ExtremeSwitching

Configuring IP Multicast Routing Protocols for VOSS

Release 6.3 (VSP 8600)
9035737 Rev.02
June 2019

© 2017-2019, Extreme Networks
All Rights Reserved.

Legal Notice

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

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

Trademarks

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

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

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

Open Source Declarations

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

Contents

Chapter 1: About this Document 10
Purpose 10
Conventions 11
Text Conventions 11
Documentation and Training 13
Getting Help 13
Providing Feedback to Us 14
Chapter 2: New in this Document 15
Notice about Feature Support 15
Chapter 3: IP multicast fundamentals 16
Enabling multicast on the switch 16
Overview of IP multicast 17
IP Multicast over Fabric Connect 21
Internet Group Management Protocol 22
IGMP Layer 2 Querier 29
IGMP Layer 2 Querier limitations 30
Multicast access control 31
Multicast stream limitation feature 34
Multicast Router Discovery protocol 35
Multicast flow distribution over MLT 35
Multicast virtualization 36
Protocol Independent Multicast-Sparse Mode 37
Rendezvous point router 38
Bootstrap router 40
Shared trees and shortest-path trees 40
Receiver joining a group 41
Receiver leaving a group 42
Source sending packets to a group 42
Required elements for PIM-SM operation 43
PIM-SM simplified example 43
PIM-SM static source groups 45
Join and prune messages 45
Register and register-stop messages 45
PIM-SMLT 45
Protocol Independent Multicast-Source Specific Multicast 47
SSM features 47
PIM-SSM architecture 48
PIM-SSM static source groups 50
Implementation of SSM and IGMP 50

Configuration limitations ... 52
PIM passive interfaces ... 53
Multicast route statistics ... 54
IP multicast network design ... 54
Multicast scalability design rules ... 55
IP multicast address range restrictions ... 56
Multicast MAC address mapping considerations ... 57
Dynamic multicast configuration changes ... 58
IGMPv3 backward compatibility ... 58
TTL in IP multicast packets ... 59
Multicast MAC filtering ... 59
Guidelines for multicast access policies ... 60
Split-subnet and multicast ... 60
Protocol Independent Multicast-Sparse Mode guidelines ... 61
Protocol Independent Multicast-Source Specific Multicast guidelines ... 76
Multicast for multimedia ... 76
Layer 3 switch clustering and multicast SMLT ... 78
Protocol Independent Multicast over IPv6 ... 85
PIM-SM over IPv6 features ... 86
Operational note for PIM-SM over IPv6 ... 86
IPv6 interface multiple addresses ... 87
IP multicast configuration and DvR ... 87
Chapter 4: IP multicast basic configuration using CLI ... 89
Configuring IP multicast in SMLT topologies ... 89
Configuring PIM-SM globally ... 92
Enabling or disabling IPv6 PIM-SM globally ... 93
Configuring global IPv6 PIM-SM properties ... 94
Configuring PIM on a VLAN ... 95
Configuring PIM on a port ... 96
Configuring SSM globally ... 98
Configuring IPv6 SSM globally ... 99
Configuring IGMP on a VLAN ... 100
Configuring IGMP ports ... 104
Configuring IGMP brouter ports ... 108
Configuring IGMP on a VRF ... 111
Chapter 5: IP multicast basic configuration using EDM ... 114
Configuring multicast on the switch ... 114
Selecting and launching a VRF context view ... 116
Enabling PIM-SM globally ... 117
Enabling IPv6 PIM-SM globally ... 119
Enabling PIM on a port ... 120
Enabling IPv6 PIM on a port ... 121
Enabling SSM globally ... 122

Enabling IPv6 SSM globally ... 123
Enabling PIM on a VLAN interface ... 124
Enabling IPv6 PIM on a VLAN interface ... 125
Configuring IGMP parameters on a port ... 126
Configuring IGMP parameters on a VLAN ... 128
Chapter 6: Multicast Listener Discovery ... 132
MLD Fundamentals ... 132
MLD versions ... 132
MLD Querier ... 133
MLD snooping ... 133
MLD configuration using the CLI ... 136
Configuring MLD trap generation ... 136
Configuring MLD log status ... 136
Configuring MLD version ... 137
Configuring the MLD last listener query interval ... 138
Configuring the MLD query interval ... 138
Configuring the MLD query maximum response time ... 139
Configuring the MLD robustness ... 140
Enabling MLD snooping on a VLAN ... 141
Enabling MLD ssm-snooping on a VLAN ... 141
Displaying MLD snooping configuration status ... 141
Displaying MLD snooping tracing information ... 142
Displaying MLD interface information ... 143
Displaying MLD system parameters ... 145
Displaying MLD cache information ... 146
Displaying the MLD group information ... 146
Viewing IPv6 MLD host cache ... 148
MLD configuration using EDM ... 149
Configuring MLD globally ... 149
Viewing the MLD SSM global information ... 149
MLD interface configuration ... 150
Configuring MLD snooping ... 155
Viewing the MLD snoop trace information ... 155
Viewing the MLD cache information ... 156
Viewing the MLD V2 cache information ... 157
Viewing IPv6 MLD host cache ... 158
Viewing the MLD source information ... 158
Viewing the MLD sender information ... 159
Viewing the MLD group information ... 160
Chapter 7: PIM configuration using the CLI ... 161
Changing the interface status to passive ... 162
Changing the interface status to active ... 163
Configuring the PIM virtual neighbor ... 165

Configuring a candidate rendezvous point 165
Configuring static RP 167
Configuring IPv6 PIM static RP 168
Configuring a candidate BSR on a port 169
Configuring a candidate BSR on a VLAN 170
Enabling square-SMLT globally 171
Chapter 8: PIM configuration using EDM 172
Enabling static RP 173
Enabling IPv6 static RP 173
Configuring a static RP 174
Configuring an IPv6 static RP entry 175
Viewing the active RP 176
Viewing the IPv6 active RP 177
Configuring a candidate bootstrap router 178
Viewing current BSR information 178
Changing VLAN interface type 179
Editing PIM interface parameters 180
Editing IPv6 PIM interface parameters 181
Configuring the PIM virtual neighbor 182
Viewing PIM-SM neighbor parameters 182
Viewing IPv6 PIM-SM neighbor parameters 183
Viewing IPv6 Neighbor Secondary Address 184
Viewing RP set parameters 184
Configuring a candidate RP 185
Enabling square-SMLT globally 186
Viewing IPv6 RP set parameters 186
Viewing IPv6 Mroute interface information 187
Viewing IPv6 Mroute next hop information 188
Configuring resource usage counter for IPv6 Mroute 189
Viewing IPv6 multicast route information 190
Chapter 9: IGMP configuration using the CLI 192
Configuring multicast stream limitation on an Ethernet port 192
Configuring multicast stream limitation on a VLAN 194
Configuring VLAN multicast stream limitation members 195
Configuring multicast router discovery options 196
Configuring explicit host tracking 198
Configuring IGMP static members 201
Configuring SSM dynamic learning and range group 202
Changing the SSM range group 203
Configuring the SSM map table 205
Configuring multicast access control for an IGMP Ethernet port 206
Configuring multicast access control for a VLAN 207
Configuring fast leave mode 208

Enabling fast leave mode on a port ... 209
Configuring IGMP fast leave members on a VLAN ... 210
Enabling IGMP Layer 2 Querier ... 210
Enabling IGMP Layer 2 Querier address ... 211
Chapter 10: IGMP configuration using EDM ... 213
Enabling IGMP snoop on a VLAN ... 213
Configuring IGMP interface static members ... 214
Configuring the SSM map table ... 215
Configuring SSM range and global parameters ... 216
Configuring multicast stream limitation on an interface ... 217
Configuring multicast stream limitation on a VLAN ... 218
Configuring multicast stream limitation on a port ... 219
Configuring multicast stream limitation members ... 219
Deleting multicast stream limitation member ... 220
Configuring the IGMP interface ... 221
Configuring IGMP sender entries ... 223
Configuring fast leave mode ... 224
Configuring multicast access control for an interface ... 225
Viewing IGMP cache information ... 227
Viewing IGMPv3 cache ... 227
Viewing and editing multicast router discovery information ... 229
Viewing the IGMP router source list ... 230
Viewing IGMP snoop information ... 231
Viewing IGMP Snoop Trace Information ... 232
Viewing IGMP Group Information ... 233
Chapter 11: Route management using the CLI ... 234
Configuring multicast stream limits ... 234
Configuring multicast static source groups ... 236
Configuring IP multicast software forwarding ... 237
Configuring the resource usage counter for multicast streams ... 238
Configuring prefix lists ... 240
Chapter 12: Route management using EDM ... 243
Viewing multicast route information ... 243
Viewing multicast next-hop information ... 244
Viewing multicast interface information ... 246
Adding new static source groups ... 247
Editing static source groups ... 247
Configuring IP multicast software forwarding ... 248
Configuring mroute stream limit ... 250
Configuring Mroute Stream Limit on an Insight Port ... 250
Configuring resource usage counter for multicast streams ... 251
Configuring a prefix list ... 252
Chapter 13: Multicast route statistics configuration using the CLI ... 254

Enabling IP multicast route statistics..... 254
Clearing IP multicast route statistics..... 257
Monitoring IP multicast route statistics..... 257
Enabling IPv6 multicast route statistics..... 260
Clearing IPv6 multicast route statistics..... 263
Monitoring IPv6 multicast route statistics..... 263
Chapter 14: Multicast route statistics configuration using EDM..... 267
Enabling IP multicast route statistics..... 267
Viewing IP multicast route statistics..... 268
Enabling IPv6 multicast route statistics..... 268
Viewing IPv6 multicast route statistics..... 269
Chapter 15: CLI show command reference..... 271
General show commands..... 271
Multicast route information..... 271
Multicast route next hop..... 272
Multicast routes on an interface..... 273
Multicast hardware resource usage..... 274
Static source groups..... 275
VLAN port data..... 276
IGMP show commands..... 276
IGMP access..... 276
IGMP cache..... 277
IGMP group..... 278
IGMP interface..... 279
IGMP multicast router discovery..... 281
IGMP multicast router discovery neighbors..... 282
IGMP router-alert..... 282
IGMP sender..... 283
IGMP snoop..... 284
IGMP static and blocked ports..... 284
Multicast group trace for IGMP snoop..... 285
SSM map information..... 286
SSM group range and dynamic learning status..... 286
PIM show commands..... 287
PIM active RP..... 287
PIM bootstrap router..... 288
PIM candidate rendezvous points..... 288
PIM interface..... 289
PIM mode..... 291
PIM neighbor..... 291
PIM route..... 291
PIM virtual neighbor..... 293
Rendezvous points (for groups)..... 294

Static RP table 294
IPv6 PIM show commands 295
IPv6 PIM mode 295
IPv6 PIM neighbor 296
IPv6 PIM interface 296
Show IPv6 PIM route 298
IPv6 PIM active RP 299
IPv6 Rendezvous points (for groups) 300
IPv6 static RP table 301
IPv6 mroute next-hop 302
IPv6 mroute route 303
IPv6 mroute interface 304
Glossary 306

Chapter 1: About this Document

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

Purpose

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

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

This document describes conceptual and procedural information to administer and configure IP Multicast Routing protocols on the switch. This includes the following operations:

- Internet Group Management Protocol (IGMP)
- Protocol Independent Multicast—Source Specific Multicast (PIM-SSM)
- Protocol Independent Multicast— Sparse Mode (PIM-SM)
- Multicast virtualization
- On the VSP 4000 Series: Multicast MAC Filtering

Configure IP multicast routing to transmit data from a source to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting. However, multicasting transmits data to specific groups, and broadcasting transmits to all devices on the network because multicasting transmits only one stream of data to many destinations, multicasting conserves bandwidth. You must configure at least one IP interface on the switch.

For more information about how to configure interfaces, see Configuring IPv4 Routing for VOSS.

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

Conventions

This section discusses the conventions used in this guide.

Text Conventions

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

Table 1: Notice Icons

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

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command. If the command syntax is `cfm maintenance-domain maintenance-level <0-7>` , you can enter `cfm maintenance-domain maintenance-level 4`.
Bold text	Bold text indicates the GUI object name you must act upon. Examples: • Click **OK**. • On the **Tools** menu, choose **Options**.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.

Table continues…

Convention	Description
	For example, if the command syntax is `ip address {A.B.C.D}`, you must enter the IP address in dotted, decimal notation.
Brackets ([ ])	Brackets ([ ]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is `show clock [detail]`, you can enter either `show clock` or `show clock detail`.
Ellipses (…)	An ellipsis (…) indicates that you repeat the last element of the command as needed. For example, if the command syntax is `ethernet/2/1 [ <parameter> <value> ]...`, you enter `ethernet/2/1` and as many parameter-value pairs as you need.
Italic Text	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
`Plain Courier Text`	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: • `show ip route` • `Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]`
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation tree, expand the **Configuration** > **Edit** folders.
Vertical Line (|)	A vertical line (|) separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is `access-policy by-mac action { allow | deny }`, you enter either `access-policy by-mac action allow` or `access-policy by-mac action deny`, but not both.

Documentation and Training

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

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for earlier versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes
Hardware/Software Compatibility Matrices	https://www.extremenetworks.com/support/compatibility-matrices/
White papers, data sheets, case studies, and other product resources	https://www.extremenetworks.com/resources/

Training

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

Getting Help

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

Extreme Portal	Search the GTAC (Global Technical Assistance Center) knowledge base, manage support cases and service contracts, download software, and obtain product licensing, training, and certifications.
The Hub	A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
Call GTAC	For immediate support: 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

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

Subscribing to Service Notifications

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

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

   Note:
   You can modify your product selections or unsubscribe at any time.
4. Click **Submit**.

Providing Feedback to Us

Quality is our first concern at Extreme Networks, and we have made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but 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, you can do so in two ways:

- Use our short online feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

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

Chapter 2: New in this Document

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 VSP 8600.

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

Chapter 3: IP multicast fundamentals

IP multicast extends the benefits of Layer 2 multicasting on LANs to WANs. Use multicasting techniques on LANs to help clients and servers find each other. With IP multicast, a source can send information to multiple destinations in a WAN with a single transmission. IP multicast results in efficiency at the source and saves a significant amount of bandwidth.

Enabling multicast on the switch

Shortest Path Bridging (SPB) and Protocol Independent Multicast (PIM) cannot interoperate with each other on the switch at the same time. To ensure that SPB and PIM stay mutually exclusive, use the boot flag called **`spbm-config-mode`**:

- The **`spbm-config-mode`** boot flag is enabled by default. This configuration enables you to configure SPB and IS-IS, but you cannot configure PIM either globally or on an interface.
- If you disable the boot flag, you can configure PIM and IGMP Snooping, but you cannot configure SPB or IS-IS.

Important:

- Any change to the **`spbm-config-mode`** boot flag requires a reboot for the change to take effect.
- If you plan to disable the boot flag, remove all SPB configurations first.
- If you plan to use the default (enabled) setting, remove all PIM configurations first.

Simplified Virtual-IST

Simplified Virtual-IST (vIST) is for conventional network deployments that use SMLT and not SPB. The Simplified vIST feature provides a single CLI command to enable the virtual IST for SMLT deployments.

- Simplified vIST is available ONLY for conventional multicast deployments with PIM and IGMP when the boot flag (**`spbm-config-mode`**) is disabled.
- When the boot flag is enabled (default setting), Simplified vIST is not available. This means that you continue to configure SPB/IS-IS for vIST.
- Simplified VIST requires that the two vIST devices be directly connected.

Note:

- PIM is supported with Simplified vIST only, not SPB vIST. However, you do not have to configure Simplified vIST to run PIM or IGMP Snooping in a **non-SMLT** topology.
- LACP is not recommended on SPB NNI MLT links or on the Simplified Virtual IST.
- ECMP is not recommended in PIM Simplified vIST scenarios. Running PIM in a Simplified vIST environment with ECMP enabled may lead to incorrect behavior since there are multiple options in terms of choosing the upstream node towards a host or source. For example, since the path chosen cannot be predicted (it is determined by the downstream PIM neighbor), we may end up not adding the Virtual IST MLT port in the PIM mroute's outgoing port list on the joined interface if the PIM Join Prune Message was received on an alternative path, different from the interface the local router considers to be the correct upstream to the source.

  Traffic loss can occur in such an environment and we recommend not enabling ECMP in PIM vIST scenarios.

After you disable the **`spbm-config-mode`** boot flag, you can configure PIM or IGMP Snooping on any VLAN including the vIST VLAN.

To configure the boot flag and Simplified vIST, see Configuring IP multicast in SMLT topologies on page 89 or Configuring multicast on the switch on page 114.

vIST VLAN IP addresses

Do not configure an RP or BSR on the vIST VLAN because you cannot ping them outside of the vIST VLAN subnet. When you enter the **ip pim enable** command on the vIST VLAN, the following message displays:

```
WARNING: Please do not use virtual IST VLAN IP address for BSR and RP
related configurations, as unicast packets to virtual IST vlan IP address
from outside of virtual IST vlan subnet will be dropped. Use Loopback or
CLIP interface IP address for BSR and RP related configurations.
```

Overview of IP multicast

IP multicast transmits messages to multiple recipients at the same time. This one-to-many delivery mechanism is similar to broadcasting, except that multicasting transmits to specific groups and broadcasting transmits to all receivers on a network. Because IP multicast transmits only one stream of data to the network where it replicates to many receivers, multicasting saves a considerable amount of bandwidth.

IP multicast services benefit applications such as video conferencing, dissemination of datagram information, and dissemination of mail or news to a large number of recipients.

Multicast protocols use different techniques to discover delivery paths.

A distribution tree is a set of multicast routers and subnetworks that permit the members of a group to receive traffic from a source. The source of the tree depends on the algorithm used by the multicast protocol. The following diagram is an example of a simple distribution tree where S is the multicast source and the arrows indicate the multicast broadcast procedure.

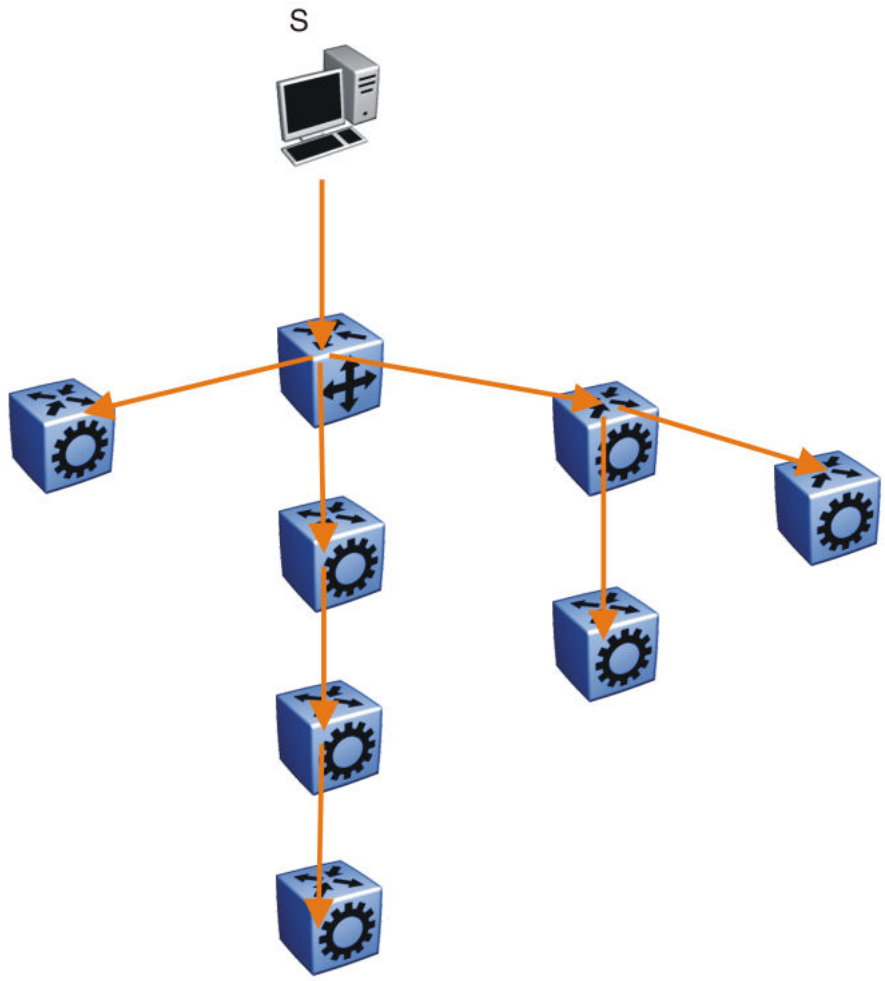

Figure 1: Multicast distribution tree and broadcasting

Broadcast and prune methods use multicast traffic to build the distribution tree. Periodically, the source sends or broadcasts data to the extremities of the internetwork to search for active group members. If no local members of the group exist, the router sends a message to the host, removing itself from the distribution tree, and thus pruning the router.

The following diagram illustrates how the host prunes routers from the distribution tree. First, the router sends a message to the source, after which the pruned routers do not receive multicast data.

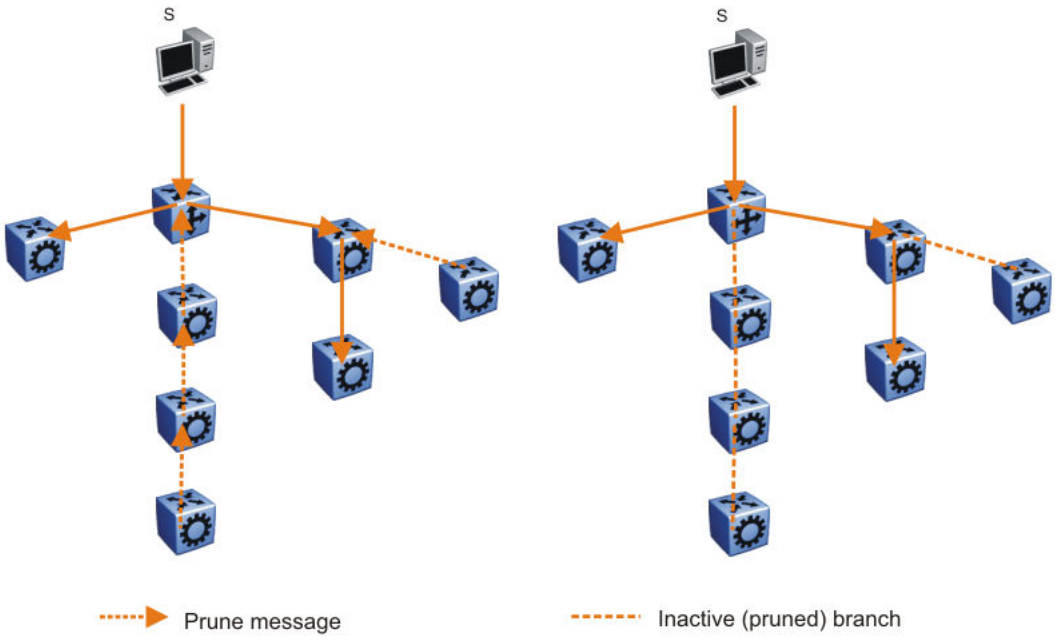

Figure 2: Pruning routers from a distribution tree

Reverse path multicast is based on the concept that a multicast distribution tree is built on the shortest path from the source to each subnetwork that contains active receivers. After a datagram arrives on an interface, the router determines the reverse path to the source of the datagram by examining the routing table of known network sources. If the datagram is not on the optimal delivery tree, the router discards it.

Multicast host groups and their group members enable the IP multicast router to transmit just to those groups interested in receiving the traffic. The switch uses the Internet Group Membership Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. A router communicates with the hosts on a local network by sending IGMP queries. Hosts respond by issuing IGMP reports. For more information about host groups, see Multicast host groups on page 19 and Multicast addresses on page 20. For more information about IGMP, see Internet Group Management Protocol on page 22.

Multicast traffic forwarding transmits frames to all interfaces or subnets for which it receives IGMP reports for the multicast group indicated in the destination IP address. Multicast packets forwarded within the same virtual LAN (VLAN) remain unchanged. The switch does not forward packets to networks that do not use members of the multicast group indicated in the destination IP address.

Multicast host groups

IP multicast is a method for addressing, routing, and delivering a datagram to a collection of receivers called a host group.

Host groups are permanent or transient, with the following characteristics:

- A permanent host group uses a well-known, administratively assigned IP multicast group address. This address is permanent and defines the group. A permanent host group can consist of zero or more members.
- A transient host group exists only as long as members need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

A host system on an IP network sends a message to a multicast group by using the IP multicast address for the group. To receive a message addressed to a multicast group, however, the host must be a member of the group and must reside on a network where that group is registered with a local multicast router.

An IP multicast host group can consist of zero or more members and places no restrictions on its membership. Host members can reside anywhere, they can join and leave the group at any time, and they can be members of more than one group at the same time.

In general, hosts that are members of the same group reside on different networks. However, a range of multicast addresses (224.0.0.x) is reserved for locally-scoped groups. All message traffic for these hosts typically remains on the local network. Hosts that belong to a group in this address range and that reside in different networks do not receive message traffic for each other.

Important:

You can apply a special set of filters (global filters) to multicast packets. You can also create, deny, or accept filters to configure the sources that can receive and send data. For more information about how to configure filters, see Configuring QoS and ACL-Based Traffic Filtering for VOSS.

