#IP REDISTRIBUTION CONFIGURATION - GlobalRouter
#
router ospf
redistribute direct
redistribute direct enable
exit

#
#I-SID CONFIGURATION
#
i-sid 300 elan
c-vid 10 port 1/1
exit

#
#VNID CONFIGURATION
#
vnid 100 i-sid 300
vtep 1
exit

#
#IP REDISTRIBUTE APPLY CONFIGURATIONS
#
ip ospf apply redistribute direct
```

Glossary

Bridge Protocol Data Unit (BPDU)	A data frame used to exchange information among the bridges in local or wide area networks for network topology maintenance.
command line interface (CLI)	A textual user interface. When you use CLI, you respond to a prompt by typing a command. After you enter the command, you receive a system response.
common and internal spanning tree (CIST)	The single spanning tree calculated by the Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP) to ensure that all LANs in a bridged Local Area Network (LAN) are simply and fully connected.
common spanning tree (CST)	The single spanning tree calculated by STP, RSTP, and MSTP to connect multiple spanning tree (MST) regions.
Enterprise Device Manager (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.
forwarding database (FDB)	A database that maps a port for every MAC address. If a packet is sent to a specific MAC address, the switch refers to the forwarding database for the corresponding port number and sends the data packet through that port.
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.
Internet Control Message Protocol (ICMP)	A collection of error conditions and control messages exchanged by IP modules in both hosts and gateways.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Local Area Network (LAN)	A data communications system that lies within a limited spatial area, uses a specific user group and topology, and can connect to a public switched telecommunications network (but is not one).

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.
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.
multiple spanning tree bridge	A bridge that supports the common spanning tree (CST) and one or more multiple spanning tree instances (MSTI) and selectively maps frames classified in a VLAN to the CST or an MSTI.
multiple spanning tree configuration identifier	A name for the revision level and summary of a given allocation of VLANs to spanning trees.
multiple spanning tree configuration table	Allocates every possible VLAN to the CST or a specific MSTI.
multiple spanning tree instance (MSTI)	One of a number of spanning trees calculated by the Multiple Spanning Tree Protocol (MSTP) within an MST region, to provide a simple and fully connected active topology for frames that belong to a VLAN mapped to the MSTI.
Multiple Spanning Tree Protocol (MSTP)	Configures multiple instances of the Rapid Spanning Tree Protocol (RSTP) on the switch.
multiple spanning tree region	A set of LANs and MST bridges physically connected by ports on the MST bridges.
Point-to-Point Protocol (PPP)	Point-to-Point Protocol is a basic protocol at the data link layer that provides its own authentication protocols, with no authorization stage. PPP is often used to form a direct connection between two networking nodes.
port	A physical interface that transmits and receives data.
quality of service (QoS)	QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.
Rapid Spanning Tree Protocol (RSTP)	Reduces the recovery time after a network breakdown. RSTP enhances switch-generated Topology Change Notification (TCN) packets to reduce network flooding.

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.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).
Simple Network Management Protocol (SNMP)	SNMP administratively monitors network performance through agents and management stations.
Source Service Access Point (SSAP)	A source service access point (SSAP) is the individual address for access into the upper layers of the network protocol stack. SSAP is an eight bit field address.
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.
trunk	A logical group of ports that behaves like a single large port.
trunk port	A port that connects to the service provider network such as the MPLS environment.
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.