Multicast addresses

Each host group uses a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP version 4 multicast address is a Class D address (the high-order bits are 1110) from 224.0.0.0 to 239.255.255.255. These addresses are assigned statically for use by permanent groups and dynamically for use by transient groups.

The block of addresses from 224.0.0.1 to 224.0.0.255 is reserved for routing protocols and other low-level protocols. Multicast routers do not forward datagrams with addresses in this range because the time-to-live (TTL) value for the packet is usually 1.

Multicast protocols

You can use the following protocols to enable multicast routing on a switch:

- Internet Group Management Protocol (IGMP)—learns the existence of host group members on directly attached subnets.
- Multicast Router Discovery (MRDISC) protocol—discovers multicast routers in a Layer 2 bridged domain configured for IGMP snoop.
- Protocol Independent Multicast (PIM)
  - Sparse Mode (PIM-SM) protocol—suitable for implementation on networks sparsely populated by receivers.
  - Source Specific Multicast (PIM-SSM) protocol—uses a one-to-many model where members can receive traffic from one or more specific sources. This protocol is suitable for television channels and other content-distribution applications.

Static source groups

Use static source groups to configure static source-group entries in the PIM-SM, or PIM-SSM multicast routing table. PIM cannot prune these entries from the distribution tree. In other words, even if no receivers for the group exist, the multicast stream for a static source-group entry stays active. PIM never prunes static forwarding entries. If you no longer need the entries, you must manually delete them.

To configure static source groups, you must first globally enable PIM. If you disable PIM, the switch saves all of the configured static source-group entries and deactivates them. After you re-enable PIM, the switch reactivates the static source groups.

Static source groups ensure that the multicast route (mroute) records remain in the distribution tree. After receivers join the group, they do not experience a delay in receiving multicast data because they do not need to graft onto the group, or start a join process in the case of PIM. This timing is essential for applications where the multicast data must send to a receiver as soon as the receiver joins the group, for example, when a switch delivers television channels to receivers. After the receiver turns the channel, which is equivalent to joining a group, the receiver can view the channel immediately.

Static entries result in continuous traffic if the source is active, even if no receivers exist. However, the system does not forward traffic with a static entry if no receivers exist, but forwards it continuously to the switch where the entry is programmed and crosses intermediate switches on the path.

You can configure static source-group entries for a specific source or subnet. If several sources on the same subnet send traffic to the same group, traffic for all these sources flows continuously when using the subnet configuration.

After you configure static source groups, keep the following points in mind:

- If you disable PIM, the switch deactivates all of the static source groups. After you re-enable PIM, the switch activates the static source groups.
- In PIM-SM configuration, the static source-group feature works for both specific source addresses and subnet addresses by using the SrcSubnetMask field.

  When the network mask is 255.255.255.255, the full source address is used to match the (S,G) which is the specific source case. When the network mask field is a subnet mask for the source, only the source subnet is used to match (S,G)s.
- In PIM-SSM configurations, static source groups have the following limitations:

  - Subnets: SSM static source groups work only with specific IP addresses. Static source groups cannot work with source subnets, so the mask must use a full 32-bit mask, 255.255.255.255, and the source must use a host address.

IP Multicast over Fabric Connect

IP Multicast over Fabric Connect introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. These extensions, combined with the use of IGMP snooping and querier functions at the edge of the SPBM cloud, efficiently transport IP multicast data by using sub-trees of the VSN shortest path tree per IP multicast group.

For more information about IP Multicast over Fabric Connect, see Configuring Fabric Multicast Services for VOSS.

Internet Group Management Protocol

A host uses IGMP to register group memberships with the local querier router to receive datagrams sent to this router targeted to a group with a specific IP multicast address.

A router uses IGMP to learn the existence of group members on networks to which it directly attaches. The router periodically sends a general query message to each of its local networks. A host that is a member of a multicasting group identifies itself by sending a response.

IGMP queries

When multiple IGMP routers operate on a network, one router is elected to send queries. This elected querier periodically sends host membership queries (also known as general queries) to the attached local subnets. The switch supports queries from all three versions of IGMP.

IGMP host reports

A host that receives a membership query from a local router can respond with a host membership report, one for each multicast group that joins. A host that receives a query delays its reply by a random interval and listens for a reply from other hosts in the same host group. For example, consider a network that includes two host members—host A and host B—of the same multicast group. The router sends out a host membership query on the local network. Both host A and host B receive the query and listen on the network for a host membership report. The delay timer for host B expires first, so it responds to the query with a membership report. Hearing the response, host A does not send a report of its own for the same group.

Each query from a router to a host includes a maximum response time field. IGMP inserts a value n into this field specifying the maximum time in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and n tenths of a second for the period that it waits before sending a response. This calculation is true for IGMP versions 2 and 3. For IGMP version 1, this field is 0 but defaults to a value of 100, that is, 10 seconds.

If at least one host on the local network specifies that it is a member of a group, the router forwards to that network all datagrams that bear the multicast address for the group.

Upon initialization, the host can immediately issue a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same as requested reports.

After hosts and routers are in a steady state, they communicate in a way that minimizes the exchange of queries and reports. The designated routers establish a path between the IP multicast stream source and the end stations and periodically query the end stations about whether to continue participation. As long as a client continues to participate, all clients, including nonparticipating end stations on the switch port, receive the IP multicast stream.

Host leave messages

If an IGMPv2 host leaves a group and it is the host that issues the most recent report, it also issues a leave group message. The multicast router on the network issues a group-specific query to determine whether other group members exist on the network. If no host responds to the query, the router assumes that no members belonging to that group exist on that interface.

Fast leave feature

The switch supports a fast leave feature that is useful for multicast-based television distribution applications. Fast leave relies on an alternative leave process where the switch stops sending traffic

for the group immediately after it receives a leave message, without issuing a query to check if other group members exist on the network. Fast leave alleviates the network from additional bandwidth demand after a customer changes television channels.

The switch provides several fast leave processes for IP multicast:

- immediate leave with one user for each interface
- immediate leave with several users for each interface
- standard IGMP leave based on a Last Member Query Interval (LMQI), which you can configure in tenths of seconds

Fast leave modifies the IGMP leave processing mechanism on an IGMP interface. After the system receives an IGMP leave message on a fast leave enabled interface, the switch does not send a group-specific query and immediately stops sending traffic to the leaving member (IGMP host) port. Without fast leave, traffic continues to forward until the group times out. This situation wastes bandwidth if no receiver that requires the group traffic exists.

Fast leave mode provides two options of the fast leave mechanism—single-user mode and multiple-users mode:

- Single-user mode: In this mode, the port stops receiving traffic immediately after a group member on that port sends a leave message. Use the single-user mode if each interface port connects to only one IGMP host.
- Multiple-users mode: Use this mode if the interface port connects to multiple IGMP hosts. In this case, the port stops receiving traffic after all members leave the IGMP group. The switch removes the leaving IGMP member and, if more group members exist on that port, the switch continues sending traffic to the port.

  When operating in multiple-users mode, the switch must use the correct membership information. To support multiple-users mode, multicast receivers on the same interface cannot use IGMP report suppression. If you must use IGMP report suppression, do not use this mode. Instead, use the LMQI (configurable in units of 1/10ths of seconds) to provide a faster leave process while still sending group-specific queries after the interface receives a leave message.

Fast leave mode applies to all fast-leave enabled IGMP interfaces.

IGMP snoop

The switch provides IP multicast capability and can support all three versions of IGMP to prune group membership for each port within a VLAN. This feature is IGMP snoop.

Important:

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

Use the IGMP snoop feature to optimize the multicast data flow, for a group within a VLAN, to only those ports that are members of the group. The switch builds a database of group members by listening to IGMP reports from each port. The switch suppresses the reports heard by not forwarding them to ports other than the one receiving the report, thus forcing the members to continuously send their own reports. The switch relays group membership from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN. Furthermore, the switch forwards multicast data only to the participating group members and to the multicast routers within the VLAN.

The multicast routing functionality can coexist with IGMP snoop on the same switch, but you can configure only one of IGMP snoop or an IP multicast routing protocol, excluding IGMP, on the same VLAN.

Multicast group trace for IGMP snoop

Use this feature to monitor the multicast group trace for an IGMP snoop-enabled switch . You can view the multicast group trace from CLI.

Multicast group trace tracks the data flow path of the multicast streams. Group trace tracks information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port.

IGMP proxy

If a switch receives multiple reports for the same multicast group, it does not transmit each report to the multicast upstream router. Instead, the switch consolidates the reports into a single report and forwards the one report. If you add another multicast group or the system receives a query since it last transmitted the report upstream, the system forwards the report onto the multicast router ports. This feature is IGMP proxy.

IGMP versions

The switch supports IGMPv1, IGMPv2, and IGMPv3. IGMPv1 and IGMPv2 are backward compatible and can exist together on a multicast network. The following list describes the purpose for each version:

- IGMPv1 provides the support for IP multicast routing. IGMPv1 specifies the mechanism to communicate IP multicast group membership requests from a host to its locally attached routers. For more information, see RFC1112.
- IGMPv2 extends the features in IGMPv1 by quickly reporting group membership termination to the routing protocol. This feature is important for multicast groups with highly volatile group membership. For more information, see RFC2236.
- IGMPv3 supports the PIM Source Specific Multicast (SSM) protocol, PIM-SM, and snooping. A host can selectively request or filter traffic from individual sources within a multicast group or from specific source addresses sent to a particular multicast group. Multicast routing protocols use this information to avoid delivering multicast packets from specific sources to networks where there are no interested receivers. For more information, see RFC3376.

For the switch implementation of PIM-SSM, each group can use multiple sources.

The following list identifies group records that a report message includes:

- current-state record
- source-list-change record
- filter-mode-change record

A current-state record is sent by a system in response to a query received on an interface. It reports the current reception state of that interface, with respect to a single multicast address.

The Record Type of a current-state record has one of the following two values:

- MODE_IS_INCLUDE — Indicates that the interface has a filter mode of include for the specified multicast address. The source address fields in this group record contain the source list of the interface for the specified multicast address.

- MODE_IS_EXCLUDE — Indicates that the interface has a filter mode of exclude for the specified multicast address. The source address fields in this group record contain the source list of the interface for the specified multicast address.

Source-List Change Record — The system sends a source-list-change record after a change of source list occurs that does not coincide with a filter-mode change on the interface for a particular multicast address. The interface on which the change occurs sends a report that includes the record. The record type of a source-list-change record can be one of the following two values:

- ALLOW_NEW_SOURCES — Indicates that the source address [i] fields in this group record contain a list of the additional sources that the system wishes to hear from, for packets sent to the specified multicast address. If the change was to an include source list, these are the addresses that were added to the list. If the change was to an exclude source list, these are the addresses that were deleted from the list.
- BLOCK_OLD_SOURCES — Indicates that the source address [i] fields in this group record contain a list of the sources that the system no longer wishes to hear from, for packets sent to the specified multicast address. If the change was to an include source list, these are the addresses that were deleted from the list; if the change was to an exclude source list, these are the addresses that were added to the list.

  If a change of source list results in both allowing new sources and blocking old sources, then two group records are sent for the same multicast address, one of type ALLOW_NEW_SOURCES and one of type BLOCK_OLD_SOURCES.

Filter Mode — The switch implements the filter-mode-change record. The system sends a filter-mode-change record whenever the filter mode changes (during a change from include to exclude, or from exclude to include) for a particular multicast address. The interface on which the change occurs sends a report that includes the record. The record type of a filter-mode-change record can be one of the following two values:

- CHANGE_TO_INCLUDE_MODE — Indicates that the interface has changed to include filter mode for the specified multicast address. The source address [i] fields in this group record contain the new source list of the interface for the specified multicast address.
- CHANGE_TO_EXCLUDE_MODE — Indicates that the interface has changed to exclude filter mode for the specified multicast address. The source address [i] fields in this group record contain the new source list of the interface for the specified multicast address.

After you enable IGMPv3, the following actions occur:

- After you change the version on an interface to or from IGMPv3, the switch experiences a disruption to existing multicast traffic on that interface but traffic does recover. Do not make this change when the system passes multicast traffic.

IGMP states

Multicast routers implementing IGMPv3 keep one state for each group for every port in every attached network. This group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network. This state consists of a set of records of the following form:

- multicast address
- group timer
- filter mode (source records)

Each source record is of the form source address or source timer. If all sources within a given group are desired, an empty source record list is kept with filter-mode set to EXCLUDE. This means hosts

on this network want all sources for this group to be forwarded. This is the IGMPv3 equivalent to a IGMPv1 or IGMPv2 group join.

Group timer

A group timer represents the time for the filter-mode to expire and switch to INCLUDE mode and is used only when a group is in EXCLUDE mode.

Group timers are updated according to the types of group records received. If a group timer is expiring when a router filter-mode for the group is EXCLUDE means, there are no listeners on the attached network in EXCLUDE mode. At this point, a router will transition to INCLUDE filter-mode.

Source timer

A source timer is maintained for every source record. Source timers are updated according to:

- the type and filter-mode of the group record received
- whenever the source is present in a received record for that group.

If a source timer expires with a router filter-mode for the group of INCLUDE, the router concludes that traffic from this particular source is no longer desired on the attached network, and deletes the associated source record.

If a source record has a running timer with a router filter-mode for the group of EXCLUDE, it means that at least one system desires the source. It should therefore be forwarded by a router on the network. If a source timer expires with a router filter-mode for the group of EXCLUDE, the router informs the routing protocol that there is no receiver on the network interested in traffic from this source. The records are deleted when the group timer expires in the EXCLUDE router filter-mode.

Processing IGMP messages for groups in SSM range

IGMP messages are processed for groups in SSM range in the following scenarios:

1. IGMPv3 interface enabled; PIM-sparse or snooping enabled
   - IGMPv3 reports that contain group records with groups within SSM range are processed with no restrictions.
   - IGMPv2 reports for groups within SSM range translate to IGMPv3 reports with one group record and type IS_EXCLUDE{NULL}. These reports are processed with no restriction as an IGMPv3 report.
   - IGMPv2 leave for groups within SSM range translate to IGMPv3 reports with one group record and type TO_INCLUDE{NULL}. These reports are processed with no restriction as an IGMPv3 report.
2. IGMPv3 interface enabled; PIM-SSM or ssm-snooping enabled
   - IGMPv3 reports that contain group records with groups within SSM range received from members in the EXCLUDE mode are discarded (eg. IS_EXCLUDE and TO_EXCLUDE messages).
   - IGMPv2 reports for groups within SSM range translate to IGMPv3 reports with one group record and type ALLOW{S1,S2,...}. The source list is obtained from the global ssm-map. If there are no sources in the global ssm-map, the message is discarded. These reports are processed with no restriction as an IGMPv3 report.
   - IGMPv2 leave for groups within SSM range translate to IGMPv3 reports with one group record and type BLOCK{S1,S2,...}. The source list is obtained from the global ssm-map. If

there are no sources in the global ssm-map, the message is discarded. These reports are processed with no restriction as an IGMPv3 report.

Note:

In order to accept v2 messages, you must enable the compatibility mode on the IGMPv3 interface.

IGMPv3 source-specific forwarding rules

After a multicast router receives a datagram from a source destined to a particular group, the router must decide to forward the datagram to the attached network. The multicast routing protocol uses IGMPv3 information to forward datagrams to all required sources or groups on a subnetwork.

The following table describes the forwarding suggestions that IGMPv3 makes to the routing protocol. The table also identifies the action taken after the source timer expires, based on the filter mode of the group.

Group filter-mode	Source-timer value	Action
INCLUDE	TIMER > 0	Forward the traffic from the source.
INCLUDE	TIMER = 0	Stop forwarding the traffic from the source, and remove the source record. If no more source records exist for the group, delete the group record.
INCLUDE	No source elements	Do not forward the source.
EXCLUDE	TIMER > 0	Forward the traffic from the source.
EXCLUDE	TIMER = 0	Do not forward the traffic from the source. If no more source records exist for the group, delete the group record.
EXCLUDE	No source elements	Forward the traffic from the source.

IGMPv3 explicit host tracking

IGMPv3 explicit host tracking enables the IGMP to track all the source and group members. To track all the source and group members, the sources that are in the include mode hold a list of members who want to receive traffic from that source.

The members that are in the exclude mode are on hold on the reporter list under the port data. By default, IGMPv3 explicit host tracking is disabled.

Important:

If explicit host tracking is enabled, you cannot downgrade the IGMPv3 interface to IGMPv1 or IGMPv2.

IGMPv3 fast leave

When a BLOCK message is received for a source, you must check if the member that sent this message is the last reporter for the source. If it is the last reporter, delete the source. Else, delete the member. No group and source specific queries are sent.

When a LEAVE message is received, you must check if the member that sent this message is the last reporter for the group. If it is the last reporter, switch to INCLUDE mode if sources are available (if no sources are available the port is deleted). Else, delete the member. No group and source specific queries or group specific queries are sent.

Important:

To use the IGMPv3 fast leave feature, you must first enable the explicit host tracking feature.

Synchronization of IGMPv3 over SMLT

The implementation of IGMPv3 offers support for IGMPv3 over SMLT. The Virtual-IST (vIST) peers must be in sync with the IGMPv3 reports received over SMLT links to ensure effective performance. The vIST protocol ensures the infrastructure to send such information from one vIST peer to the other.

The synchronization of IGMPv3 members and their advertised sources is different from IGMPv1 and IGMPv2. Because of IGMPv3 compatibility mode, you must consider the IGMP member version. If you have version 1 or 2 members, you must synchronize the IGMP information as IGMPv1 or IGMPv2 reports, so the peer can build an accurate database. In particular, if members with version 1 or 2 exist, the group filter mode is exclude and the exclude source list is empty. Also no v1 or v2 member will be present on any source from include list.

Each member sends IGMP reports in the same manner for all IGMP versions. The sending mechanism depends on the SMLT state.

After a vIST peer receives an IGMPv3 report over an SMLT link, it must pass the message to its peer. If the SMLT state is up, the vIST peer sends the message encapsulated in an vIST IGMPv3 message. If the SMLT state is down, the vIST peer sends the message as a plain IGMPv3 report.

In both cases the IGMPv3 message is not altered and the receiving vIST peer processes it as expected in SMLT conditions (translating the receiving port to SMLT port if applicable).

Note:

If you enable compatibility mode and the member sends an IGMPv1 or IGMPv2 report, the message is either a vIST IGMPv1 or v2 encapsulated Message or a plain IGMPv1 or IGMPv2 report.

After SMLT up or down events occur, the vIST peer must synchronize its IGMPv3 database to its peer, taking into account the new state of the SMLT link.

If you enable IGMP explicit host tracking, each include source stores information for each member that advertises that particular source in an include list. This information is synchronized with the vIST peer.

If you do not enable explicit host tracking, each source from include list contains only information related to the last member that sent an IGMPv3 report. Only this information is synchronized with the vIST peer.

Backward compatibility

IGMPv3 for PIM-SSM is backward compatible with IGMPv2. You can configure the switch to operate in v3-only mode or in v2-v3 compatibility mode. If you configure the switch to use v3-only mode, it ignores all v2 and v1 messages except the query message.

If you configure the switch to operate in v2-v3 compatibility mode, the switch supports all IGMPv1, v2, and v3 messages. The switch parses the group address of the messages. If the group address is out of SSM range and it is a v3 message, the switch drops the message; if it is a v2 message, PIM-SM or IGMP snoop processes handle the message.

After the switch receives an IGMPv2 leave message and the group address in it is within SSM range, the switch sends the group-and-source specific query. If the group address is not within the SSM range, the switch sends the group specific query.

According to RFC3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.
- If a router that is not explicitly configured to use IGMPv1 or IGMPv2 hears an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure if the switch dynamically downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

In v2-v3 compatibility mode, an IGMPv2 host can only join if you configure a static entry in SSM map and if the interface operates in PIM-SSM mode or IGMP SSM-Snoop mode.

You can use the compatibility mode with Split MultiLink Trunking (SMLT). One core switch sends an SMLT message to the other core switch after it receives an IGMPv3 message. This action synchronizes the IGMP host information.

Implementation of IGMP

You can enable and disable multicast routing on an interface basis. If you disable multicast routing on an interface, the interface does not generate IGMP queries. If the switch or interface is in IGMP router behavior mode, for example, PIM enabled, you cannot configure IGMP snoop. The switch still learns the group membership and snoops multicast receivers on the switch VLAN or ports.

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.

IGMP Layer 2 Querier limitations

The following limitations apply to IGMP Layer 2 Querier.

- IGMP Layer 2 Querier is based on IGMP Snoop. If you disable IGMP Snoop, IGMP Layer 2 Querier does not work until you enable IGMP Snoop and IGMP Layer 2 Querier.
- After you enable IGMP Snoop and IGMP Layer 2 Querier on an interface, if the system receives no IGMP query messages, it becomes the querier.

IGMP Layer 2 Querier limitations and DvR

The following limitations apply when you configure IGMP Layer 2 Querier on DvR enabled nodes.

- You can configure IGMP Layer 2 Querier only on the DvR Controllers in a DvR domain. When you configure the following parameters on the Controllers, the configuration is automatically pushed to the DvR Leaf nodes within the domain.
  - IGMP version
  - IGMP query interval

  - IGMP query maximum response time
  - IGMP robustness value
  - IGMP last member query interval
  - IGMP compatibility mode
- You cannot configure IGMP snooping on DvR enabled Layer 2 VSNs.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Multicast access control

Multicast access control is a set of features that operate with standard existing multicast protocols. You can configure multicast access control for an IP multicast-enabled port or VLAN with an access control policy that consists of several IP multicast groups.

You can use this feature to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams). For example, in a television distribution application, instead of applying a filter to each channel (multicast group), you can apply a multicast access policy to a range of channels (groups), thereby reducing the total number of filters and providing a more efficient and scalable configuration. Also, if you want to add or remove television channels from a package, you can modify the multicast access policy; you do not need to change filters for individual VLANs or ports. Multicast access policies contain an ID and a name (for example, PremiumChannels), the list of IP multicast addresses, and the subnet mask.

Multicast access control is not a regular filtering configuration. Multicast access control is for multicast streams and relies on handling multicast control and initial data to prevent hosts from sending or receiving specified multicast streams; it does not use filters. Also, multicast access control provides a list of multicast groups in one configuration using the same routing policy prefix list configuration. For information about prefix lists, see Configuring IPv4 Routing for VOSS. You can configure multicast access control and change it dynamically to support changes in the configuration without restarting the protocol. You can change the access capabilities of a user or service subscriber without loss of service.

The following paragraph describes a typical application.

The local cable television company offers three packages; each one includes 35 channels (35 multicast groups). The company configures each package in an access control policy. This policy applies to a set of VLANs or ports to prevent users from viewing the channels on those VLANs. Use the same policy to prevent users from sending traffic to those groups (also known as spoofing) by specifying the deny-tx option for that port. After you define the packages, you can use them for access policy configuration. You can easily change the package by changing the group range, without changing all the port configurations.

The multicast access control functionality applies to an IP multicast application where you must control user access. You can use it in financial-type applications and other enterprise applications, such as multicast-based video conferencing.

Six types of multicast access control policies exist:

- deny-tx

- deny-rx
- deny-both
- allow-only-tx
- allow-only rx
- allow-only-both

The tx policies control the sender and ingress interface for a group; the rx policies control the receivers and egress interface for a group.

deny-tx

Use the deny-tx access policy to prevent a matching source from sending multicast traffic to the matching group on the interface where you configure the deny-tx access policy. Configure this policy on the ingress interface to the multicast source. The deny-tx access policy performs the opposite function of the allow-only-tx access policy. Therefore, the deny-tx access policy and the allow-only-tx access policy cannot exist on the same interface at the same time.

For example, in Figure 3: Data flow using deny-tx policy on page 32, a VLAN 1, the ingress VLAN, uses a deny-tx access policy. This policy prevents multicast traffic sent by Sender from forwarding from VLAN 1 to a receiver, consequently preventing Receiver 1 and Receiver 2 from receiving data from the multicast group. You can create receive-only VLANs, such as VLAN 1, with the deny-tx policy.

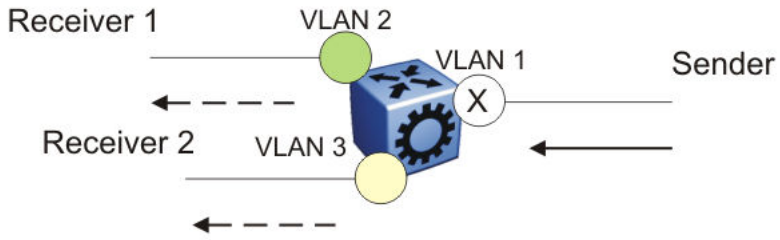

Figure 3: Data flow using deny-tx policy

deny-rx

Use the deny-rx access policy to prevent a matching group from receiving IGMP reports from the matching receiver on the interface where you configure the deny-rx access policy. The deny-rx access policy performs the opposite function of the allow-only-rx access policy. Therefore, the deny-rx access policy and the allow-only-rx access policy cannot exist on the same interface at the same time.

For example, in Figure 4: Data flow using deny-rx policy on page 33, a VLAN 2 uses a deny-rx access policy, preventing IGMP reports sent by Receiver 1 from receiving on VLAN 2. You can deny a multicast group access to a specific VLAN or receiver using the deny-rx policy.

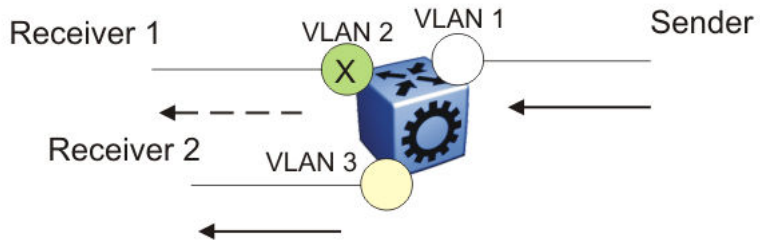

Figure 4: Data flow using deny-rx policy

deny-both

Use the deny-both access policy to prevent a matching IP address from both sending multicast traffic to, and receiving IGMP reports from, a matching receiver on an interface where you configure the deny-both policy. You can use this policy to eliminate all multicast activity for a receiver or source in a specific multicast group. The deny-both access policy performs the opposite function of the allow-only-both access policy. Therefore, the deny-both access policy and the allow-only-both access policy cannot exist on the same interface at the same time.

For example, in Figure 5: Data flow using deny-both policy on page 33, a VLAN 2 uses a deny-both access policy, preventing VLAN 2 from receiving IGMP reports sent by Receiver 2, and preventing multicast traffic sent by Sender 2 from forwarding from VLAN 2. You can prevent certain VLANs from participating in an activity involving the specified multicast groups with the deny-both policy.

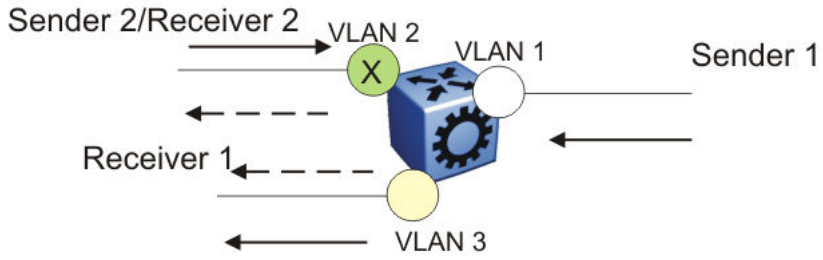

Figure 5: Data flow using deny-both policy

allow-only-tx

Use the allow-only-tx policy to allow only the matching source to send multicast traffic to the matching group on the interface where you configure the allow-only-tx policy. The interface discards all other multicast data it receives. The allow-only-tx access policy performs the opposite function of the deny-tx access policy. Therefore, the allow-only-tx access policy and the deny-tx access policy cannot exist on the same interface at the same time.

allow-only-rx

Use the allow-only-rx policy to allow only the matching group to receive IGMP reports from the matching receiver on the interface where you configure the allow-only-rx access policy. The interface discards all other multicast data it receives. The allow-only-rx access policy performs the opposite function of the deny-rx access policy. Therefore, the allow-only-rx access policy and the deny-rx access policy cannot exist on the same interface at the same time.

allow-only-both

Use the allow-only-both policy to allow only the matching IP address to both send multicast traffic to, and receive IGMP reports from, the matching receiver on the interface where you configure the allow-only-both access policy. The interface discards all other multicast data and IGMP reports. The allow-only-both access policy performs the opposite function of the deny-both access policy. Therefore, the allow-only-both access policy and the deny-both access policy cannot exist on the same interface at the same time.

Host addresses and masks

When you configure multicast access policies, you must specify the host (IP) address and host (subnet) mask of the host to filter (the host that sends multicast traffic).

You can use the host subnet mask to restrict access to a portion of the host network. For example, if you configure the host subnet mask as 255.255.255.255, you use the full host address. To restrict access to a portion of the network of a host, use a subnet mask such as 255.255.255.0. Access control applies to the specified subnet only.

Multicast stream limitation feature

You can configure the multicast stream limitation feature to limit the number of multicast groups that can join a VLAN. By limiting the number of concurrent multicast streams, a service provider can, for example, protect the bandwidth on a specific interface and control access to multicast streams.

Use multicast stream limitation in an environment where you want to limit users to a certain number of multicast streams simultaneously. For example, a television service provider can limit the number of television channels a user can watch at a time. (To a television service provider, a multicast stream is synonymous with a television channel.) If a user purchases a service contract for two single-tuner television receivers, they can use two channels flowing at the same time, but not a third. The service provider can control the bandwidth usage in addition to preventing users from watching more than the allowed number of channels at a point in time.

You can enable the multicast stream limitation feature on the switch by using one of the following methods:

- for each interface—This limitation controls the total number of streams for all clients on this brouter port.
- for each VLAN—This limitation controls the total number of streams for all clients on this VLAN. This method is equivalent to the interface stream limitation.
- for each VLAN port—This limitation controls the number of streams for all clients on this VLAN port. This method is equivalent to the interface port stream limitation.

You can configure the maximum number of streams for each limit independently. After the number of streams meets the limit, the interface drops additional join reports for new streams. The maximum number of streams for each limit is 65535 and the default is 4.

Multicast Router Discovery protocol

The Multicast Router Discovery (MRDISC) protocol can automatically discover multicast-capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and IGMP host membership reports. This feature is useful in a Layer 2 bridging domain that you configure for IGMP snoop.

IGMP multicast router discovery consists of three message types that discover multicast routers on the network:

- Multicast router advertisements: routers advertise that IP multicast forwarding is enabled on an interface.
- Multicast router solicitations: routers solicit a response of multicast router advertisements from all multicast routers on a subnet.
- Multicast router termination messages: a router terminates its multicast routing functions.

Multicast routers send multicast router advertisements periodically on all interfaces where you enable multicast forwarding. Multicast routers also send advertisements in response to multicast router solicitations.

Multicast router solicitations transmit to the IGMP-MRDISC all-routers multicast group that uses a multicast address of 224.0.0.2. Multicast router solicitations do not transmit if a router needs to discover multicast routers on a directly attached subnet.

Multicast router termination messages transmit after a router terminates its multicast routing functions. Other non-IP forwarding devices, such as Layer 2 switches, can send multicast router solicitations to solicit multicast router advertisements.

To function MRDISC on IGMP snoop interface, you must explicitly enable MRDISC. The Solicitation messages are sent only if IGMP snoop and MRDISC are enabled on the switch.

Multicast flow distribution over MLT

MultiLink Trunking (MLT) is a mechanism to distribute multicast streams over a multilink trunk and achieve an even distribution of the streams. The distribution is based on source-subnet and group addresses. In applications like television distribution, multicast traffic distribution is particularly important because the bandwidth requirements are substantial when you use a large number of television streams.

The switch enables this feature by default and you can not change the configuration.

Traffic distribution

Traffic distribution distributes the streams on the multilink trunk links if an MLT configuration change occurs. For example, you can add or delete ports.

This feature distributes active streams according to the distribution algorithm on the multilink trunk links. This distribution can cause minor traffic interruptions. To minimize the effect of distribution of multicast traffic on the multilink trunks, the implementation does not move the streams to the

appropriate links at the same time. Instead, it distributes a few streams at every time tick of the system.

To that end, after a multilink trunk port becomes inactive, this feature distributes all the streams on the multilink trunk ports based on the assignment provided by the distribution algorithm.

By default, distribution is enabled and you can not change the configuration.

For more information about MLT, see Configuring Link Aggregation, MLT, SMLT and vIST for VOSS .

Multicast virtualization

Multicast provides simplified extension of internal video and data delivery to remote locations.

Virtualized multicast enables multiple VPN routing instances on devices and supports various unicast routing protocols so that you can provide the services of many virtual routers from one physical device.

You can configure multicast routing support with the Virtual Routing and Forwarding (VRF) Lite feature and you can use VRF Lite to emulate many virtual routers with one router.

Multicast virtualization support includes:

- IGMP snooping
- IGMP in Layer 2 virtual services networks (VSN)
- IGMP in Layer 3 VSNs

To implement multicast virtualization, you must perform the following tasks:

1. Create a VRF. For more information about how to create and configure a VRF, see Configuring IPv4 Routing for VOSS.
2. Create a VLAN and associate it with the VRF.
3. Enable one of the following: IGMP snooping on the VLAN, Layer 2 VSN, or Layer 3 VSN.

If you use IGMP snooping on the VLAN, ensure the IGMP version on the multicast hosts or other network devices is either the same as the version on the VLAN, or enable compatibility mode.

Multicast virtualization does not support PIM. The switch supports IGMP with PIM only in the Global Router.

VRF Lite background

VRF Lite provides independent IPv4 forwarding instances and independent routing instances (contexts), which can reside on the same or different VLANs and ports.

While forwarding and routing instances are mapped to IP interfaces, incoming traffic is classified into a VLAN and IP interface and, depending on the IP interface, routed context traffic is forwarded.

Protocol Independent Multicast-Sparse Mode

PIM-SM, as defined in RFC2362, supports multicast groups spread out across large areas of a company or the Internet. PIM-SM sends multicast traffic only to routers that specifically join a multicast group. This technique reduces traffic flow over WAN links and overhead costs for processing unwanted multicast packets.

Dense-mode protocols use a flood-and-prune technique, which is efficient with densely-populated receivers. However, for sparsely populated networks, PIM-SM is more efficient because it sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic.

PIM-SM is independent of a specific unicast routing protocol, but it does require the presence of a unicast routing protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). PIM-SM uses the information from the unicast routing table to create and maintain multicast trees that enable PIM-enabled routers to communicate.

Typically, a PIM-SM network consists of several multipoint data streams, each targeted to a small number of LANs in the internetwork. For example, customers whose networks consist of multiple hosts on different LANs can use PIM-SM to simultaneously access a video data stream, such as video conferencing, on a different subnet.

Important:

In some cases, PIM stream initialization can take several seconds.

Hosts

A host is a source, a receiver, or both:

- A source, also known as a sender, sends multicast data to a multicast group.
- A receiver receives multicast data from one or several sources that sends data to a multicast group.

PIM-SM domain

PIM-SM operates in a domain of contiguous routers on which PIM-SM is enabled.

Each PIM-SM domain requires the following routers:

- designated router (DR)
- rendezvous point (RP) router
- bootstrap router (BSR)

Although a PIM-SM domain can use only one active RP router and one active BSR, you can configure additional routers as a candidate RP (C-RP) router and as a candidate BSR (C-BSR). Candidate routers provide backup protection in case the primary RP router or BSR fails.

As a redundancy option, you can configure several RPs for the same group in a PIM domain. As a load sharing option, you can have several RPs in a PIM-SM domain map to different groups. The switch devices use the hash function defined in the PIM-SM standard to elect the active RP.

Designated router

The designated router (DR), the router with the highest IP address on a LAN, performs the following tasks:

- sends register messages to the RP router on behalf of directly connected sources
- sends join and prune messages to the RP router on behalf of directly connected receivers
- maintains information about the status of the active RP router for local sources in each multicast group

Important:

The DR is not a required configuration. Switches act automatically as the DR for directly attached sources and receivers.

Rendezvous point router

PIM-SM builds a shared multicast distribution tree within each domain, and the RP router is at the root of this shared tree. Although you can physically locate the RP anywhere on the network, it must be as close to the source as possible. Only one active RP router exists for a multicast group.

At the RP router, receivers meet new sources. Sources use the RP to identify themselves to other routers on the network; receivers use the RP to learn about new sources.

The RP performs the following tasks:

- registers a source that wants to announce itself and send data to group members
- joins a receiver that wants to receive data for the group
- forwards data to group

Candidate rendezvous point router

You can configure a set of routers as C-RP routers that serve as backup to the RP router. If an RP fails, all the routers in the domain apply the same algorithm to elect a new RP from the group of C-RP routers. To make sure that the routers use a complete list of C-RP routers, the C-RP router periodically sends unicast advertisement messages to the BSR. The most common implementation is to configure a PIM-SM router as both a C-RP router and a C-BSR.

The switch devices use the hash function defined in the PIM-SM standard to elect the active RP.

Static rendezvous point router

You can configure a static entry for an RP router with static RP. This feature avoids the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. Static RP-enabled switches cannot learn about RPs through the BSR because the switch loses all dynamically learned BSR information and ignores BSR messages. After you configure static RP entries, the switch adds them to the RP set as if they were learned through the BSR.

Important:

In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interface configured as an RP.

When you configure a PIM static RP in a switch, the next hop of the unicast route toward the PIM static RP must be a PIM neighbor. The PIM protocol fails to work, due to a route change, if the next hop toward an already configured static RP becomes a non-PIM neighbor. If a PIM neighbor cannot reach the configured RP, the RP does not activate and its state remains invalid.

A static RP-enabled switch can communicate with switches from other vendors that do not use the BSR mechanism. Some vendors use either early implementations of PIM-SM v1 that do not support the BSR or proprietary mechanisms like the Cisco Auto-RP. For a network to work properly with static RP, you must map all the switches in the network (including switches from other vendors) to the same RP or RPs, if several RPs exist in the network.

To avoid a single point of failure, you can also configure redundant static RPs.

Use the static RP feature when you do not need dynamic learning mode, typically in small networks, or for security reasons, where RPs are forced to devices in the network so that they do not learn other RPs.

Static RP configuration considerations

Before you can configure a static RP, you must enable PIM-SM and enable static RP.

After you meet these prerequisites, keep in mind the following configuration considerations:

- You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.
- All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.
- Static RPs do not age, that is, they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless you configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP.
- In a PIM domain with both static and dynamic RP switches, the static RP switches cannot use a local interfaces configured as an RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of vendor switches across the network, you must ensure that all switches and routers use the same active RP because other vendors can use different algorithms to elect the active RP. The switch devices use the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.

  Important:

  To reduce convergence times, create only one static RP for each group. The more static RPs you configure for redundancy, the more time PIM requires to rebuild the mroute table and associate RPs.

- Static RP configured on the switch is active as long as the switch uses a unicast route to the static RP network. If the switch loses this route, the static RP is invalidated and the hash

algorithm remaps all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm remaps the affected groups.

Bootstrap router

The BSR receives RP router advertisement messages from the candidate RPs. The BSR adds the RP router with its group prefix to the RP set. Only one BSR exists for each PIM-SM domain.

The BSR periodically sends bootstrap messages containing the complete RP set to all routers in the domain. The BSR ensures that all PIM-SM routers send join, prune, and register packets.

Within a PIM-SM domain, you can configure a small set of routers as C-BSRs. The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Important:

Configure C-BSRs on routers that are central to all candidate RPs.

Shared trees and shortest-path trees

A PIM-SM domain uses shared trees and shortest-path trees to deliver data packets to group members. This section describes both trees.

Shared trees

Group members in a PIM-SM domain receive the first packet of data from sources across a shared tree. A shared tree consists of a set of paths that connect all members of a multicast group to the RP. PIM creates a shared tree when sources and receivers send messages toward the RP.

Shortest-path trees

After receiving a certain number of packets from the RP, the DR changes from a shared tree to an SPT. Switching to an SPT creates a direct route between the receiver and the source. The switch changes to the SPT after it receives the first packet from the RP.

Figure 6: Shared tree and shortest-path tree on page 41 shows a shared tree and an SPT.

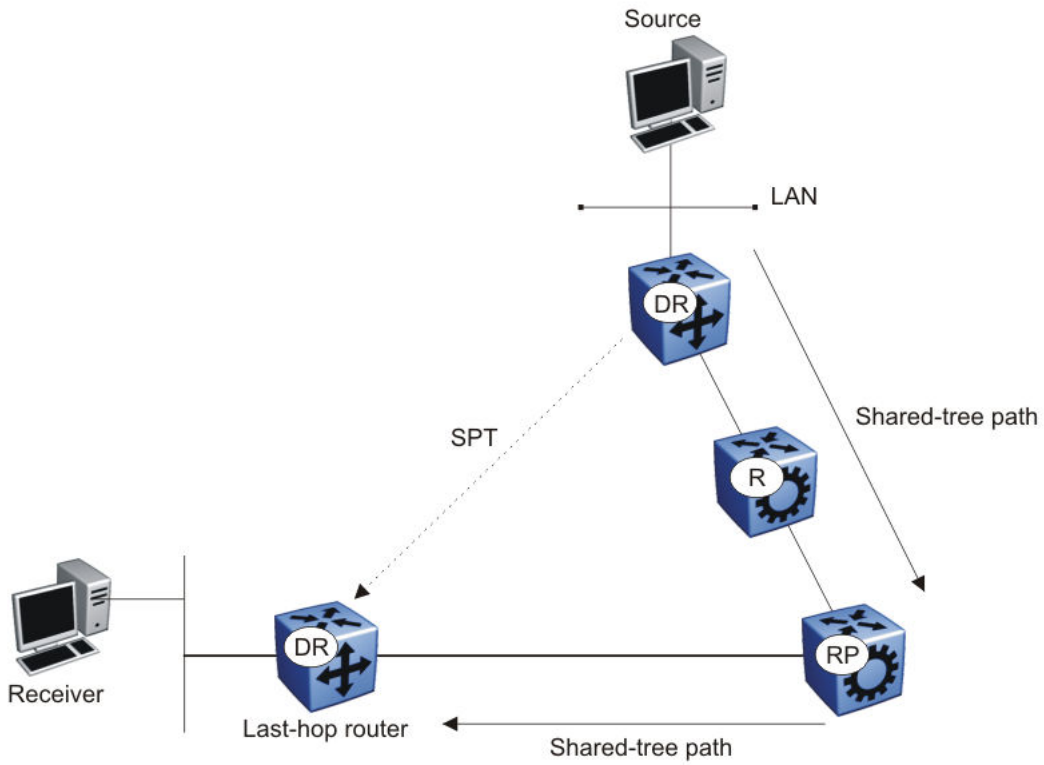

Figure 6: Shared tree and shortest-path tree

Receiver joining a group

The following steps describe how a receiver joins a multicast group:

1. A receiver multicasts an IGMP host membership message to the group that it wants to join.
2. After the last-hop router (the DR), normally the PIM router with the highest IP address for that VLAN, receives the IGMP message for a new group join, the router looks up the associated elected RP with responsibility for the group.
3. After it determines the RP router for the group, the last-hop router creates a (*,G) route entry in the multicast forwarding table and sends a (*,G) join message to the RP. After the last-hop router receives data packets from the RP, if the multicast packet arrival rate exceeds the DR threshold, the last-hop router switches to the SPT by sending an (S,G) join message to the source. (S denotes the source unicast IP address, and G denotes the multicast group address.)
4. If the last-hop router switches to the SPT, the following actions occur:
   - All intermediate PIM routers along the path to the source create the (S,G) entry.
   - To trim the shared tree, the router sends an (S,G) prune message to the RP.

Receiver leaving a group

Before it leaves a multicast group, a receiver sends an IGMP leave message to the DR. If all directly connected members of a multicast group leave or time out, and no downstream members remain, the DR sends a prune message upstream and PIM-SM deletes the route entry after that entry times out.

When the system ages PIM mroutes, it does not clear the (S,G) entry for an inactive route immediately after the expiration period. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.

Source sending packets to a group

The following steps describe how a source sends multicast packets to a group:

1. A source directly attached to a VLAN bridges the multicast data to the DR. The DR for the VLAN (the router with the highest IP address) encapsulates each packet in a register message and sends a unicast message directly to the RP router to distribute to the multicast group.
2. If a downstream group member chooses to receive multicast traffic, the RP router sends a join or prune message toward the source DR and forwards the data down the RP tree after it obtains the data natively.
3. After the receiver DR obtains the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.
4. If no downstream members want to receive multicast traffic, the RP router sends a register-stop message (for the source) to the DR.

   The DR starts the register suppression timer after it receives the first register-stop message. During the register suppression timeout period (the default is 60 seconds), the following events occur:

   - The DR for the source sends a probe packet to the RP router before the register suppression timer expires. The probe packet prompts the RP router to determine whether new downstream receivers joined the group.
   - If no new receivers joined the group, the RP router sends another register-stop message to the DR for the source, and its register suppression timer restarts.
   - After the RP router no longer responds with a register-stop message to the source DR probe message, the register suppression timer expires and the DR sends encapsulated multicast packets to the RP router. The RP router uses this method to tell the DR that new members joined the group.

The RP sends a register-stop message to the DR immediately after it receives the first multicast data packet.

Required elements for PIM-SM operation

For PIM-SM to operate, the following elements must exist in the PIM-SM domain:

- You must enable an underlying unicast routing protocol for the switch to provide routing table information to PIM-SM.
- You must configure an active BSR to send bootstrap messages to all PIM-v2 configured switches and routers to enable them to learn group-to-RP mapping. If you configure several BSRs in a network, an active BSR is elected based on priority and IP address (if priority is equal, the BSR with the higher IP address is elected).
- You must include an RP to perform the following tasks:
  - manage one or several IP multicast groups
  - become the root for the shared tree to these groups
  - accept join messages from receiver switches for groups that it manages
  - elect an active RP based on priority and IP address (if priority is equal, the RP with the higher IP address is elected)

PIM-SM simplified example

Figure 7: PIM-SM simplified example on page 44 shows a simplified example of a PIM-SM configuration.

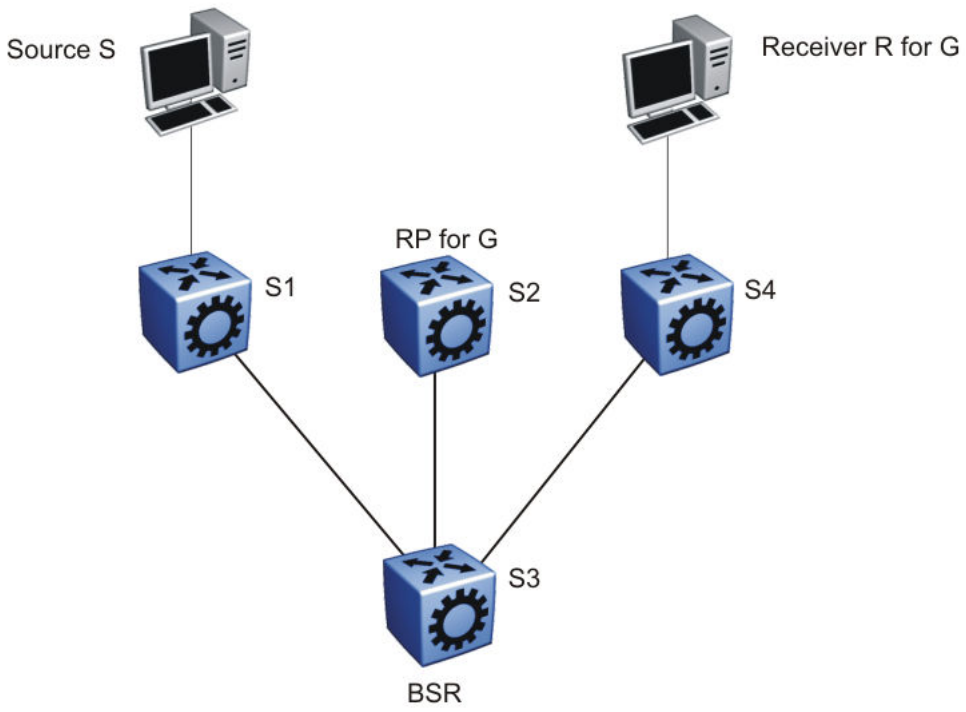

Figure 7: PIM-SM simplified example

In the sample configuration, the following events occur:

1. The BSR distributes RP information to all switches in the network.
2. R sends an IGMP membership report to S4.
3. Acting on this report, S4 sends a (*,G) join message to RP.
4. S sends data to G.
5. The DR (S1 in this example) encapsulates the data that it unicasts to RP (S2) in register messages.
6. S2 decapsulates the data, which it forwards to S4.
7. S4 forwards the data to R.
8. If the packet rate exceeds the DR threshold, S4 sends S1 an (S,G) join message.
9. S1 forwards data to S4. After S4 receives data from S1, it prunes the stream from the RP.

Important:

Figure 7: PIM-SM simplified example on page 44 is a simplified example and is not the best design for a network if you locate the source and receiver as shown. In general, place RPs as close as possible to sources.

PIM-SM static source groups

You can configure static source groups as static source-group entries in the PIM-SM multicast routing table. PIM-SM cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 20.

Join and prune messages

The DR sends join and prune messages from a receiver toward an RP for the group to either join the shared tree or remove (prune) a branch from it. A single message contains both a join and a prune list. This list includes a set of source addresses that indicate the shortest-path trees or the shared trees that the host wants to join. The DR sends join and prune messages hop-by-hop to each PIM router on the path to the source or the RP.

Register and register-stop messages

The DR sends register messages to the RP for a directly connected source. The register message informs the RP of a new source, causing the RP to send join or prune messages back toward the DR of the source, which forwards the data down the RP tree after it obtains the data natively. After the receiver DR obtains the first packet, it switches to the shortest-path tree (SPT) and continues receiving data through the SPT path.

The DR stops sending encapsulated packets to the RP after it receives a register-stop message. This traffic stops without delay because the RP sends a register-stop message immediately after it receives the first multicast data packet, and joins the shortest-path tree.

PIM-SMLT

IP multicast routing support with Split MultiLink Trunking (SMLT) builds a virtual switch that represents the two switches of the split multilink trunk core.

When switches use PIM in the core, they need to exchange protocol-related updates as part of the interswitch trunking (IST) protocol. IST hides the fact that the edge switch attaches to two physical switches.

PIM-SMLT can work in triangular, square, and full mesh configurations with Layer 3 IP multicast. However, PIM-SSM in square or full mesh SMLT topologies is not supported.

The following rules apply:

- If a VLAN receives traffic from the IST link, it cannot forward on the split multilink trunk link or the edge for the same VLAN.
- If one side of the SMLT link toward the receiver is down, such that the traffic cannot be forwarded directly down the SMLT link from the router on which traffic is ingressing, the IST Peer MUST forward that traffic it receives over the IST link down its side of the SMLT toward the receiver. The decision of whether the IST Peer needs to forward traffic received over the IST to SMLT receivers is made in the datapath, which has full knowledge of the remote SMLT link state.
- Traffic can use the IST to route between VLANs if the forwarding decision for the multicast protocol requires that the other side of the core forwards the multicast traffic (follow the IP multicast routing and forwarding rules for routed traffic). Other VLANs that are not part of SMLT continue to behave in the same way.
- To create a temporary default route pointing to a peer IST, you must enable PIM on the IST VLAN.
- In a scaled multicast environment, if you must reconfigure the members of an MLT link, either SMLT or IST, by removing the ports from the MLT membership list, you must first shutdown the port by using the **shutdown** command at the port configuration level. Let the unicast and multicast traffic subside, and then remove the port from the MLT membership list. If you reconfigure the MLT without first shutting down the port, it can lead to excessive hardware updates to multicast forwarding records and can result in high utilization of the CPU.

  Note:

  In a scaled PIM over Simplified vIST deployment, disabling all the PIM interfaces (**no ip routing**) causes the VLACP ports to bounce. With no user intervention, the packets start getting processed again in approximately 10 seconds. VLACP enables the ports and full functionality is restored.

SMLT provides for fast failover in all cases, but does not provide a functionality similar to Routed SMLT (RSMLT).

Important:

You must enable square SMLT globally before you configure square or full-mesh configurations.

Traffic delay with PIM while restarting peer SMLT switches

If you restart peer SMLT switches, you can lose, or experience a delay in, PIM traffic. The local and remote SMLT links must be up to forward traffic. If a remote SMLT link is down, you can experience a traffic delay.

PIM uses a DR to forward data to receivers on a VLAN. If you restart the DR in an SMLT VLAN, you can lose data because of the following actions:

- If the DR is down, the non-DR switch assumes the role and starts forwarding data.
- After the DR comes back up, it takes priority (higher IP address) to forward data so the non-DR switch stops forwarding data.
- The DR is not ready to forward traffic due to protocol convergence and because it takes time to learn the RP set and create the forwarding path. This situation can result in a traffic delay of 2 to 3 minutes because the DR learns the RP set after Open Shortest Path First (OSPF) converges.

A workaround to this delay is to a configure the static RP router on the peer SMLT switches. This feature avoids the process of selecting an active RP router from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism. After the DR comes back up, traffic resumes as soon as OSPF converges. This workaround reduces the traffic delay to approximately 15 to 65 seconds.

Protocol Independent Multicast-Source Specific Multicast

Source Specific Multicast optimizes PIM-SM by simplifying the many-to-many model. Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM only builds source-based SPTs. Whereas PIM-SM always joins a shared tree first, and then switches to the source tree, SSM eliminates the need to start with a shared tree by immediately joining a source through the SPT. SSM avoids using an RP and RP-based shared trees, which can be a potential problem.

Until now only one channel for one group was allowed to exist in ssm map. From now on multiple channels for the members of the SSM group are allowed to be configured in this map.

This configuration is ideal for applications like television channel distribution and other content-distribution businesses. Banking and trade applications can also use SSM as it provides more control over the hosts receiving and sending data over their networks.

When a v2 report in SSM range is received it is translated to an igmpv3 report message with one group record with type ALLOW and the source lists copied from the igmp ssm map static entries and passed to igmpv3 module. When a v2 leave in SSM range is received it is translated to an igmpv3 report message with one group record with type BLOCK and the source lists copied from the igmp ssm map static entries and passed to igmpv3 module. This behaviour is displayed only when PIM-SSM mode is enabled.

SSM applications use IP addresses reserved by the Internet Assigned Numbers Authority (IANA) in the 232/8 range (232.0.0.0 to 232.255.255.255). SSM recognizes packets in this range and controls the behavior of multicast routing devices and hosts that use these addresses. When a source (S) transmits IP datagrams to an SSM destination address (G), a receiver can receive these datagrams by subscribing to the (S,G) channel.

A channel is a source-group (S,G) pair where S is the source that sends to the multicast group and G is an SSM group address. SSM defines channels on an individual or multiple source basis, which enforces the one-to-many concept of SSM applications. In an SSM channel, each group is associated with multiple sources.

SSM features

PIM-SM requires a unicast protocol to forward multicast traffic within the network to perform the Reverse Path Forwarding (RPF) check. PIM-SM uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree that PIM-enabled routers use to

communicate. The unicast routing table must contain a route to every multicast source in the network as well as routes to PIM entities like the RPs and BSR.

SSM uses only a subset of the PIM-SM features such as the SPT, DR, and some messages (hello, join, prune, and assert). However, some features are unique to SSM. These features, described in the following sections, are extensions of the IGMP and PIM protocols.

PIM-SSM architecture

The following diagram illustrates how the PIM-SSM architecture requires routers to perform the following actions:

- support IGMPv3 source-specific host membership reports and queries at the edge routers
- initiate PIM-SSM (S,G) joins directly and immediately after receiving an IGMPv3 join report from the designated router
- restrict forwarding to SPTs within the SSM address range by all PIM-SSM routers

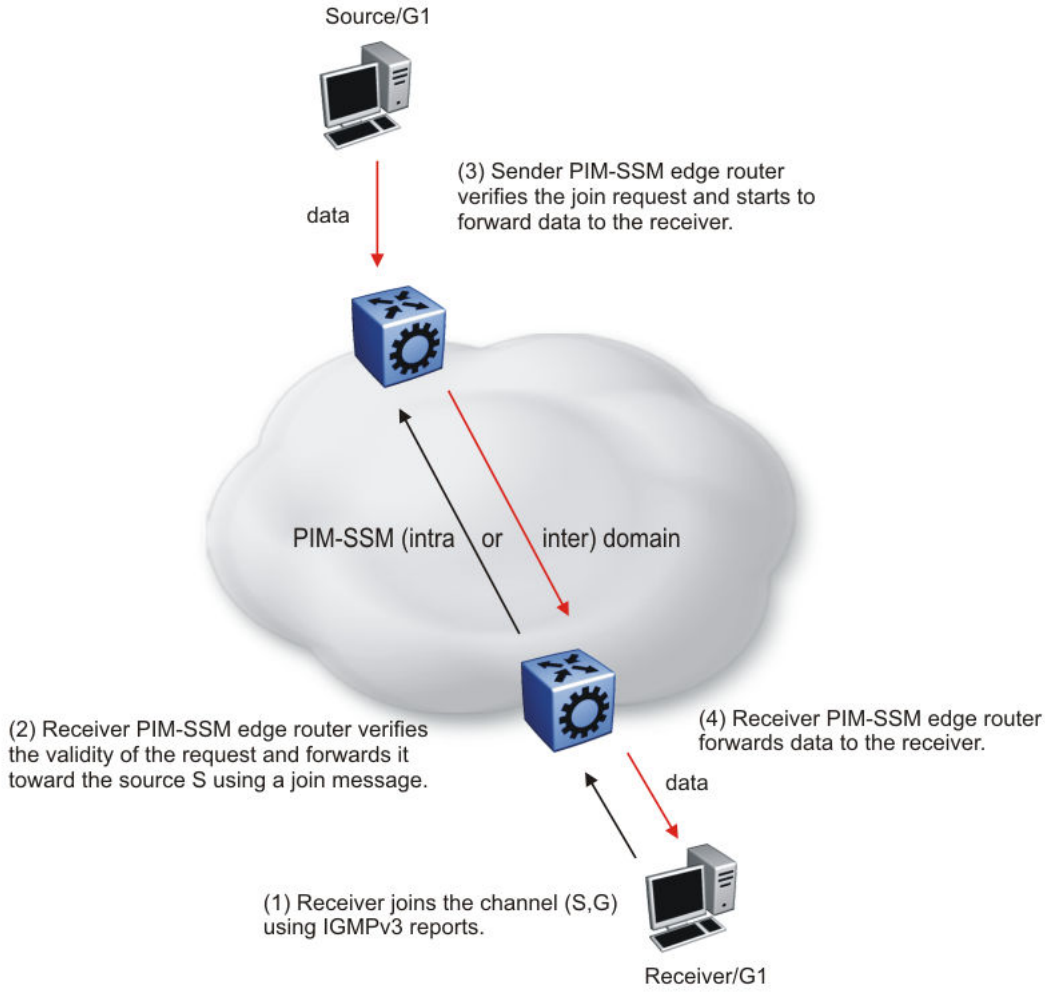

Figure 8: PIM-SSM architecture

The following rules apply to Layer 3 devices with SSM enabled:

- Receive IGMPv3 membership join reports in the SSM range and, if no entry (S,G) exists in the SSM channel table, create one.
- Receive IGMPv2 membership join reports, but only for groups that already use a static (S,G) entry in the SSM channel table.
- Send periodic join messages to maintain a steady SSM tree state.
- Use standard PIM-SM SPT procedures for unicast routing changes, but ignore rules associated with the SPT for the (S,G) route entry.
- Receive prune messages and use standard PIM-SM procedures to remove interfaces from the source tree.
- Forward data packets to interfaces from the downstream neighbors that sent an SSM join, or to interfaces with locally attached SSM group members.

- Drop data packets that do not use an exact-match lookup (S,G) in their forwarding database for S and G.

PIM-SSM static source groups

You can configure static source group entries in the PIM-SSM multicast routing table with static source groups. PIM-SSM cannot prune these entries from the distribution tree. For more information about static source groups, see Static source groups on page 20.

Implementation of SSM and IGMP

The following sections describe how the switch implements PIM-SSM and IGMP.

SSM range

The standard SSM range is 232/8, but you can extend the range to include an IP multicast address. Although you can configure the SSM range, you cannot configure it for all multicast groups (224/4 or 224.0.0.0/240.0.0.0 or 224.0.0.0/255.0.0.0).

You can extend the SSM range to configure existing applications without changing their group configurations.

SSM channel table

You can use the SSM channel to manually configure (S,G) entries that map existing groups to their sending source. These table entries apply to the whole switch, not for each interface, and both IGMPv2 and IGMPv3 hosts use the SSM channel table.

The following rule applies to an SSM channel table for an individual switch:

- You can map one source to multiple groups.
- You can allow multiple sources to the same group.

Important:

Different switches can use different mappings for groups to sources, for example, different channels map differently even if they are on the same network.

SSM and IGMPv2

SSM-configured switches can accept reports from IGMPv2 hosts on IGMPv2 interfaces if the group uses an SSM channel table entry. However, the IGMPv2 host groups must exist in the SSM range defined on the switch, which is 232/8 by default.

- After the SSM switch receives an IGMPv2 report for a group that is in the SSM channel table, it joins the specified source immediately.
- After the SSM switch receives an IGMPv2 report for a group that uses an enabled static SSM channel table entry, it triggers PIM-SSM processing as if it received an equivalent IGMPv3 report.
- After the SSM switch receives an IGMPv2 report for a group out of the SSM range, it processes the report as if it is in PIM-SM mode.

SSM and IGMPv3

The switch supports IGMPv3 for SSM. With IGMPv3, a host can selectively request or filter traffic from sources within the multicast group. IGMPv3 is an interface-level configuration.

Important:

IGMPv3 works without PIM-SSM or SSM-snoop enabled on the interface.

The following rules apply to IGMPv3-enabled interfaces:

- Send only IGMPv3 (source-specific) reports for addresses in the SSM range.
- Accept IGMPv3 reports.
- Drop IGMPv2 reports.

  The IGMPv2 report mentioned in SSM and IGMPv2 on page 50 is processed because it is an IGMPv2 report received on an IGMPv2 interface. If an IGMPv2 interface receives an IGMPv3 report, it drops the report even if PIM-SSM is enabled and the entry is in the SSM channel table. The IGMP versions must match.
- Discard IGMP packets with a group address out of the SSM range.

The switch implements IGMPv3 in one of two modes: dynamic and static.

In dynamic mode, the switch learns about new (S,G) pairs from IGMPv3 reports and adds them to the SSM channel table. If you do not enable dynamic mode and an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report.

In static mode, you can statically configure (S,G) entries in the SSM channel table. If an IGMPv3-enabled interface receives a report that includes a group not listed in the SSM channel table, it ignores the report. The interface also ignores the report if the group is in the table, but the source or mask does not match what is in the table.

Important:

After you enable IGMPv3, changes to the query interval and robustness values on the querier switch propagate to other switches on the same VLAN through IGMP query.

Both IGMPv2 and IGMPv3 hosts use the SSM channel table:

- An IGMPv2 host (with an IGMPv2 VLAN) must use an existing SSM channel entry if the group is in the SSM range.
- If you enable dynamic learning for an IGMPv3 host, the SSM channel automatically learns the group. Otherwise, the SSM channel also needs a static entry.

The following table summarizes how a switch in PIM-SSM mode works with IGMP if you disable IGMPv3 compatibility. In the following table, references to matching a static SSM channel entry assumes that the entry is enabled. If an entry is disabled, it is treated as though it is disallowed.

Table 3: PIM-SSM interaction with IGMPv2 and v3 with IGMPv3 compatibility disabled

Host	VLAN	SSM range	Action
IGMPv2 host	IGMPv3 VLAN	In or out of range	Drop report.
IGMPv3 host	IGMPv2 VLAN	In or out of range	Drop report.

Table continues…

Host	VLAN	SSM range	Action
IGMPv2 host	IGMPv2 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 host	IGMPv2 VLAN	Out of range	Ignore the SSM channel table and process the report as if it is in PIM-SM mode.
IGMPv3 host	IGMPv3 VLAN	Out of range	Process the report.
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic enabled. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and matches an existing SSM channel entry. Create (S,G).
IGMPv3 host	IGMPv3 VLAN	In range	Dynamic disabled and does not match an existing SSM channel entry. Drop report.

The following table summarizes how a switch in PIM-SSM mode works with IGMP if you enable IGMPv3 compatibility.

Table 4: PIM-SSM interaction with IGMPv2 and v3 with IGMPv3 compatibility enabled

Host	VLAN	SSM range	Action
IGMPv2 Host	IGMPv3 VLAN	In range	If the report matches an existing static SSM channel entry, create (S,G). If the report does not match an existing static SSM channel entry, drop it.
IGMPv2 Host	IGMPv3 VLAN	Out of range	Process the report as in PIM-SM mode.

If an IGMPv3 group report enters the VLAN port and the port must discard one or more of the groups in that packet after the application of IGMP access controls, the port drops the entire packet and does not forward it on to other ports of the VLAN.

If an IGMPv3 interface receives an IGMPv2 or v1 query, the interface backs down to IGMPv2 or v1. As a result, the interface flushes all senders and receivers on the interface.

Configuration limitations

Run PIM-SSM on either all switches in the domain or only on the edge routers. If you use a mix of PIM-SSM and PIM-SM switches in the domain, run PIM-SSM on all the edge routers and run PIM-SM on all the core routers.

Important:

A PIM domain with edge routers running PIM-SM and core routers running PIM-SSM does not operate properly. If you prefer or require a mixed PIM-SM and PIM-SSM topology, run PIM-SSM on the edge switches and PIM-SM in the core. Ensure a valid RP configuration exists for groups that exist outside of the SSM range. If a valid RP configuration exists, the SSM switches process the joins in SM mode. If no RP exists, the SSM switches drop the reports.

Static source groups cannot conflict with SSM channels. If you configure a static source group or an SSM channel, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) to different sources or multiple groups to a single source for both static source group and an SSM channel.

PIM passive interfaces

You can configure the PIM interface as active or passive. The default is active. With an active interface, you can configure transmit and receive PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you use a high number of PIM interfaces and these interfaces connect to end users, not to other switches.

A PIM passive interface does not transmit and drops messages of the following type:

- hello
- join
- prune
- register
- register-stop
- assert
- candidate-RP-advertisement
- bootstrap

If a PIM passive interface receives these types of messages, it drops them and the switch logs a message, detailing the type of protocol message and the IP address of the sending device. These log messages help to identify the device that performs routing on the interface, which is useful if you must disable a device that does not operate correctly.

Important:

A device can send register and register-stop messages to a PIM passive interface, but these messages cannot be sent out of that interface.

The PIM passive interface maintains information about hosts, through IGMP, that are related to senders and receivers, but the interface does not maintain information about PIM neighbors. You can configure a BSR or an RP on a PIM passive interface.

You can also use the PIM passive interface feature as a security measure to prevent routing devices from becoming attached and participating in the multicast routing of the network.

Important:

Before you change the state (active or passive) of a PIM interface, disable PIM on that interface. This action prevents instability in the PIM operations, especially when neighbors exists or the interface receives streams. After you disable PIM, the switch loses traffic for approximately 80 seconds.

Multicast route statistics

Note:

This feature does not apply to all hardware platforms. For more information about feature support, see Release Notes for VSP 8600.

The multicast route statistics feature provides statistics for multicast streams through the switch. Using the Command Line Interface (CLI), Simple Network Management Protocol (SNMP) or Enterprise Device Manager (EDM), you can track the number of senders sending multicast streams to a particular group address. You can also obtain a count of the packets or bytes being received for a particular multicast group address and the average size of the frames. Multicast route statistics are supported for both IPv4 and IPv6 group addresses.

Determining the route statistics is especially useful when debugging a multicast network and also when administering the network.

Multicast route statistics and DvR

When you enable or clear IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

IP multicast network design

Use multicast routing protocols to efficiently distribute a single data source among multiple users in the network. This section provides information about how to design networks that support IP multicast routing.

For more design guidelines, conceptual, and configuration information about IP Multicast over Fabric Connect, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Multicast scalability design rules

The following section lists the design rules to increase multicast route scaling.

Important:

The switch does not support High Availability (HA).

The switch software supports the following:

- Protocol-Independent Multicast (PIM)
- Split MultiLink Trunking (SMLT) and Routed-SMLT (RSMLT)

Multicast scalability design rules

1. Whenever possible, use simple network designs that do not use VLANs that span several switches. Instead, use routed links to connect switches.
2. Whenever possible, group sources sending to the same group in the same subnet. The switch uses a single egress forwarding pointer for all sources in the same subnet sending to the same group. Be aware that these streams have separate hardware forwarding records on the ingress side.
3. Do not configure multicast routing on edge switch interfaces that do not contain multicast senders or receivers. By following this rule, you:
   - Provide secure control over multicast traffic that enters or exits the interface.
   - Reduce the load on the switch, as well as the number of routes. This improves overall performance and scalability.
4. Avoid initializing many (several hundred) multicast streams simultaneously. Initial stream setup is a resource-intensive task, and initializing a large number can increase the setup time. In some cases, this delay can result in stream loss.
5. Whenever possible, do not connect IP multicast sources and receivers by using VLANs that interconnect switches (see the following figure). In some cases, this can result in excessive hardware record use. By placing the source on the interconnected VLAN, traffic takes two paths to the destination, depending on the reverse path forwarding (RPF) checks and the shortest path to the source.

   For example, if a receiver is on VLAN 1 on switch S1 and another receiver is on VLAN 2 on switch S1, traffic can be received from two different paths to the two receivers, which results in the use of two forwarding records. If the source on switch S2 is on a different VLAN than VLAN 3, traffic takes a single path to switch S1 where the receivers are located.

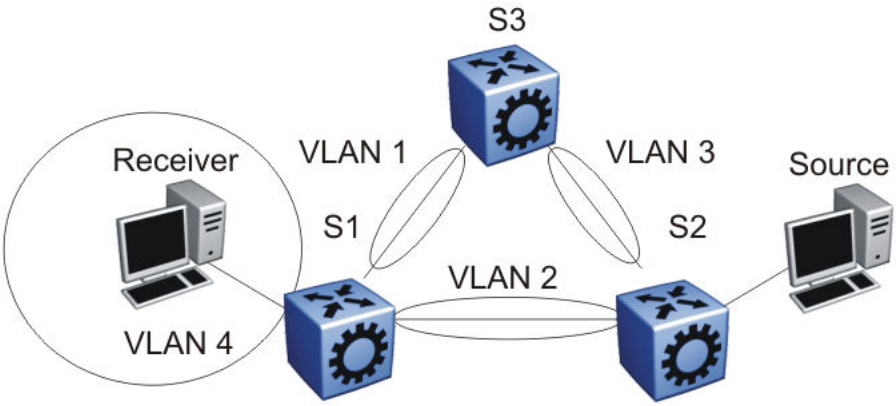

Figure 9: IP multicast sources and receivers on interconnected VLANs

IP multicast address range restrictions

IP multicast routers use D class addresses, which range from 224.0.0.0 to 239.255.255.255. Although you can use subnet masks to configure IP multicast address ranges, the concept of subnets does not exist for multicast group addresses. Consequently, the usual unicast conventions —where you reserve the all 0s subnets, all 1s subnets, all 0s host addresses, and all 1s host addresses—do not apply.

Internet Assigned Numbers Authority (IANA) reserves addresses from 224.0.0.0 through 224.0.0.255 for link-local network applications. Multicast-capable routers do not forward packets with an address in this range. For example, Open Shortest Path First (OSPF) uses 224.0.0.5 and 224.0.0.6, and Virtual Router Redundancy Protocol (VRRP) uses 224.0.0.18 to communicate across local broadcast network segments.

IANA also reserves the range of 224.0.1.0 through 224.0.1.255 for well-known applications. IANA assigns these addresses to specific network applications. For example, the Network Time Protocol (NTP) uses 224.0.1.1, and Mtrace uses 224.0.1.32. RFC1700 contains a complete list of these reserved addresses.

Multicast addresses in the 232.0.0.0/8 (232.0.0.0 to 232.255.255.255) range are reserved only for source-specific multicast (SSM) applications, such as one-to-many applications. While this range is the publicly reserved range for SSM applications, private networks can use other address ranges for SSM.

Finally, addresses in the range 239.0.0.0/8 (239.0.0.0 to 239.255.255.255) are administratively scoped addresses; they are reserved for use in private domains. Do not advertise these addresses outside the private domain. This multicast range is analogous to the 10.0.0.0/8, 172.16.0.0/20, and 192.168.0.0/16 private address ranges in the unicast IP space.

In a private network, only assign multicast addresses from 224.0.2.0 through 238.255.255.255 to applications that are publicly accessible on the Internet. Assign addresses in the 239.0.0.0/8 range to multicast applications that are not publicly accessible.

Although you can use a multicast address you choose on your own private network, it is generally not good design practice to allocate public addresses to private network entities. Do not use public addresses for unicast host or multicast group addresses on private networks.

Multicast MAC address mapping considerations

Like IP, Ethernet has a range of multicast MAC addresses that natively support Layer 2 multicast capabilities. While IP has a total of 28 addressing bits available for multicast addresses, Ethernet has only 23 addressing bits assigned to IP multicast. The Ethernet multicast MAC address space is much larger than 23 bits, but only a subrange of that larger space is allocated to IP multicast. Because of this difference, 32 IP multicast addresses map to one Ethernet multicast MAC address.

IP multicast addresses map to Ethernet multicast MAC addresses by placing the low-order 23 bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01:00:5E:00:00:00. Thus, more than one multicast address maps to the same Ethernet address (see the following figure). For example, all 32 addresses 224.1.1.1, 224.129.1.1, 225.1.1.1, 225.129.1.1, 239.1.1.1, 239.129.1.1 map to the same 01:00:5E:01:01:01 multicast MAC address.

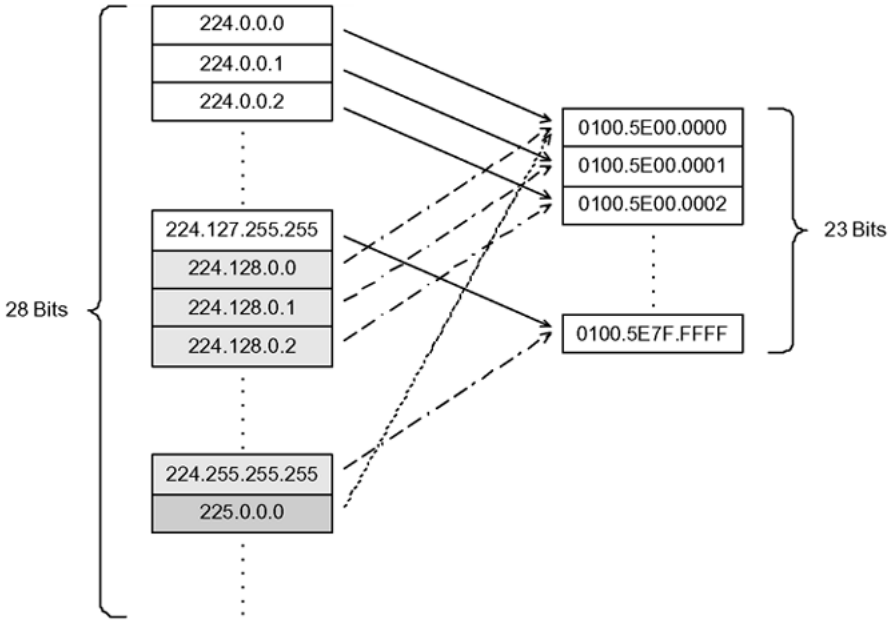

Figure 10: Multicast IP address to MAC address mapping

Most Ethernet switches handle Ethernet multicast by mapping a multicast MAC address to multiple switch ports in the MAC address table. Therefore, when you design the group addresses for multicast applications, take care to efficiently distribute streams only to hosts that are receivers.

The VSP 4000 Series devices switch IP multicast data based on the IP multicast address, not the MAC address, and thus, do not have this issue.

As an example, consider two active multicast streams using addresses 239.1.1.1 and 239.129.1.1. Suppose that two Ethernet hosts, receiver A and receiver B, connect to ports on the same switch and only want the stream addressed to 239.1.1.1. Suppose also that two other Ethernet hosts,

receiver C and receiver D, also connect to the ports on the same switch as receiver A and B, and want to receive the stream addressed to 239.129.1.1. If the switch uses the Ethernet multicast MAC address to make forwarding decisions, then all four receivers receive both streams—even though each host only wants one stream. This transmission increases the load on both the hosts and the switch. To avoid this extra load, ensure that you manage the IP multicast group addresses used on the network.

The VSP 4000 Series switches do not forward IP multicast packets based on multicast MAC addresses—even when bridging VLANs at Layer 2. Thus, the platform does not encounter this problem. Instead, the platform internally maps IP multicast group addresses to the ports that contain group members.

When an IP multicast packet is received, the lookup is based on the IP group address, regardless of whether the VLAN is bridged or routed. While the problem described in the previous example does not affect the VSP 4000 Series switches, other switches in the network may be affected. This problem is particularly true of pure Layer 2 switches.

In a network that includes multiple hardware platforms, the easiest way to ensure that this issue does not arise is to use only a consecutive range of IP multicast addresses that correspond to the lower-order 23 bits of that range. For example, use an address range from 239.0.0.0 through 239.127.255.255. A group address range of this size can still easily accommodate the needs of even the largest private enterprise.

Dynamic multicast configuration changes

You must not perform dynamic multicast configuration changes when multicast streams flow in a network. For example, do not change the routing protocol that runs on an interface, or the IP address, or the subnet mask for an interface until multicast traffic ceases.

For such changes, ensure that you temporarily stop all multicast traffic. If the changes are necessary and you have no control over the applications that send multicast data, you can disable the multicast routing protocols before you perform the change. For example, consider disabling multicast routing before making interface address changes. In all cases, these changes result in traffic interruptions because they affect neighbor-state machines and stream-state machines.

In addition, when removing port members of an MLT group you must first disable the ports. Changing the group set without first shutting the ports down can result in high-CPU utilization and processing in a scaled multicast environment due to the necessary hardware reprogramming on the multicast records.

IGMPv3 backward compatibility

IGMPv3 for PIM is backward compatible with IGMPv1/v2. According to RFC3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.

- If a router that is not explicitly configured to use IGMPv1 or IGMPv2, detects an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure the IGMP version of an interface to version 3 regardless of the PIM or snooping mode.

You can configure whether the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

Note:

If you enable the explicit host tracking option on an IGMPv3 interface, you cannot downgrade to IGMPv1 or IGMPv2. You must disable explicit host tracking to downgrade the IGMP version.

TTL in IP multicast packets

The switch treats multicast data packets with a time-to-live (TTL) of 1 as expired packets and sends them to the CPU before dropping them. To avoid this issue, ensure that the originating application uses a hop count large enough to enable the multicast stream to traverse the network and reach all destinations without reaching a TTL of 1. Ensure that you use a TTL value of 33 or 34 to minimize the effect of looping in an unstable network.

Multicast MAC filtering

Certain network applications require multiple hosts to share a multicast MAC address. Instead of flooding all ports in the VLAN with this multicast traffic, you can use Multicast MAC filtering to forward traffic to a configured subset of the ports in the VLAN. This multicast MAC address is not an IP multicast MAC address.

At a minimum, map the multicast MAC address to a set of ports within the VLAN. In addition, if traffic is routed on the local host, you must configure an Address Resolution Protocol (ARP) entry to map the shared unicast IP address to the shared multicast MAC address. You must configure an ARP entry because the hosts can also share a virtual IP address, and packets addressed to the virtual IP address need to reach each host.

Ensure that you limit the number of such configured multicast MAC addresses to a maximum of 100. This number is related to the maximum number of possible VLANs you can configure, because for every multicast MAC filter that you configure the maximum number of configurable VLANs reduces by one. Similarly, configuring large numbers of VLANs reduces the maximum number of configurable multicast MAC filters downward from 100.

Although you can configure addresses starting with 01.00.5E, which are reserved for IP multicast address mapping, do not enable IP multicast with streams that match the configured addresses. This configuration can result in incorrect IP multicast forwarding and incorrect multicast MAC filtering.

Guidelines for multicast access policies

Use the following guidelines when you configure multicast access policies:

- Use masks to specify a range of hosts. For example, 10.177.10.8 with a mask of 255.255.255.248 matches hosts addresses 10.177.10.8 through 10.177.10.15. The host subnet address and the host mask must be equal to the host subnet address. An easy way to determine this is to ensure that the mask has an equal or fewer number of trailing zeros than the host subnet address. For example, 3.3.0.0/255.255.0.0 and 3.3.0.0/255.255.255.0 are valid. However, 3.3.0.0/255.0.0.0 is not.
- Apply receive-access policies to all eligible receivers on a segment. Otherwise, one host joining a group makes that multicast stream available to all.
- Receive access policies are initiated after the switch receives reports with addresses that match the filter criteria.
- Transmit access policies apply after the switch receives the first packet of a multicast stream.

Multicast access policies can apply to a routed PIM interface if Internet Group Management Protocol (IGMP) reports the reception of multicast traffic.

The following rules and limitations apply to IGMP access policy parameters when you use them with IGMP instead of PIM:

- The static member parameter applies to IGMP snooping and PIM on both interconnected links and edge ports.
- The Static Not Allowed to Join parameter applies to IGMP snooping and PIM on both interconnected links and edge ports.
- For multicast access control, the denyRx parameter applies to IGMP snooping and PIM. The DenyTx and DenyBoth parameters apply only to IGMP snooping.

Split-subnet and multicast

The split-subnet issue arises when you divide a subnet into two unconnected sections in a network. This division results in the production of erroneous routing information about how to reach the hosts on that subnet. The split-subnet problem applies to all types of traffic, but it has a larger impact on a PIM-SM network.

To avoid the split-subnet problem in PIM networks, ensure that the RP router is not in a subnet that can become a split subnet. Also, avoid having receivers on this subnet. Because the RP is an entity that must be reached by all PIM-enabled switches with receivers in a network, placing the RP on a split-subnet can impact the whole multicast traffic flow. Traffic can be affected even for receivers and senders that are not part of the split-subnet.

Protocol Independent Multicast-Sparse Mode guidelines

Protocol Independent Multicast-Sparse Mode (PIM-SM) uses an underlying unicast routing information base to perform multicast routing. PIM-SM builds unidirectional shared trees rooted at a RP router for each group and can also create shortest-path trees for each source.

PIM-SM and PIM-SSM Scalability

For more information on interface scaling, see the Release Notes for VSP 8600.

The software does not support virtualized PIM. PIM is supported in the Global Routing Table only.

Interfaces that run PIM must also use a unicast routing protocol (PIM uses the unicast routing table), which puts stringent requirements on the system. With a high number of interfaces, take special care to reduce the load on the system.

Use few active IP routed interfaces. You can use IP forwarding without a routing protocol enabled on the interfaces, and enable only one or two with a routing protocol. You can configure proper routing by using IP routing policies to announce and accept routes on the switch. Use PIM passive interfaces on the majority of interfaces.

Important:

For information on the maximum values for total PIM interfaces and active interfaces, see the Release Notes for VSP 8600. If you configure the maximum number of active interfaces, all remaining interfaces must be passive.

When you use PIM-SM, the number of routes can scale up to the unicast route limit because PIM uses the unicast routing table to make forwarding decisions. For higher route scaling, use OSPF instead of Routing Information Protocol (RIP).

As a general rule, a well-designed network does not have many routes in the routing table. For PIM to work properly, ensure that all subnets configured with PIM are reachable and that PIM uses the information in the unicast routing table. For the RPF check, to correctly reach the source of any multicast traffic, PIM requires the unicast routing table.

PIM General Requirements

Design simple PIM networks where VLANs do not span several switches.

PIM relies on unicast routing protocols to perform its multicast forwarding. As a result, include in your PIM network design, a unicast design where the unicast routing table has a route to every source and receiver of multicast traffic, as well as a route to the RP router and Bootstrap router (BSR) in the network. Ensure that the path between a sender and receiver contains PIM-enabled interfaces. Receiver subnets are not always required in the routing table.

Use the following guidelines:

- Ensure that every PIM-SM domain is configured with an RP, either by static definition or via BSR.
- Ensure that every group address used in multicast applications has an RP in the network.

- As a redundancy option, you can configure several RPs for the same group in a PIM domain.
- As a load sharing option, you can have several RPs in a PIM-SM domain map to different groups.
- In order to configure an RP to cover the entire multicast range, configure an RP to use the IP address of 224.0.0.0 and the mask of 240.0.0.0.
- Configure an RP to handle a range of multicast groups by using the mask parameter. For example, an entry for group value of 224.1.1.0 with a mask of 255.255.255.192 covers groups 224.1.1.0 to 224.1.1.63.
- In a PIM domain with both static and dynamic RP switches, you cannot configure one of the (local) interfaces for the static RP switches as the RP. For example, in the following scenario:

  (static RP switch) Sw1 ------ Sw2 (BSR/Cand-RP1) -----Sw3

  You cannot configure one of the interfaces on switch Sw1 as static RP because the BSR cannot learn this information and propagate it to Sw2 and Sw3. PIM requires that you consistently configure RP on all the routers of the PIM domain, so you can only add the remote interface Candidate-RP1 (Cand-RP) to the static RP table on Sw1.
- If a switch needs to learn an RP-set, and has a unicast route to reach the BSR through this switch, you cannot enable or configure static RP on a switch in a mixed mode of candidate RP and static RP switches. For examples, see the following two figures.

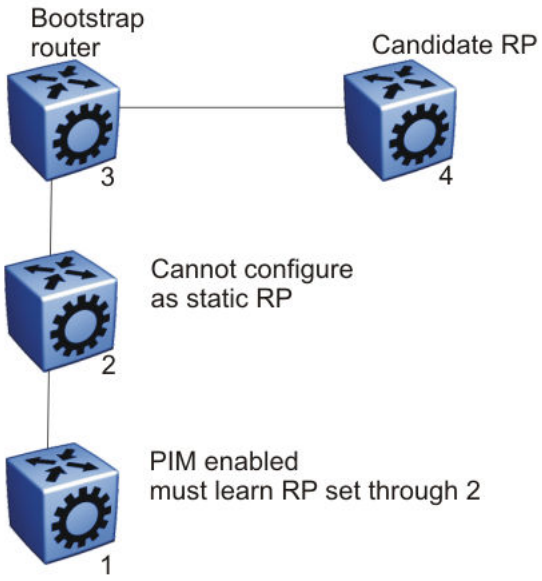

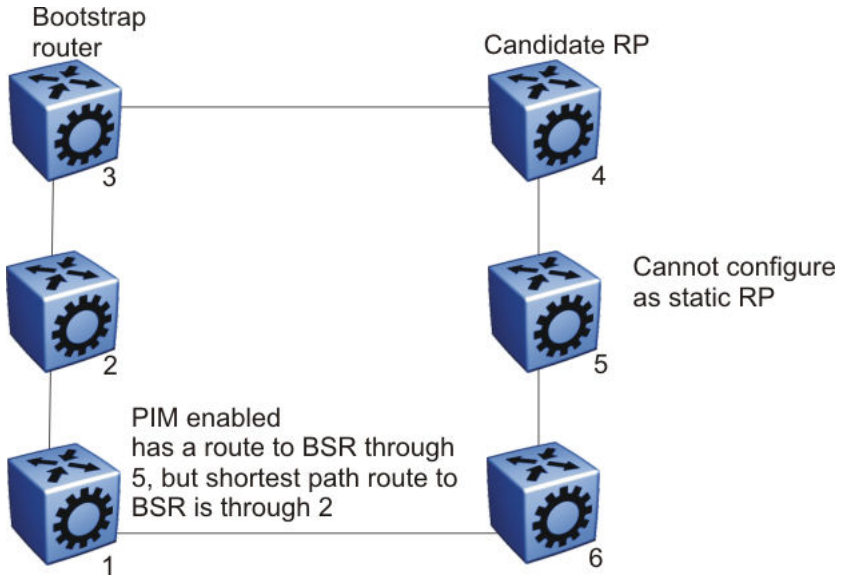

PIM and Shortest Path Tree Switchover

When an IGMP receiver joins a multicast group, PIM on the leaf router first joins the shared tree. After the first packet is received on the shared tree, the router uses the source address information in the packet to immediately switch over to the shortest path tree (SPT).

To guarantee a simple, yet high-performance implementation of PIM-SM, the switch does not support a threshold bit rate in relation to SPT switchover. Intermediate routers (that is, not directly connected IGMP hosts) do not switch over to the SPT until directed to do so by the leaf routers.

Other vendors can offer a configurable threshold, such as a certain bit rate at which the SPT switch-over occurs. Regardless of their implementation, no interoperability issues with the switch result. Switching to and from the shared and shortest path trees is independently controlled by each downstream router. Upstream routers relay joins and prunes upstream hop-by-hop, building the desired tree as they go. Because a PIM-SM compatible router already supports shared and shortest path trees, no compatibility issues arise from the implementation of configurable switchover thresholds.

PIM Traffic Delay and SMLT Peer Reboot

PIM uses a designated router (DR) to forward data to receivers on the DR VLAN. The DR is the router with the highest IP address on a LAN. If this router is down, the router with the next highest IP address becomes the DR. However, if the VLAN is an SMLT VLAN, the DR is not a factor in determining which switch forwards the data down to the receiver. Either aggregate switch can forward data to the receiver, because the switches act as one. The switch that forwards depends on where the source is located (on another SMLT/vIST link or on a non-SMLT/non-vIST link) and whether either side of the receiver SMLT link is up or down. If the forwarder switch is rebooted, traffic loss occurs until protocol convergence is completed.

Consider the following cases:

- If the source is on an SMLT link that is not the receiver SMLT, the switch that directly received the data on its side of the source SMLT link forwards it down to the receiver on the receiver SMLT regardless of which switch is the DR for the receiver VLAN. The forwarding switch sends a copy of the data over the vIST link to the peer switch, which drops the data because it knows that the remote SMLT is up and therefore the remote peer has already forwarded the data. If the forwarding switch goes down, the other switch receives the data directly over its source SMLT link and takes over forwarding to the receivers. After the original switch comes back up, the original switch again receives the data directly over its source SMLT. The original switch may not be ready to forward the data because of the protocol reconvergence, so the original switch loses traffic until reconvergence is complete.
- If the source is not learned on another SMLT link or the vIST link on each aggregate switch; they have a route to the source which is not on an SMLT or across the vIST. The switches must choose which one forwards the data down the receiver SMLT link; which one is the designated forwarder, so that duplicate data does not occur. The highest IP address is the designated forwarder. If the designated forwarder becomes disabled, the other takes over. When it is reenabled, the other switch sees that it is no longer the highest IP address and it sees that the remote SMLT link comes up. The other switch then assumes that the vIST peer is capable of being the designated forwarder and it stops forwarding down to the receivers. If the original switch is not ready to forward the data due to reconvergence, traffic loss occurs.

In either case, configuring a static RP helps the situation. To avoid this traffic delay, a workaround is to configure a static RP on the peer SMLT switches. This configuration avoids the process of selecting an active RP router from the list of candidate RPs, and also of dynamically learning about RPs through the BSR mechanism. Then, when the DR comes back, traffic resumes as soon as OSPF converges. This workaround reduces the traffic delay.

Circuitless IP for PIM-SM

Use CLIP to configure a resilient RP and BSR for a PIM network. When you configure an RP or BSR on a regular interface, if it becomes nonoperational, the RP and BSR also become nonoperational. This status results in the election of other redundant RPs and BSRs, and can disrupt IP multicast traffic flow in the network. As a best practice for multicast networks design, always configure the RP and BSR on a CLIP interface to prevent a single interface failure from causing these entities to fail.

Also, configure redundant RPs and BSRs on different switches such that these entities are on CLIP interfaces. For the successful setup of multicast streams, ensure that a unicast route exists to all CLIP interfaces from all locations in the network. A unicast route is mandatory because, for proper RP learning and stream setup on the shared RP tree, every switch in the network needs to reach the RP and BSR. You can use PIM-SM CLIP interfaces only for RP and BSR configurations, and are not intended for other purposes.

It is not recommended to have non-SMLT IGMP leaf ports on a router configured to be one of the redundant RP CLIP devices. It is possible that these IGMP hosts can become isolated from the multicast data stream(s).

If you configure dual-redundant RPs (vIST peers with the same CLIP interface IP address used for the RP), the topology in the following figure does not work in link-failure scenarios. Use caution if

you design a network with this topology where the vIST peers are PIM enabled, and the source and receiver edges are Layer 2.

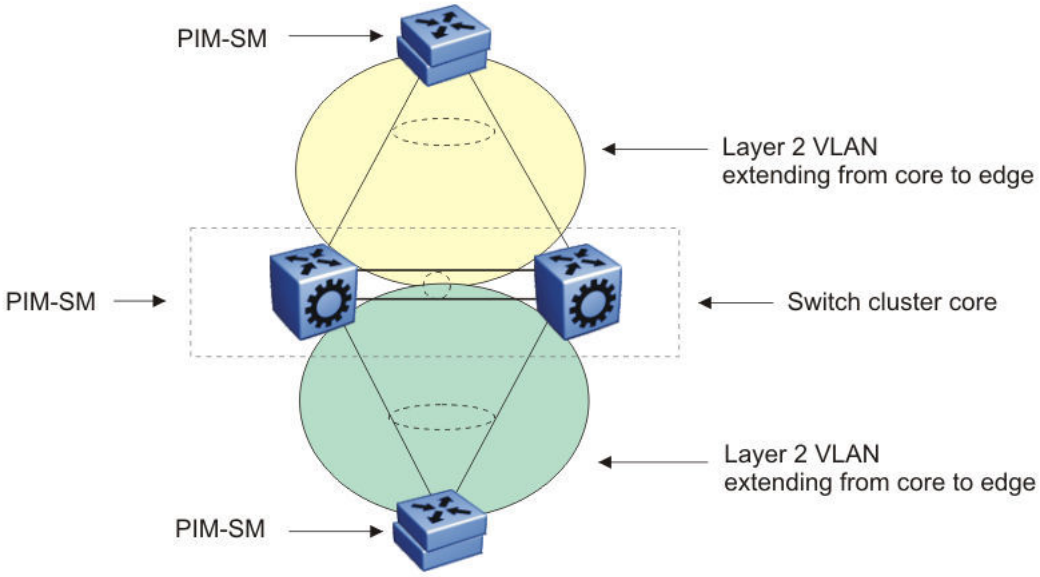

Consider an example where one of the peers, vIST-A, is the PIM DR for the source VLAN, and the source data is hashed to vIST-A from the Layer 2 source edge. vIST-A forwards traffic to the receiver edge using the SMLT link from vIST-A to the receiver edge. If the SMLT link fails, vIST-A does not forward traffic over the vIST link to vIST-B, and the receiver edge does receive the data.

In this topology, the receiver edge sends an IGMP membership report for a group, which is recorded on both vIST peers as an IGMP LEAF on the receiver SMLT port on the receiver VLAN.

Because both of the vIST peers are the RP for the group, they do not send a (*,g) PIM JOIN message toward the other RP. The (*,g) PIM mroute does not record the vIST port as a JOIN port on either vIST device. The PIM (*,g) mroute records only a LEAF on the SMLT receiver port.

Because the source is local (Layer 2 edge), there is no PIM (s,g) JOIN message toward the source and the (s,g) PIM mroute does not record the vIST port as a JOIN port on either vIST device. The PIM (s,g) mroute records only a LEAF on the SMLT receiver port.

If the source is hashed to vIST-A, the PIM DR for the incoming VLAN, traffic is forwarded to the receiver correctly. vIST-A does not forward traffic over the vIST to vIST-B, because no JOIN exists on the vIST port. If the receiver SMLT link from the vIST-A peer is down, the traffic is not forwarded to vIST-B, and is not received by the receiver edge. Traffic resumes after the link is restored. If the source data hashes to the non-DR peer, vIST-B, no problem occurs because the non-DR always forwards traffic to the DR.

A similar situation exists in this topology when vIST-A is both the RP and the DR for the Layer 2 receiver edge. The vIST port is not in the outgoing port list because there is no JOIN message from the peer toward the source (which is not PIM enabled). Therefore, if the SMLT link from vIST-A to the receiver edge is down, the system does not forward traffic to the peer vIST-B and down to the receiver.

You can avoid the preceding problems with this topology by performing one of the following actions:

- Enable PIM on the source edge.

  The vIST peers send PIM joins toward the source and the JOIN is recorded on the vIST port for the (s,g). Data is forwarded to the peer.
- Do not configure dual redundant RPs.

  One vIST peer is the RP for a group.
- Do not configure one vIST peer as both the DR for the source VLAN and the RP for the receiver group.

  The system forwards the traffic to the RP or to the DR, depending on which peer receives the source, and, if the SMLT link to the receiver goes down there will be no data loss.

PIM-SM and Static RP

Use static RP to provide security, interoperability, and redundancy for PIM-SM multicast networks. Consider if the administrative ease derived from using dynamic RP assignment is worth the security risks involved. For example, if an unauthorized user connects a PIM-SM router that advertises itself as a candidate RP (C-RP), it can possibly take over new multicast streams that otherwise distribute through an authorized RP. If security is important, use static RP assignment.

You can use the static RP feature in a PIM environment with devices that run legacy PIM-SMv1 and Cisco Auto-RP. For faster convergence, you can also use static RP in a PIM-SMv2 environment. If you configure static RP with PIM-SMv2, the BSR is not active.

Static RP and Auto-RP

Some legacy PIM-SMv1 networks use the auto-RP protocol. Auto-RP is a Cisco proprietary protocol that provides equivalent functionality to the legacy platform supported PIM-SM RP and BSR. You can use the static RP feature to interoperate in this environment. For example, in a mixed-vendor network, you can use auto-RP among routers that support the protocol, while other routers use static RP. In such a network, ensure that the static RP configuration mimics the information that is dynamically distributed to guarantee that multicast traffic is delivered to all parts of the network.

In a mixed auto-RP and static RP network, ensure that the legacy platform does not serve as an RP because it does not support the auto-RP protocol. In this type of network, the RP must support the auto-RP protocol.

Static RP and RP Redundancy

You can provide RP redundancy through static RPs. To ensure consistency of RP selection, implement the same static RP configuration on all PIM-SM routers in the network. In a mixed vendor network, ensure that the same RP selection criteria is used among all routers. For example, to select the active RP for each group address, the switch uses a hash algorithm defined in the PIM-SMv2 standard. If a router from another vendor selects the active RP based on the lowest IP address, then the inconsistency prevents stream delivery to certain routers in the network.

If a group address-to-RP discrepancy occurs among PIM-SM routers, network outages occur. Routers that are unaware of the true RP cannot join the shared tree and cannot receive the multicast stream.

Failure detection of the active RP is determined by the unicast routing table. As long as the RP is considered reachable from a unicast routing perspective, the local router assumes that the RP is fully functional and attempts to join the shared tree of that RP.

The following figure shows a hierarchical OSPF network where a receiver is in a totally stubby area. If RP B fails, PIM-SM router A does not switch over to RP C because the injected default route in the unicast routing table indicates that RP B is still reachable.

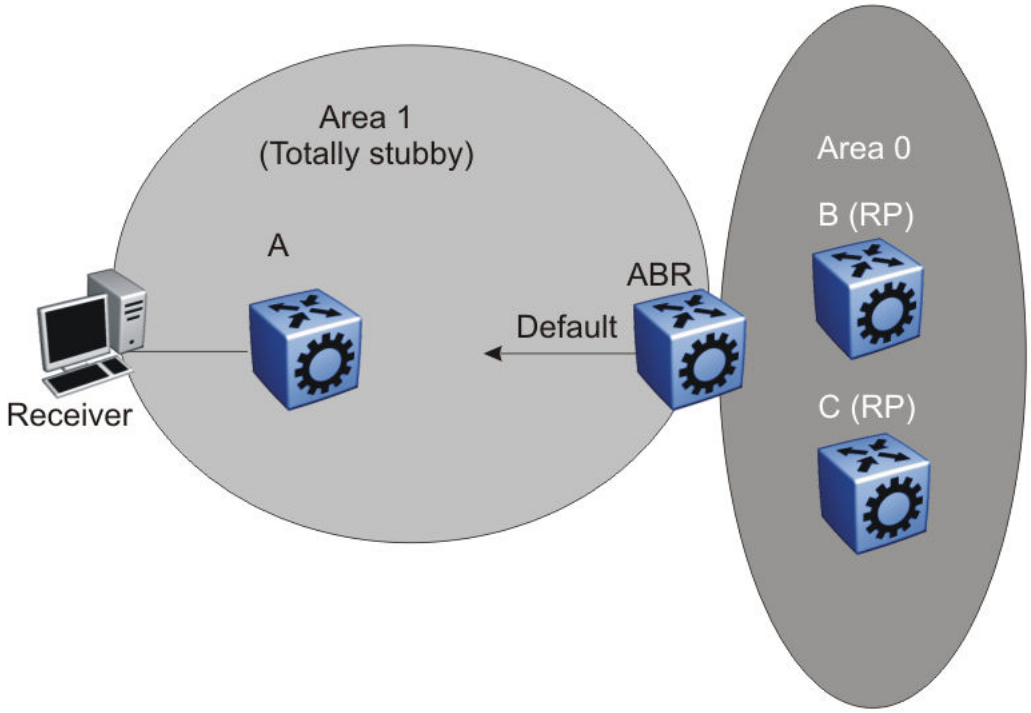

Because failover is determined by unicast routing behavior, carefully consider the unicast routing design, as well as the IP address you select for the RP. Static RP failover performance depends on the convergence time of the unicast routing protocol. For quick convergence, ensure that you use a link state protocol, such as OSPF. For example, if you use RIP as the routing protocol, an RP failure can take minutes to detect. Depending on the application, this situation can be unacceptable.

Static RP failover time does not affect routers that have already switched over to the SPT; failover time only affects newly-joining routers.

Unsupported Static RP Configurations

If you use static RP, you disable dynamic RP learning. The following figure shows an unsupported configuration for static RP. In this example because of inter-operation between static RP and dynamic RP, no RP exists at switch 2. However, (S,G) creation and deletion occurs every 210 seconds at switch 16.

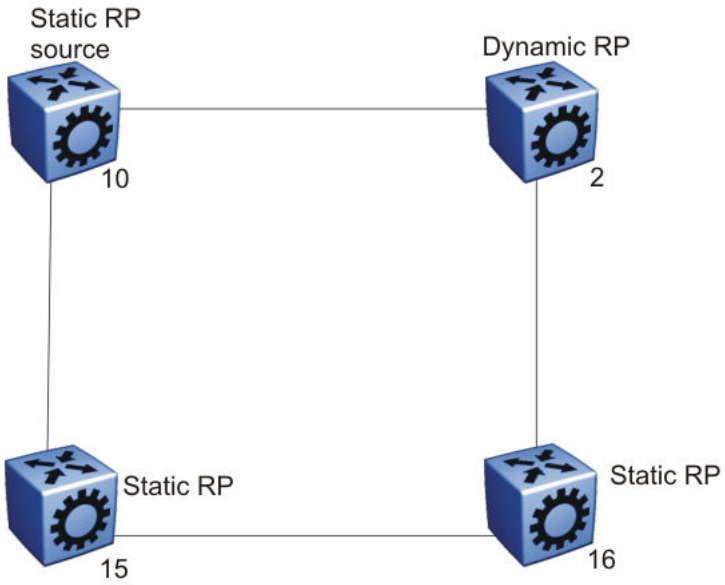

Switches 10, 15, and 16 use static RP, whereas switch 2 uses dynamic RP. The source is at switch 10, and the receivers are switches 15 and 16. The RP is at switch 15 locally. The receiver on switch 16 cannot receive packets because its SPT goes through switch 2.

Switch 2 is in a dynamic RP domain, so it cannot learn about the RP on switch 15. However, (S, G) records are created and deleted on switch 16 every 210 seconds.

Rendezvous Point Router Considerations

You can place an RP on a switch when VLANs extend over several switches. However, when you use PIM-SM, ensure that you do not span VLANs on more than two switches.

Use static group-range-to-RP mappings in an SMLT topology as opposed to RP set learning using the Bootstrap Router (BSR) mechanism. Static RP allows for faster convergence in box failure, reset and HA failover scenarios, whereas there are inherent delays in the BSR mechanism as follows:

- When a router comes back up after a failover or reset, to accept and propagate (*,g) join requests from surrounding routers (either PIM join messages or local IGMP membership reports) to the RP, a PIM router must determine the address of the RP for each group for which they desire (*,g) state. The PIM router must know the unicast route to the RP address. The route to the RP address is learned by using a unicast routing protocol such as OSPF, and the RP address is either statically configured or dynamically learned using the BSR mechanism.
- When a box comes up after a reset, if the RP is not statically configured, it must wait for the BSR to select the RP from candidate RP routers, and then propagate the RP set hop-by-hop to all PIM routers. This must be done before a join message can be processed. If the PIM router receives a join message before it learns the RP set, it drops the join message, and the router waits for another join or prune message to arrive before it creates the multicast route and propagates the join message to the RP. The default Join/Prune timer is 60 seconds, and because of this and the delays inherent in BSR RP-set learning, significant multicast traffic

interruptions can occur. If the RP is statically configured, the only delay is in the unicast routing table convergence and the arrival of the Join/Prune messages from surrounding boxes.

Layer 3 Multicast Extended VLANS

Avoid using a Layer 3 multicast extended VLAN topology without SMLT.

Do not connect non-SMLT PIM routers in a linear fashion on the same VLAN. This topology is called an extended VLAN. Unlike a shared VLAN topology where all routers on the same VLAN are physically one hop away from each other, a VLAN router at one end of the extended VLAN has one or more routers in between it and the router at the far end of the extended VLAN. The following figure shows an extended VLAN.

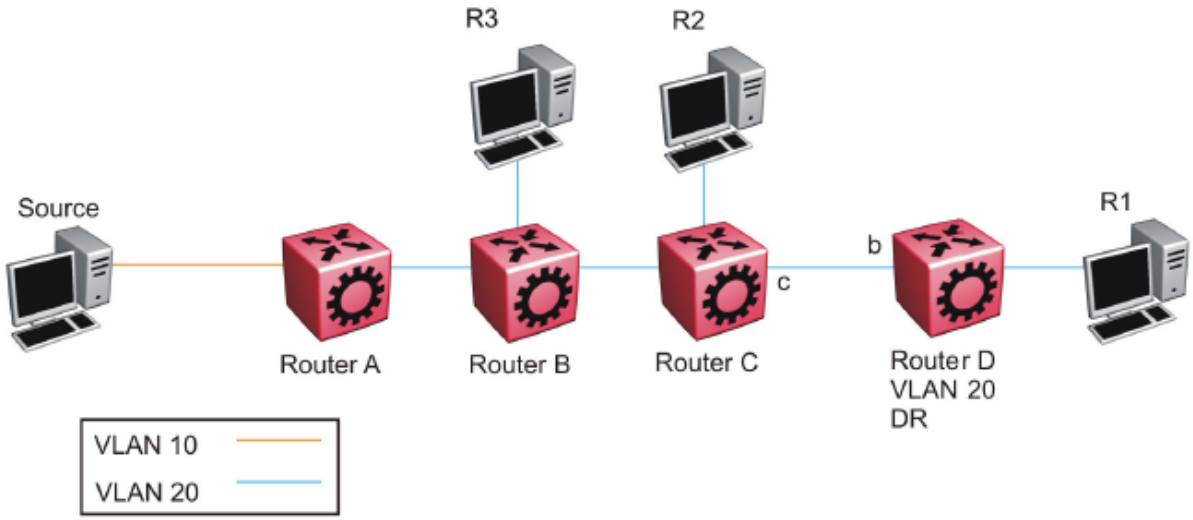

In the preceding figure, all routers use PIM-SSM. The source connects to Router A on VLAN 10. All routers and receiver hosts connect on the same extended VLAN, VLAN 20. All routers have a receiver in VLAN 20. Router D is the PIM DR for VLAN 20 and the source host is not on VLAN 20. PIM-SSM does not require a Rendezvous Point (RP).

In this topology, each router receives an IGMP membership report from its local receiver host, and then sends a PIM SG join message towards the source on VLAN 20. VLAN flooding propagates the PIM SG join message through to Router A, the PIM DR for the source VLAN 10. Each router from Router D to Router A records a PIM join on the port on which the join message was received, and then sends out its own join message toward the source. Data then flows from the source to the receiver, as long as a join exists on those ports.

Because all routers are in the same VLAN 20, they receive joins from one another due to flooding in the VLAN. For example, Router D receives join messages from Router C on its port ‘b’, and Router C receives join messages from Router B on its port toward Router B. On VSP 9000 routers, PIM join suppression is always enabled. In accordance with the PIM protocol rules, suppression causes Router D to stop sending a join towards the source because it receives a join for the same group and same RP on the port (port b) of the upstream neighbor (the router towards the source). Router D does not need to send a redundant join on the same VLAN. Router D stops sending a join, and the join that is recorded on port c of Router C eventually times out and is removed from the egress

list of the (s,g) multicast route entry on Router C. This removal causes Router C to stop forwarding multicast traffic to Router D, and to the receiver (R1).

The purpose of join suppression is to suppress joins on a shared VLAN, such that if all routers on the shared VLAN want to receive data from the same RP and group, then only one of them needs to send the join on the VLAN. One join is enough to pull the data from the source router to the shared VLAN for all routers to receive. The other routers can suppress sending their own joins when they see such a join on the port toward the upstream router. In this way, less protocol message congestion exists in the shared VLAN. In the following figure, Router D sends the initial join message, which is seen by Router B and Router C. Router B and Router C suppress their own join messages. Router A (the PIM DR for the source VLAN 10) sends the data to VLAN 20, which is received by Routers B, C, and D due to the shared (non-extended) VLAN topology, and traffic is forwarded to all receiver hosts.

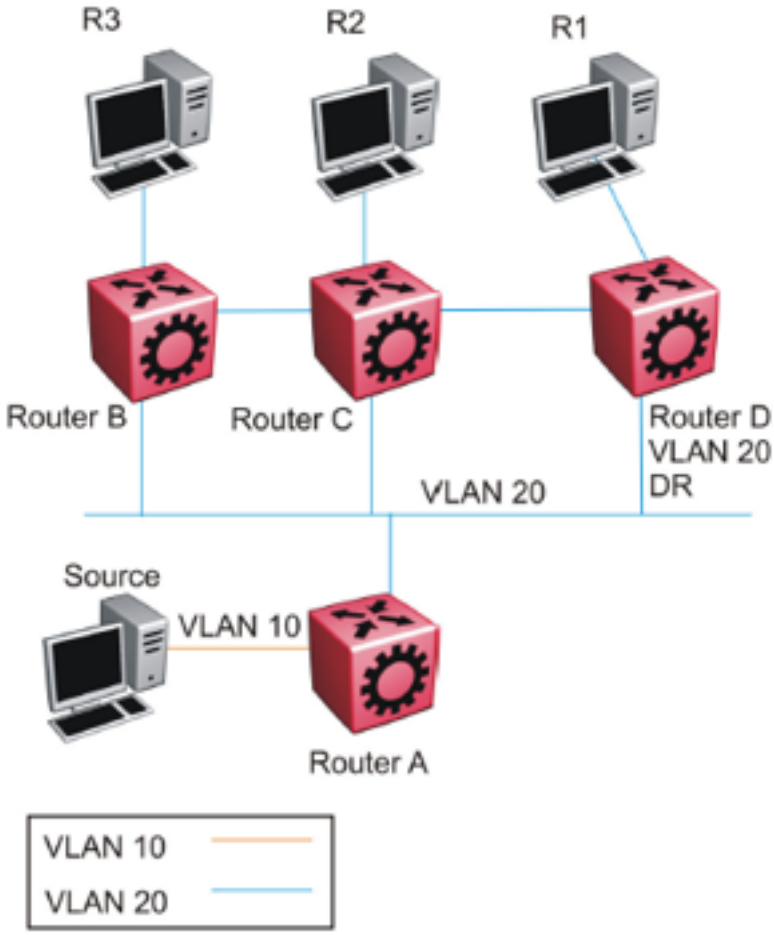

The extended VLAN topology looks exactly like the non-extended shared VLAN topology to the router, which cannot distinguish between the two.

In the current release, you cannot disable join suppression on a router. This enhancement will be added in a future release. Until this enhancement is included, you can perform the following actions:

1. Avoid this type of extended VLAN topology, and instead use Layer 3 routing between the routers. Do not extend VLAN 20 throughout, but rather, create a different VLAN between each router.
2. Configure the PIM DR for VLAN 20 to be the router closer to the source (Router B) so that any join received on the VLAN 20 DR (Router B) will be recorded as an IGMP local leaf on VLAN 20 as opposed to a PIM join, which does not time out until the receiver host stops sending IGMP membership reports.

PIM-SM Design and the BSR Hash Algorithm

To optimize the flow of traffic down the shared trees in a network that uses a BSR to dynamically advertise candidate RPs, consider the hash function. The BSR uses the hash function to assign multicast group addresses to each C-RP.

The BSR distributes the hash mask used to compute the RP assignment. For example, if two RPs are candidates for the range 239.0.0.0 through 239.0.0.127, and the hash mask is 255.255.255.252, that range of addresses is divided into groups of four consecutive addresses and assigned to one or the other C-RP.

The following figure illustrates a suboptimal design where Router A sends traffic to a group address assigned to RP D. Router B sends traffic assigned to RP C. RP C and RP D serve as backups for each other for those group addresses. To distribute traffic, it is desirable that traffic from Router A use RP C and that traffic from Router B use RP D.

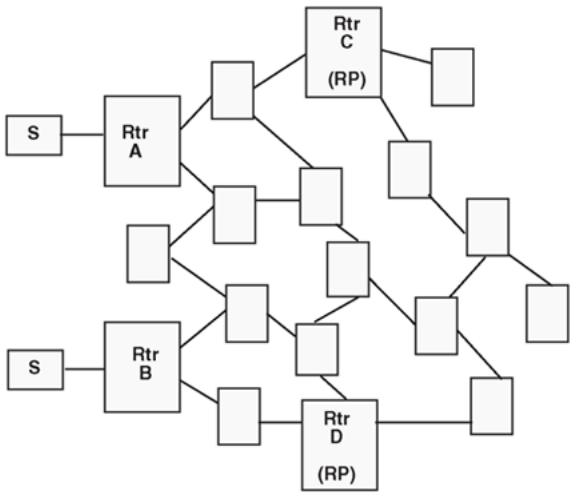

While still providing redundancy in the case of an RP failure, you can ensure that the optimal shared tree is used by using the following methods.

1. Use the hash algorithm to proactively plan the group-address-to-RP assignment.

   Use this information to select the multicast group address for each multicast sender on the network and to ensure optimal traffic flows. This method is helpful for modeling more

complex redundancy and failure scenarios, where each group address has three or more C-RPs.

2. Allow the hash algorithm to assign the blocks of addresses on the network, and then view the results using the command **`show ip pim active-rp`** .

   Use the command output to assign multicast group addresses to senders that are located near the indicated RP. The limitation to this approach is that while you can easily determine the current RP for a group address, the backup RP is not shown. If more than one backup for a group address exists, the secondary RP is not obvious. In this case, use the hash algorithm to reveal which of the remaining C-RPs take over for a particular group address in the event of primary RP failure.

The hash algorithm works as follows:

1. For each C-RP router with matching group address ranges, a hash value is calculated according to the formula:

   Hash value [G, M, C(i)] = {1 103 515 245 * [(1 103 515245 * (G&M) +12 345) XOR C(i)] + 12 345} mod 2^31

   The hash value is a function of the group address (G), the hash mask (M), and the IP address of the C-RP C(i). The expression (G&M) guarantees that blocks of group addresses hash to the same value for each C-RP, and that the size of the block is determined by the hash mask.

   For example, if the hash mask is 255.255.255.248, the group addresses 239.0.0.0 through 239.0.0.7 yield the same hash value for a given C-RP. Thus, the block of eight addresses are assigned to the same RP.

2. The C-RP with the highest resulting hash value is chosen as the RP for the group. In the event of a tie, the C-RP with the highest IP address is chosen.

   This algorithm runs independently on all PIM-SM routers so that every router has a consistent view of the group-to-RP mappings.

Candidate RP Considerations

The C-RP priority parameter determines an active RP for a group. The hash values for different RPs are only compared for RPs with the highest priority. Among the RPs with the highest priority value and the same hash value, the C-RP with the highest RP IP address is chosen as the active RP.

You cannot configure the C-RP priority. Each RP has a default C-RP priority value of 0, and the algorithm uses the RP if the group address maps to the grp-prefix that you configure for that RP. If a different router in the network has a C-RP priority value greater than 0, the switch uses this part of the algorithm in the RP election process.

Currently, you cannot configure the hash mask used in the hash algorithm. Unless you configure a different PIM BSR in the network with a nondefault hash mask value, the default hash mask of 255.255.255.252 is used. Static RP configurations do not use the BSR hash mask; they use the default hash mask.

For example:

RP1 = 128.10.0.54 and RP2 = 128.10.0.56. The group prefix for both RPs is 238.0.0.0/255.0.0.0. Hash mask = 255.255.255.252.

The hash function assigns the groups to RPs in the following manner:

The group range 238.1.1.40 to 238.1.1.51 (12 consecutive groups) maps to 128.10.0.56. The group range 238.1.1.52 to 238.1.1.55 (4 consecutive groups) maps to 128.10.0.54. The group range 238.1.1.56 to 238.1.1.63 (8 consecutive groups) maps to 128.10.0.56.

PIM-SM RP Selection Algorithm Inconsistency between Platforms

In topologies where VOSS platforms are interoperating with ERS or VSP9000 platforms, the selection of the RP from multiple candidate RPs may produce different results on VOSS than it does on ERS or VSP9000. VOSS platforms conform to PIM RFC 4601, while ERS and VSP9000 platforms conform to RFC 2362. RFC 4601 is not backward compatible with RFC 2362 regarding how it defines the selection algorithm for an RP, specifically when there are several candidate RPs for the same group, but with different prefix lengths. Both RFCs have the RP selection mechanism based on a specific hash function, common to all routers in PIM domain, however there are differences in determining the pool of candidate RPs to which the hash function will be applied. In RFC 4601, only the RP of the group range with the longest prefix match for the group range will be chosen to apply the hash function and thus participate in the actual election. In RFC 2362, longest prefix match is not part of the selection criteria, and therefore ERS/VSP9000 could potentially choose a different RP, since they will apply the hash function on a different pool of candidate RPs. This would cause inconsistencies in the PIM-SM network.

In order to work around this issue, please define RP group ranges with the same prefix length, such that the next RFC-defined match rule will be applied equally across all platforms in the network.

PIM-SM Receivers and VLANs

Some designs cause unnecessary traffic flow on links in a PIM-SM domain. In these cases, traffic is not duplicated to the receivers, but wastes bandwidth.

The following figure shows such a situation. Switch B is the DR between switches A and B. Switch C is the RP. A receiver R is on the VLAN (V1) that connects switches A and B. A source sends multicast data to the receiver.

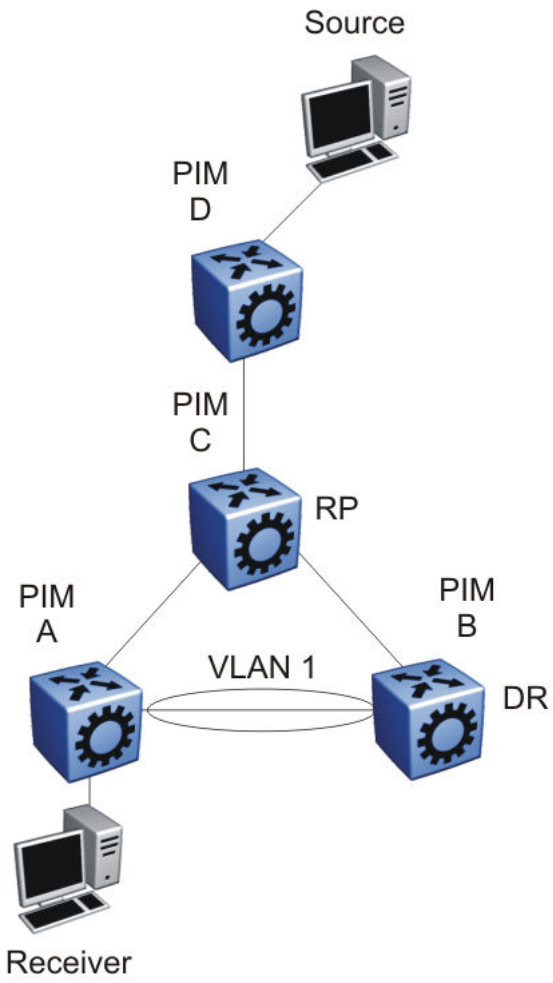

IGMP reports that the messages that the receiver sends are forwarded to the DR, and both A and B create (*,G) records. Switch A receives duplicate data through the path from C to A, and through the second path from C to B to A. Switch A discards the data on the second path (assuming the upstream source is A to C).

To avoid this waste of resources, do not place receivers on V1. This configuration guarantees that no traffic flows between B and A for receivers attached to A. In this case, the existence of the receivers is only learned through PIM join messages to the RP [for (*,G)] and of the source through SPT joins.

PIM Network with Non-PIM Interfaces

For proper multicast traffic flow in a PIM-SM domain, as a general rule, enable PIM-SM on all interfaces in the network (even if paths exist between all PIM interfaces). Enable PIM on all interfaces because PIM-SM relies on the unicast routing table to determine the path to the RP, BSR, and multicast sources. Ensure that all routers on these paths have PIM-SM enabled interfaces.

The following figure provides an example of this situation. If A is the RP, then initially the receiver receives data from the shared tree path (that is, through switch A).

If the shortest path from C to the source is through switch B, and the interface between C and B does not have PIM-SM enabled, then C cannot switch to the SPT. C discards data that comes through the shared path tree (that is, through A). The simple workaround is to enable PIM on VLAN1 between C and B.

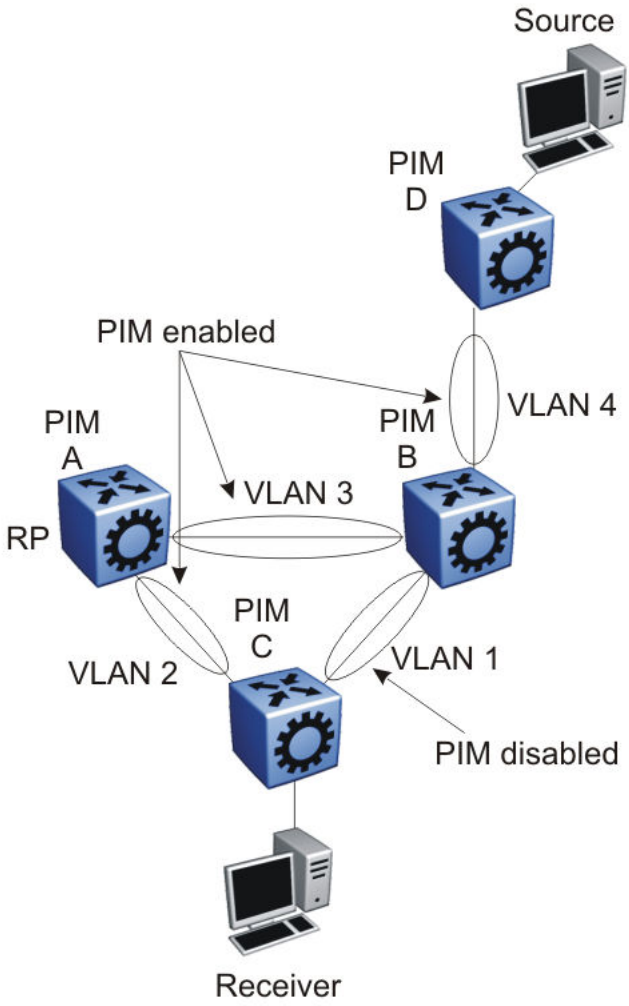

Source Filtering

The system can report interest in receiving packets from *only* a specific source address (INCLUDE), from all *but* specific source addresses (EXCLUDE), or sent to specific multicast addresses. IGMPv3 interacts with PIM-SM, PIM-SSM, and snooping to provide source filtering.

Protocol Independent Multicast-Source Specific Multicast guidelines

PIM-Source Specific Multicast (SSM) is a one-to-many model that uses a subset of the PIM-SM features. In this model, members of an SSM group can only receive multicast traffic from a specific source or sources, which is more efficient and puts less load on multicast routing devices.

IGMPv3 supports PIM-SSM by enabling a host to selectively request traffic from individual sources within a multicast group. The system can report interest in receiving packets from only specific source addresses (INCLUDE). IGMPv3 interacts with PIM-SM, PIM-SSM, and snooping to provide source filtering.

IGMPv2 SSM extensions

Virtual Services Platform 4000 processes messages according to the following rules:

- After IGMPv3 receives an IGMPv2 report in the SSM range, the system translates the report to an IGMPv3 report message.
- After an IGMPv2 router sends queries on an IGMPv3 interface, the switch downgrades this interface to IGMPv2 (backward compatibility).

  This can cause traffic interruption, but the switch recovers quickly.

PIM-SSM design considerations

Use the following information when you design an SSM network:

- If you configure SSM, it affects SSM groups only. The switch handles other groups in sparse mode (SM) if a valid RP exists on the network.
- You can configure PIM-SSM only on switches at the edge of the network. Core switches use PIM-SM if they do not have receivers for SSM groups.
- For networks where group addresses are already in use, you can change the SSM range to match the groups.
- One switch has a single SSM range.
- You can have different SSM ranges on different switches.

  Configure the core switches that relay multicast traffic so that they cover all of these groups in their SSM range, or use PIM-SM.
- One group in the SSM range can have multiple sources for a given SSM group.

Multicast for multimedia

The switch provides a flexible and scalable multicast implementation for multimedia applications. Several features are dedicated to multimedia applications and in particular to television distribution.

Join and leave performance

For TV applications, you can attach several TVs directly to the switch, or through an IGMP-capable Ethernet switch. Base this implementation on IGMP; the set-top boxes use IGMP reports to join a TV channel and IGMP leaves to exit the channel. After a viewer changes channels, the switch

issues an IGMPv2 leave for the old channel (multicast group), and sends a membership report for the new channel. If viewers change channels continuously, the number of joins and leaves can become large, particularly if many viewers attach to the switch.

The switch supports more than a thousand joins and leaves per second, which is well adapted to TV applications.

Important:

For IGMPv3, ensure a join rate of 1000 per second or less. This ensures the timely processing of join requests.

If you use the IGMP proxy functionality at the receiver edge, you reduce the number of IGMP reports received by switch. This provides better overall performance and scalability.

Fast Leave

IGMP Fast Leave supports two modes of operation: single-user mode and multiple-user mode.

In single-user mode, if more than one member of a group is on the port and one of the group members leaves the group, everyone stops receiving traffic for this group. Single-user mode does not send a group-specific query before the effective leave takes place.

Multiple-user mode allows several users on the same port or VLAN. If one user leaves the group and other receivers exist for the same stream, the stream continues. The switch tracks the number of receivers that join a given group. For multiple-user mode to operate properly, do not suppress reports. This ensures that the switch properly tracks the correct number of receivers on an interface.

The Fast Leave feature is particularly useful in IGMP-based TV distribution where only one receiver of a TV channel connects to a port. If a viewer changes channels quickly, you create considerable bandwidth savings if you use Fast Leave.

You can implement Fast Leave on a VLAN and port combination; a port that belongs to two different VLANs can have Fast Leave enabled on one VLAN (but not on the other). Thus, with the Fast Leave feature enabled, you can connect several devices on different VLANs to the same port. This strategy does not affect traffic after one device leaves a group to which another device subscribes. For example, you can use this feature when two TVs connect to a port through two set-top boxes, even if you use the single-user mode.

To use Fast Leave, you must first enable explicit host tracking. IGMP uses explicit host tracking to track all source and group members. Explicit host tracking is disabled by default. For configuration information, see Configuring fast leave mode on page 208.

Last member query interval tuning

If an IGMPv2 host leaves a group, it notifies the router by using a leave message. Because of the IGMPv2 report suppression mechanism, the router cannot access information of other hosts that require the stream. Thus, the router broadcasts a group-specific query message with a maximum response time equal to the last member query interval (LMQI).

Because this timer affects the latency between the time that the last member leaves and the time the stream actually stops, you must properly tune this parameter. This timer can especially affect TV delivery or other large-scale, high-bandwidth multimedia applications. For instance, if you assign a value that is too low, this can lead to a storm of membership reports if a large number of hosts are subscribed. Similarly, assigning a value that is too high can cause unwanted high-bandwidth stream propagation across the network if users change channels rapidly. Leave latency also depends on the robustness value, so a value of 2 equates to a leave latency of twice the LMQI.

Determine the proper LMQI value for your particular network through testing. If a very large number of users connect to a port, assigning a value of 3 can lead to a storm of report messages after a group-specific query is sent. Conversely, if streams frequently start and stop in short intervals, as in a TV delivery network, assigning a value of 10 can lead to frequent congestion in the core network.

Another performance-affecting factor that you need to be aware of is the error rate of the physical medium. For links that have high packet loss, you can find it necessary to adjust the robustness variable to a higher value to compensate for the possible loss of IGMP queries and reports.

In such cases, leave latency is adversely affected as numerous group-specific queries are unanswered before the stream is pruned. The number of unanswered queries is equal to the robustness variable (default 2). The assignment of a lower LMQI can counterbalance this effect. However, if you configure the LMQI too low, it can actually exacerbate the problem by inducing storms of reports on the network. LMQI values of 3 and 10, with a robustness value of 2, translate to leave latencies of 6/10 of a second and 2 seconds, respectively.

When you choose an LMQI, consider all of these factors to determine the best configuration for the given application and network. Test that value to ensure that it provides the best performance.

Important:

In networks that have only one user connected to each port, use the Fast Leave feature instead of LMQI, because no wait is required before the stream stops. Similarly, the robustness variable does not affect the Fast Leave feature, which is an additional benefit for links with high loss.

Layer 3 switch clustering and multicast SMLT

Switch clustering is the logical aggregation of two nodes to form one logical entity known as the switch cluster. The two peer nodes in a switch cluster connect using a virtual interswitch trunk (vIST). The vIST exchanges forwarding and routing information between the two peer nodes in the cluster. This section provides guidelines for switch clusters that use multicast and Split Multilink Trunking (SMLT).

General guidelines

The following list identifies general guidelines to follow if you use multicast and switch clustering:

- Enable Protocol Independent Multicast - Sparse Mode (PIM-SM) on the vIST VLAN for fast recovery of multicast. A unicast routing protocol is not required.
- Enable Internet Group Management Protocol (IGMP) snooping and proxy on the edge switches.

The following figure shows multicast behavior in an SMLT environment. The configuration in the following figure provides fast failover if the switch or rendezvous point (RP) fails.

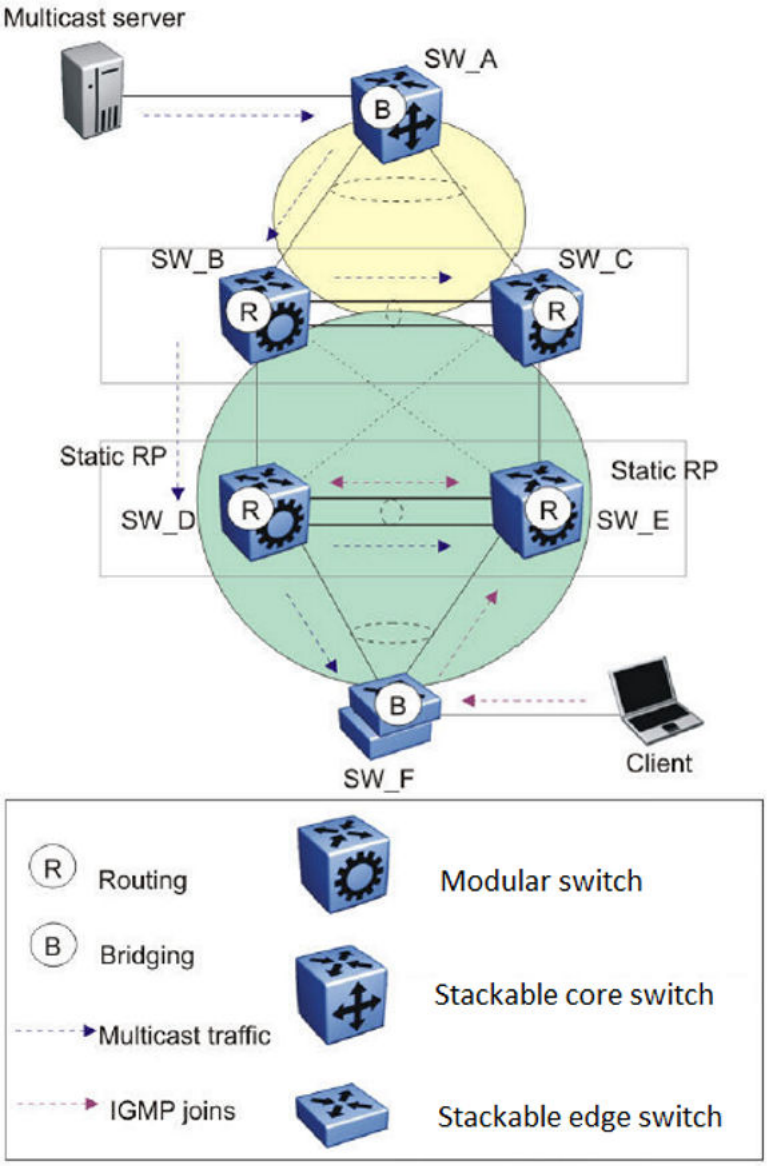

Figure 11: Multicast behavior in SMLT environment

In Multicast behavior in SMLT environment the following actions occur:

1. The multicast server sends multicast data towards the source designated router (DR).
2. The source DR sends register messages with encapsulated multicast data towards the RP.
3. After the client sends IGMP membership reports towards the multicast router, the router creates a (*,G) entry.
4. The RP sends join messages towards the source DR on the reverse path.
5. After the source DR receives the join messages, it sends native multicast traffic.
6. After SW_B or SW_D receives multicast traffic from upstream, it forwards the traffic on the vIST as well as on the SMLT link. Other aggregation switches drop multicast traffic received

over the vIST at egress. This action provides fast failover for multicast traffic. Both SW_D and SW_E (Aggregation switches) have similar (S,G) records.

7. In case of SW_D or RP failure, SW_B changes only the next-hop interface towards SW_E. Because the circuitless IP (CLIP) RP address is the same, SW_B does not flush (S,G) entries and achieves fast failover.

Multicast triangle topology

A triangle design is an SMLT configuration that connects edge switches or SMLT clients to two aggregation switches. Connect the aggregation switches together with a vIST that carries all the SMLT trunks configured on the switches.

The switch supports the following triangle configurations:

- a configuration with Layer 3 PIM-SM routing on both the edge and aggregation switches
- a configuration with Layer 2 snooping on the client switches and Layer 3 routing with PIM-SM on the aggregation switches

To avoid using an external query device to provide correct handling and routing of multicast traffic to the rest of the network, use the triangle design with IGMP Snoop at the client switches. Use multicast routing at the aggregation switches as shown in the following figure.

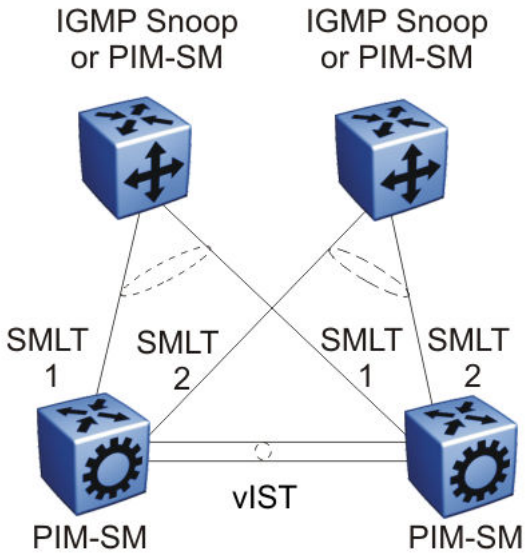

Figure 12: Multicast routing using PIM-SM

Client switches run IGMP Snoop or PIM-SM, and the aggregation switches run PIM-SM. This design is simple and, for the rest of the network, PIM-SM performs IP multicast routing. The aggregation switches are the query devices for IGMP, so an external query device is not required to activate IGMP membership. These switches also act as redundant switches for IP multicast.

Multicast data flows through the vIST link when receivers are learned on the client switch and senders are located on the aggregation switches, or when sourced data comes through the aggregation switches. This data is destined for potential receivers attached to the other side of the

vIST. The data does not reach the client switches through the two aggregation switches because only the originating switch forwards the data to the client switch receivers.

Note:

Always place multicast receivers and senders on the core switches on VLANs different from those that span the vIST.

The following figure shows a switch clustering configuration with a single switch cluster core and dual-connected edge devices. This topology represents different VLANs spanning from each edge device and those VLANs routed at the switch cluster core. You can configure multiple VLANs on the edge devices, 802.1Q tagged to the switch cluster core.

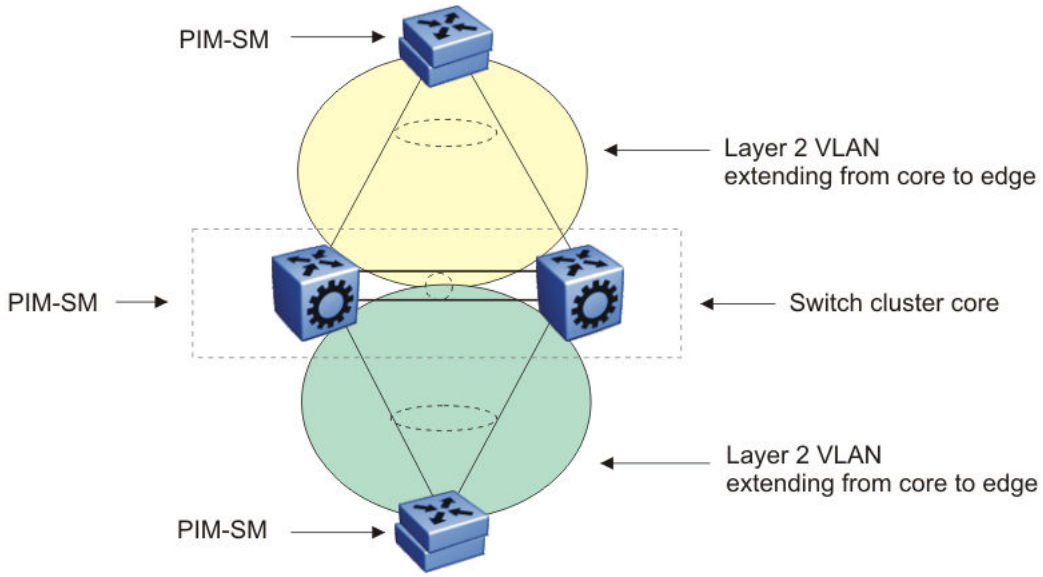

Figure 13: Multicast SMLT triangle

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either Virtual Router Redundancy Protocol (VRRP) BackupMaster or Routed SMLT (RSMLT) Layer 2 Edge on the switch cluster core.

Square and full-mesh topology multicast guidelines

A square design connects a pair of aggregation switches to another pair of aggregation switches. A square design becomes a full-mesh design if the aggregation switches are connected in a full-mesh. The switch supports Layer 3 IP multicast (PIM-SM only) over a full-mesh SMLT or RSMLT configuration.

In a square design, configure all switches with PIM-SM. Place the bootstrap router (BSR) and RP in one of the four core switches; and place the RP closest to the source. If using PIM-SM over a square or full-mesh configuration, enable the **`multicast smlt-square`** flag.

The following three figures show switch clustering configurations with two-switch cluster cores and dual-connected edge devices.

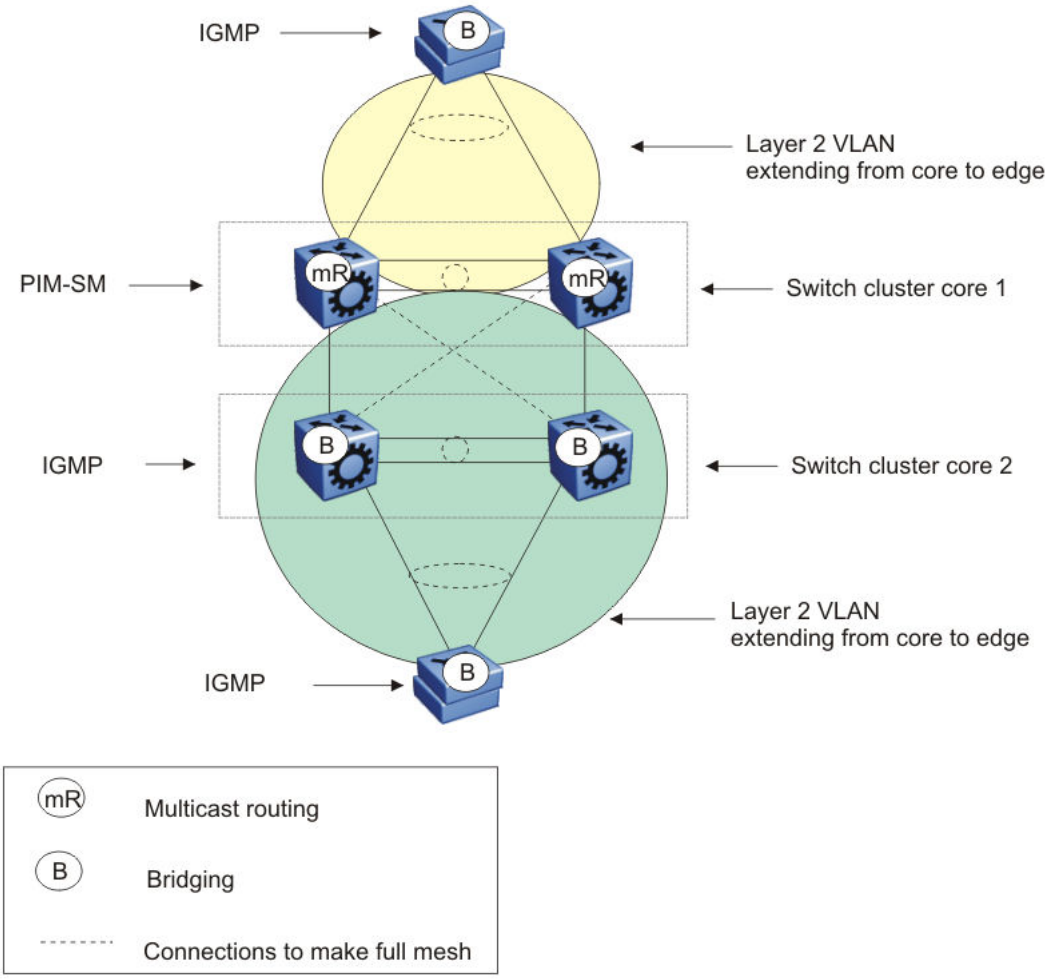

Figure 14: Multicast SMLT square 1

In the preceding figure, only one of the switch cluster cores performs Layer 3 multicast routing while the other is strictly Layer 2. Configure multiple VLANs on the edge devices, 802.1Q tagged to the switch cluster cores.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster core.

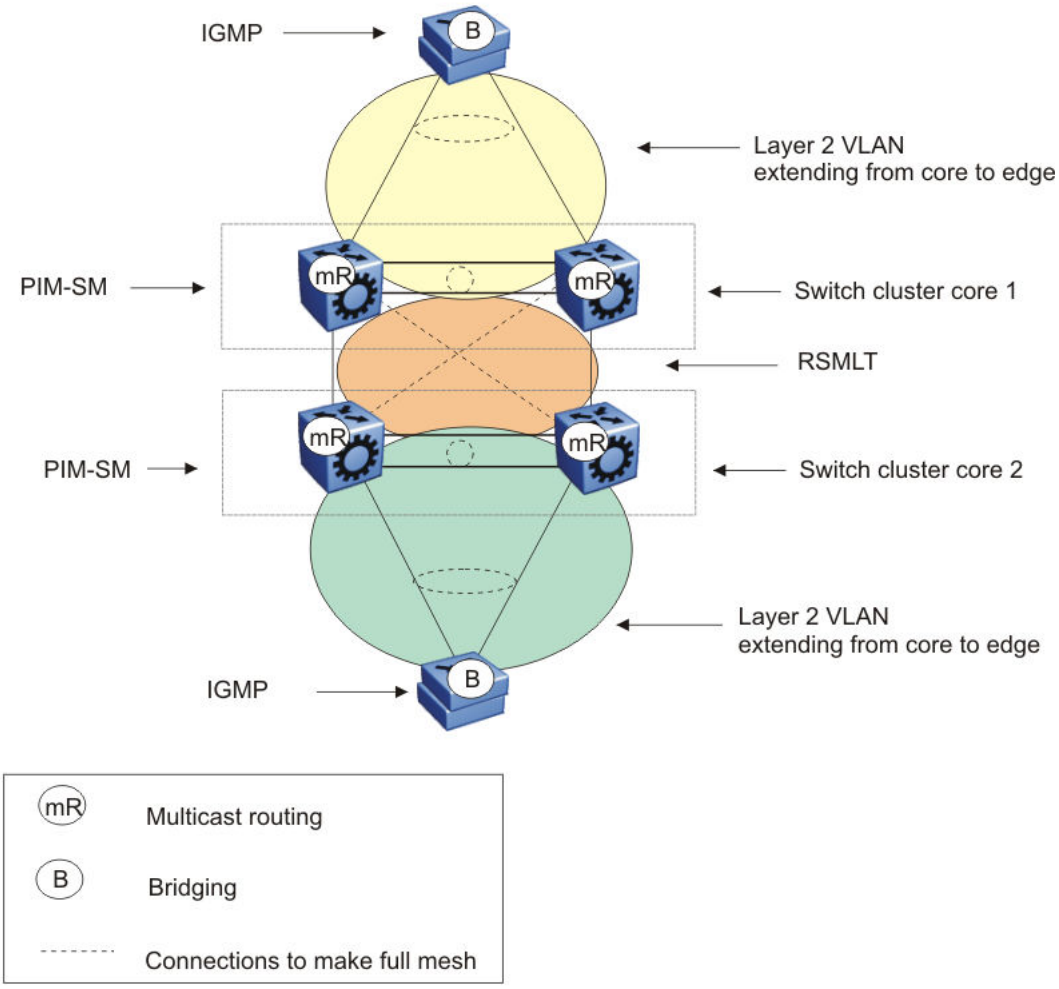

Figure 15: Multicast SMLT square 2

In the preceding figure, both of the switch cluster cores performs Layer 3 multicast routing, while the edge devices are Layer 2 IGMP.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster cores. Do not enable VRRP on the RSMLT VLAN between switch cluster cores.

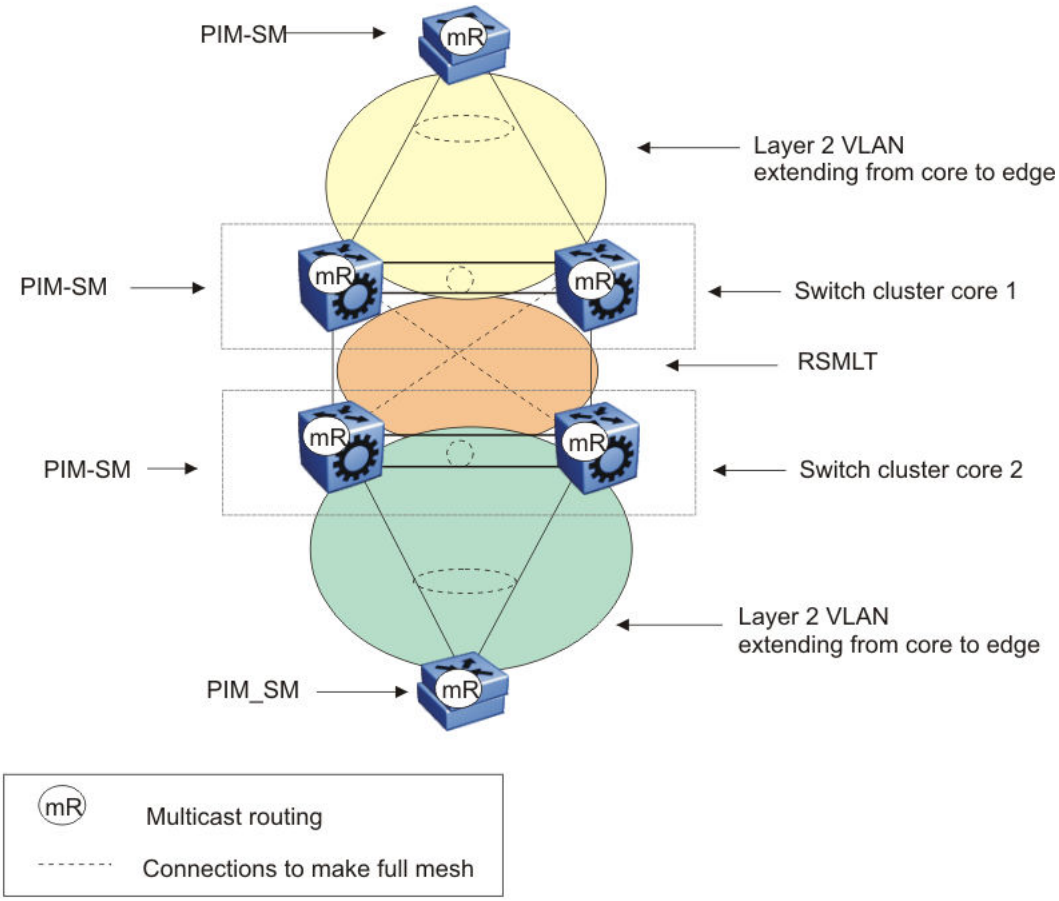

Figure 16: Multicast SMLT square 3

In the preceding figure, both of the switch cluster cores and the edge devices perform Layer 3 multicast routing.

Use an edge device that supports a form of link aggregation. Disable spanning tree on the link aggregation group on the edge devices. Enable either the VRRP BackupMaster or RSMLT Layer 2 Edge on the switch cluster cores. Do not enable VRRP on the RSMLT VLAN between switch cluster cores.

SMLT and multicast traffic issues

If PIM-SM or other multicast protocols are used in an SMLT environment, enable the protocol on the vIST. Routing protocols in general are not run over an vIST but multicast routing protocols are an exception. When using PIM-SM and a unicast routing protocol, ensure the unicast route to the BSR and RP has PIM-SM active and enabled. If multiple OSPF paths exist and PIM-SM is not active on each pair, the BSR is learned on a path that does not have PIM-SM active. The following figure demonstrates this issue.

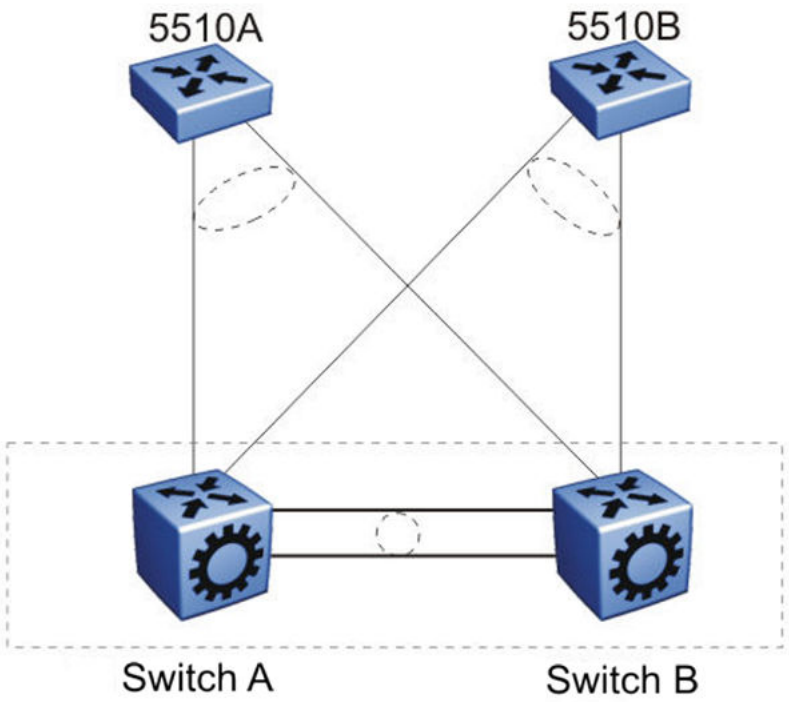

Figure 17: Unicast route example

The network configuration in the preceding figure is as follows:

- 5510A is on VLAN 101.
- 5510B is on VLAN 102.
- Switch B is the BSR.
- Switch A and Switch B have OSPF enabled.
- PIM is enabled and active on VLAN 101.
- PIM is either disabled or passive on VLAN 102.

In this example, the unicast route table on Switch A learns the BSR on Switch B through VLAN 102 using OSPF. The BSR is either not learned or does not provide the RP to Switch A.

Protocol Independent Multicast over IPv6

Several multicast protocols are used to enable IP multicast.

Hosts use the Internet Group Management Protocol (IGMP) for IPv4 and Multicast Listener Discovery (MLD v1/v2) for IPv6 to report multicast group memberships of directly attached multicast

listeners to neighboring multicast routers. MLD is the direct IPv6 replacement for the IGMP protocol used in IPv4.

Routers use Protocol Independent Multicast-Sparse Mode (PIM-SM) and PIM source Specific Mode (SSM) to exchange multicast routing information. The PIM-SM protocol is the multicast routing protocol that uses the underlying unicast routing information base to build unidirectional shared trees to group members rooted at the RP per group, and creates shortest-path trees (SPT) per source. The router forwards multicast packets along these trees. PIM-SSM does not require RP and only supports SPT.

PIM over IPv6 uses the IPv6 unicast routing table for reverse path information about source and RP.

Note:

IPv4 and IPv6 multicast streams cannot interact. To configure an end-to-end PIM IPv6 network, all nodes from sender to receiver must support PIM IPv6.

PIM-SM over IPv6 features

The following are features of PIM-SM over IPv6:

- Compliant with RFC 4601
- Multicast networks built by PIM IPv4 and PIM IPv6 do not overlap
- IPv4 receiver hosts cannot receive data from IPv6 source hosts and vice versa
- IPv4 and IPv6 multicast protocols can be enabled at the same time on the same VLAN
- PIM IPv4 and PIM IPv6 can be configured on the same VLAN
- PIM IPv4 and PIM IPv6 must be configured separately
- Supports sparse and ssm modes

Operational note for PIM-SM over IPv6

The following are operational considerations when deploying PIM-SM over IPv6:

- You can only configure PIM-SM if you configure the spbm_config_mode boot flag to false
- The following HELLO messages options are not supported:
  - GENid
  - DR priority
  - LAN-PRUNE delay
  - T-bit
- IPv6 multicast is not supported over SPBM

- IPv6 multicast routing is not virtualized, it is supported only on GRT
- IPv6 multicast configuration on SMLT VLAN is not supported. vIST peers cannot form PIM-SM over IPv6 neighbor adjacencies. Senders and receivers on the vIST peers (SMLT and non-SMLT) cannot communicate. MLT and LACP is supported.
- The switch does not support the following features:
  - Static entries
  - Bootstrap message (BSR)
  - Anycast RP
  - Virtual PIM neighbors
  - Fast join prune
  - Software forwarding
  - Passive PIM interfaces
  - IP mroute stream limit
  - Bidirectional PIM
  - Multicast Border Router (PMBR)
  - VRF support for PIM (GRT only)
  - IGMP and PIM mtrace capability
  - High Availability (HA)

IPv6 interface multiple addresses

IPv6 interfaces can have multiple addresses associated with them. A router running PIM for IPv6 has a network unique domain-wide reachable IPv6 VLAN address used for multiple hop messages. A link local address is associated with the VLAN. The link local address is a non-routable unicast IPv6 address used as source address (primary interface address) for transmitting different types of PIM messages.

IP multicast configuration and DvR

Configuration of IPv4 multicast is supported only on the Controller nodes of a DvR domain. You cannot configure IP multicast on the DvR Leaf nodes. The following sections detail IP multicast configuration support on DvR enabled nodes (Controllers or Leaf nodes).

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Multicast configuration that is pushed from DvR Controllers to DvR Leaf nodes

When you perform the following multicast configuration on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- IP multicast over Fabric Connect:

  For more information on IP multicast over Fabric Connect on a DvR enabled interface, see Configuring Fabric Multicast Services for VOSS.
- IGMP Layer 2 Querier parameters such as the IGMP Layer 2 Querier version, query interval, query maximum response time, robustness value, last member query interval and compatibility mode.
- Enabling and clearing of multicast route statistics

Multicast configuration that is not supported on DvR enabled Layer 2 VSNs

- IGMP Snooping on DvR enabled Layer 2 VSNs
- SPB-PIM Gateway

  For more information on SPB-PIM Gateway, see Configuring Fabric Multicast Services for VOSS.

Chapter 4: IP multicast basic configuration using CLI

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts subscribe to multicast services using a host membership protocol. The Internet Group Management Protocol (IGMP) is an example of an IPv4 host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast–Sparse Mode (PIM–SM) and Protocol Independent Multicast–Source Specific Multicast (PIM–SSM).

Configuring IP multicast in SMLT topologies

This procedure shows how to configure PIM and IGMP Snooping in an SMLT environment. The configuration steps show how to enable multicast, and then configure the usual PIM and IGMP Snooping related VLANs and global attributes. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

Before you begin

SPBM must not be enabled on the vIST peers or any router that participates in the PIM network.

About this task

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 4094 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.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable the boot flag:

   ```
   no boot config flags spbm-config-mode
   ```

   The system responds with these messages:

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

   Warning: Please carefully save your configuration file before
   rebooting the switch. Saving configuration file when spbm-config-
   mode is changed to disable, removes SPBM configurations from the
   configuration file.
   ```

3. Save the configuration and, then reboot the switch.

   Important:

   Any change to the **spbm-config-mode** boot flag requires a reboot for the change to take effect.

4. Create the vIST VLAN:

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

   interface vlan <1-4059>

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

5. Configure the vIST peer address and VLAN:

   ```
   virtual-ist peer-ip <A.B.C.D> vlan <1-4059>
   ```

6. Configure the SMLT MLT:

   ```
   mlt <1-512> enable

   mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]]
   [,...]}

   interface mlt <1-512>

   smlt
   ```

7. Configure the vIST MLT:

   ```
   mlt <1-512> enable

   mlt <1-512> member {slot/port[/sub-port] [-slot/port[/sub-port]]
   [,...]}

   mlt <1-512> encapsulation dot1q

   interface mlt <1-512>

   virtual-ist enable
   ```

Note:

The **virtual-ist enable** command enables Simplified vIST and is only available when the **spbm-config-mode** boot flag is disabled.

8. Create a customer VLAN and assign the SMLT MLT ID:

   ```
   vlan create <2-4059>
   vlan mlt <1-4059> <1-512>
   interface vlan <1-4059>
   ip address <A.B.C.D/X>
   ```

9. Configure PIM or IGMP Snooping on the SMLT VLAN:

   ```
   interface vlan <1-4059>
   ip pim enable
   ```
   or
   ```
   ip igmp snooping
   ```

10. Configure PIM on the vIST VLAN:

    ```
    interface vlan <1-4059>
    ip pim enable
    ```

11. Enable PIM globally:

    ```
    ip pim enable
    ```

Note:

You can also configure other global PIM attributes such as **ip pim join-prune-interval**.

Example

```
enable
configure terminal
no boot config flags spbm-config-mode
```

Save the configuration and reboot the switch.

```
virtual-ist peer-ip 198.51.100.0 vlan 50

mlt 3 enable
mlt 3 member 1/35,1/36
interface mlt 3
smlt
exit
mlt 5 enable
mlt 5 member 2/15,2/17
mlt 5 encapsulation dot1q
interface mlt 5
virtual-ist enable
exit
vlan create 50 type port-mstprstp 0
interface vlan 50
ip address 198.51.100.0 255.255.255.0 1
exit
vlan create 100
vlan mlt 100 3
```

```
interface vlan 100
ip address 192.0.2.0 255.255.255.0 2
exit
interface vlan 100
ip pim enable (or ip igmp snooping)
exit
interface vlan 50
ip pim enable
exit
ip pim enable
```

Configuring PIM-SM globally

Configure PIM-SM to enable or disable PIM-SM globally on the switch and change default global parameters.

About this task

PIM-SM is the default mode so you do not need to configure the PIM mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable PIM-SM:

   ```
   ip pim enable
   ```

3. Configure the time between bootstrap messages:

   ```
   ip pim bootstrap-period <5–32757>
   ```

4. Configure the timeout to discard data:

   ```
   ip pim disc-data-timeout <5–65535>
   ```

5. Enable the fast join prune interval:

   ```
   ip pim fast-joinprune
   ```

6. Configure the forward cache timeout:

   ```
   ip pim fwd-cache-timeout <10–86400>
   ```

7. Configure the interval for join and prune messages:

   ```
   ip pim join-prune-interval <1–18724>
   ```

8. Specify how long to suppress register messages:

   ```
   ip pim register-suppression-timeout <6–65535>
   ```

9. Specify how often the candidate-rendezvous point (C-RP) sends advertisements:

```
ip pim rp-c-adv-timeout <5-26214>
```

10. Configure the polling interval for the routing table manager (RTM):

```
ip pim unicast-route-change-timeout <2-65535>
```

11. Verify the configuration changes:

```
show ip pim
```

Example

Verify the configuration changes:

```
Switch:1(config)#show ip pim
```

```
================================================================
                    Pim General Group - GlobalRouter
================================================================
PimStat                : enabled
Mode                   : sparse
StaticRP               : disabled
FastJoinPrune          : disabled
BootstrapPeriod        : 60
CRPAdvTimeout          : 60
DiscDataTimeout        : 60
FwdCacheTimeout        : 210
RegSupprTimeout        : 60
UniRouteChangeTimeout  : 5
JoinPruneInt           : 60
```

Enabling or disabling IPv6 PIM-SM globally

About this task

Use this procedure to enable or disable IPv6 PIM-SM globally. By default, IPv6 PIM-SM is disabled.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable IPv6 PIM-SM:

```
ipv6 pim enable
```

3. Disable IPv6 PIM-SM:

```
no ipv6 pim enable
```

4. Set IPv6 PIM-SM status to default:

```
default ipv6 pim enable
```

Configuring global IPv6 PIM-SM properties

About this task

Use this procedure to configure the global IPv6 PIM-SM parameters on the switch.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the timeout to discard data:

   ```
   ipv6 pim disc-data-timeout <5-65535>
   ```

3. Configure the forward cache timeout:

   ```
   ipv6 pim fwd-cache-timeout <10-86400>
   ```

4. Configure the interval for join and prune messages:

   ```
   ipv6 pim join-prune-interval <1-18724>
   ```

5. Specify how long to suppress register messages:

   ```
   ipv6 pim register-suppression-timeout <10-65535>
   ```

6. Configure the polling interval for the routing table manager (RTM):

   ```
   ipv6 pim unicast-route-change-timeout <2-65535>
   ```

7. Configure the PIM mode:

   ```
   ipv6 pim mode <sparse> <ssm>
   ```

8. Verify the configuration changes:

   ```
   show ipv6 pim
   ```

Example

Verify the configuration changes:

```
Switch:1(config)#show ipv6 pim
================================================================================
                        Pim General Group - GlobalRouter

================================================================================
PimStat               : disabled
Mode                  : sparse
StaticRP              : disabled
FwdCacheTimeout       : 210
DiscDataTimeout       : 60
RegSupprTimeout       : 60
UniRouteChangeTimeout : 5
JoinPruneInt          : 60
```

Variable definitions

The following table describes the variables for the `ipv6 pim` command.

Variable	Description
disc-data-timeout <5-65535>	Specifies the duration in seconds to discard data until the switch receives the join message from the rendezvous point (RP). An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The default value is 60.
enable	Enables PIM globally on the switch. The default is disabled.
fwd-cache-timeout <10-86400>	Specifies the forward cache timeout value. The default value is 120.
join-prune-interval <1-18724>	Specifies the duration in seconds before the PIM router sends out the next join or prune message to its upstream neighbors. The default value is 60.
mode <sparse> <ssm>	Configures PIM mode on the switch. The default value is sparse.
register-suppression-timeout <10-65535>	Specifies the duration in seconds the designated router (DR) suppresses sending registers to the RP. The timer starts after the DR receives a register-stop message from the RP. The default value is 60.
static-rp	Add new static-rp entries and enable static-rp.
unicast-route-change-timeout <2-65535>	Specifies the duration in seconds the switch polls the RTM for unicast routing information updates for PIM. The default value is 5.

Configuring PIM on a VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you configure PIM on a VLAN.
- The interface uses a valid IP address.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Create a PIM interface on a VLAN:

   ```
   ip pim enable
   ```

   This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

   ```
   ip pim join-prune-interval <1-18724>
   ```

4. Configure the time between hello messages:

   ```
   ip pim hello-interval <0-18724>
   ```

5. Verify the configuration:

   ```
   show ip pim interface vlan [<1-4059>]
   ```

Example

Configure the interval for join and prune messages, the time between hello messages, and then verify the configuration.

```
Switch:1(config-if)#ip pim join-prune-interval 60
Switch:1(config-if)#ip pim hello-interval 30
Switch:1(config-if)#show ip pim interface vlan 10
================================================================================
                                 Vlan Ip Pim
================================================================================

VLAN-ID   PIM-ENABLE MODE    HELLOINT  JPINT   CBSRPREF      INTF TYPE
--------------------------------------------------------------------------------
10         enable    sparse  30        60      -1   (disabled) active
```

Configuring PIM on a port

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you configure it on an interface.
- The interface uses a valid IP address.

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. Create a PIM interface on a port:

```
ip pim enable
```

This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

```
ip pim join-prune-interval <1-18724>
```

4. Configure the time between hello messages:

```
ip pim hello-interval <0-18724>
```

Example

Configure the interval for join and prune messages and the time between hello messages:

```
Switch(config-if)#ip pim join-prune-interval 60
Switch(config-if)#ip pim hello-interval 30
```

Configuring IPv6 PIM on a port or VLAN

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- Enable IPv6 interface.

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. Create a PIM interface on a port or VLAN:

```
ipv6 pim enable
```

This command creates an active interface, by default.

3. Configure the interval for join and prune messages:

```
ipv6 pim join-prune-interval <1-18724>
```

4. Configure the time between hello messages:

```
ipv6 pim hello-interval <0-18724>
```

Example

```
Switch:1(config-if)#ipv6 pim join-prune-interval 60
Switch:1(config-if)#ipv6 pim hello-interval 30
```

Variable definitions

The following table describes the variables for the `ipv6 pim` command.

Variable	Description
hello-interval <0-18724>	Specifies the duration in seconds before the PIM router sends out the next hello message to neighboring switches. The default value is 30 seconds.
join-prune-interval <1-18724>	Specifies the duration in seconds before the PIM router sends out the next join or prune message to its upstream neighbors. The default value is 60 seconds.

Configuring SSM globally

Configure SSM to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Before you begin

- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP and OSPF, see Configuring OSPF and RIP for VOSS.
- Enable PIM globally.

About this task

Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range.

For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure PIM-SSM:

   ```
   ip pim mode ssm
   ```

Configuring IPv6 SSM globally

Configure IPv6 SSM to optimize IPv6 PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Before you begin

- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM.

  For more information about RIPng and OSPFv3, see Configuring IPv6 Routing for VOSS.

- Enable IPv6 PIM globally.

About this task

Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model which requires only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On a SSM-enabled switch, SSM behavior is limited to the SSM group range.

For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure IPv6 PIM-SSM:

   ```
   ipv6 pim mode ssm
   ```

Configuring IGMP on a VLAN

Configure IGMP for each interface to change default multicasting operations.

Note:

When you configure the following IGMP parameters on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- **ip igmp version**
- **ip igmp query-interval**
- **ip igmp query-max-response**
- **ip igmp robust-value**
- **ip igmp last-member-query-interval**
- **ip igmp compatibility-mode**

IGMP snooping is not supported on DvR enabled Layer 2 VSNs.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Before you begin

- For PIM interfaces, you must enable PIM globally and on the VLAN. For snooping interfaces, do not enable PIM.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable IGMP v2-v3 compatibility mode:

   ```
   ip igmp compatibility-mode
   ```

3. Configure the system to downgrade the version of IGMP:

   ```
   ip igmp dynamic-downgrade-version
   ```

4. Configure message intervals and response times:

   ```
   ip igmp last-member-query-interval <0-255> [query-interval <1-65535>] [query-max-response <0-255>]
   ```

5. Configure expected packet loss and IGMP version:

   ```
   ip igmp robust-value <2-255> [version <1-3>]
   ```

6. Add multicast router ports:

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

7. Enable proxy-snoop:

   ```
   ip igmp proxy
   ```

8. Enable router alert:

   ```
   ip igmp router-alert
   ```

9. Enable snooping:

   ```
   ip igmp snooping
   ```

10. Enable SSM-snooping:

    ```
    ip igmp ssm-snoop
    ```

Example

Enter VLAN Interface Configuration Mode for VLAN 10:

```
Switch:1(config)# interface vlan 10
```

Configure the last member query interval to 15 tenths of a second (equal to 1.5 seconds).

```
Switch:1(config-if)# ip igmp last-member-query-interval 15
```

Configure the query interval to 100 seconds.

```
Switch:1(config-if)# ip igmp query-interval 100
```

Configure the query maximum response time to 15 tenths of a second (equal to 1.5 seconds).

```
Switch:1(config-if)# ip igmp query-max-response 15
```

Configure the robustness value to 4 seconds.

```
Switch:1(config-if)# ip igmp robust-value 4
```

Enable proxy snoop for the VLAN.

```
Switch:1(config-if)# ip igmp proxy
```

Enable snoop for the VLAN.

```
Switch:1(config-if)# ip igmp snooping
```

Enable support for SSM on the snoop interface.

```
Switch:1(config-if)# ip igmp ssm-snoop
```

Enable IGMPv3.

```
Switch:1(config-if)# ip igmp version 3
```

Variable definitions

Use the definitions in the following table to use the **ip igmp** command.

Variable	Value
access-list *WORD<1–64>* {A.B.C.D/X} <eny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Specifies the name of the access list from 1–64 characters. Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host. Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic

compatibility-mode	Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command:
	`default ip igmp compatibility-mode`
	, or use the no option to disable compatibility mode:
	`no ip igmp compatibility-mode`

dynamic-downgrade-version	Configures if the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:
	`default ip igmp dynamic-downgrade-version`
	or use the no option to disable downgrade:
	`no ip igmp dynamic-downgrade-version`

igmpv3-explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is disabled.
immediate-leave	Enables fast leave on a VLAN.
immediate-leave-members *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Configures IGMP fast leave members on a VLAN to specify fast-leave-capable ports. 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.
last-member-query-interval <0–255>	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1.

Table continues…

Variable	Value
	Decreasing the value reduces the time to detect the loss of the last member of a group. The default is 10 tenths of a second. You should configure this value between 3–10 (equal to 0.3 – 1.0 seconds).
mrdisc [maxadvertinterval <2–180>] [maxinitadvertinterval <2–180>] [maxinitadvertisements <2–15>] [minadvertinterval <3–180>] [neighdeadinterval <2–180>]	Configure the multicast router discovery options to enable the automatic discovery of multicast capable routers. The default parameter values are: • maxadvertinterval: 20 seconds • maxinitadvertinterval: 2 seconds • maxinitadvertisements: 3 • minadvertinterval: 15 seconds • neighdeadinterval: 60 seconds
mrouter *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Adds multicast router ports. 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.
proxy	Activates the proxy-snoop option globally for the VLAN.
query-interval <1–65535>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
query-max-response <0–255>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robust-value <2–255>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.
router-alert	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option.

Table continues…

Variable	Value
	Important: To maximize network performance, configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
snoop-querier	Enables the IGMP Layer 2 Querier feature on the VLAN. The default is disabled.
snoop-querier-addr {A.B.C.D}	Specifies the IGMP Layer 2 Querier source IP address.
snooping	Activates the snoop option for the VLAN.
ssm-snoop	Activates support for PIM-SSM on the snoop interface.
static-group {A.B.C.D} {A.B.C.D}*{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* [static|blocked]	Configures IGMP static members to add members to a snoop group. {A.B.C.D} {A.B.C.D} indicates the IP address range of the selected multicast group. *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* adds ports to a static group entry. [static|blocked] configures the route to static or blocked.
stream-limit stream-limit-max-streams <*0-65535*>	Configure multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. The default is 4.
stream-limit-group *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}* enable max-streams <*0-65535*>	Configure multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN. The default max-streams value is 4.
version <1–3>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configuring IGMP ports

Configure IGMP for each interface to change default multicasting operations.

Note:

When you configure the following IGMP parameters on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- **`ip igmp version`**

- **ip igmp query-interval**
- **ip igmp query-max-response**
- **ip igmp robust-value**
- **ip igmp last-member-query-interval**
- **ip igmp compatibility-mode**

For more information on DvR, see Configuring IPv4 Routing for VOSS.

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 IGMP v2-v3 compatibility mode:

   ```
   ip igmp compatibility-mode
   ```

3. Configure the system to downgrade the version of IGMP:

   ```
   ip igmp dynamic-downgrade-version
   ```

4. Configure message intervals and response times:

   ```
   ip igmp last-member-query-interval <0-255> [query-interval <1-65535>] [query-max-response <0-255>]
   ```

5. Configure expected packet loss and IGMP version:

   ```
   ip igmp robust-value <2-255> [version <1-3>]
   ```

6. Configure IGMP for a specific port:

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

7. Enable router alert:

   ```
   ip igmp router-alert
   ```

Example

Configure message intervals and response times:

```
Switch(config-if)#ip igmp last-member-query-interval 30 query-interval 60 query-max-response 90
```

Configure expected packet loss and IGMP version:

```
Switch(config-if)#ip igmp robust-value 2 version 3
```

Configure IGMP for a specific port:

```
Switch(config-if)#ip igmp port 1/4
```

Enable router alert:

```
Switch(config-if)#ip igmp router-alert
```

Variable definitions

Use the definitions in the following table to use the **`ip igmp`** command.

Variable	Value
access-list *WORD<1–64>* {A.B.C.D/X} <eny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>	Specifies the name of the access list from 1–64 characters. Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host. Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic

compatibility-mode	Activates v2-v3 compatibility mode. The default value is disabled, which means IGMPv3 is not compatible with IGMPv2. To use the default configuration, use the default option in the command:
	`default ip igmp compatibility-mode`
	, or use the no option to disable compatibility mode:
	`no ip igmp compatibility-mode`

dynamic-downgrade-version	Configures if the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMPv3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:
	`default ip igmp dynamic-downgrade-version`
	or use the no option to disable downgrade:
	`no ip igmp dynamic-downgrade-version`

igmpv3-explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is disabled.

Table continues…

Variable	Value
immediate-leave	Enables fast leave on a port.
last-member-query-interval <0–255>	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The default is 10 tenths of a second. You should configure this value between 3–10 (equal to 0.3 – 1.0 seconds).
port *{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}*	Configures IGMP for a specific 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.
query-interval <1–65535>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
query-max-response <0–255>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the query-interval.
robust-value <2–255>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.
router-alert	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. **Important:** To maximize network performance, configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable

Table continues…

Variable	Value
stream-limit stream-limit-max-streams *<0-65535>*	Configure multicast stream limitation on a port to limit the number of concurrent multicast streams on the port. The default is 4.
version <1–3>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configuring IGMP brouter ports

Configure IGMP for each interface to change default multicasting operations.

Note:

When you configure the following IGMP parameters on the DvR enabled interface of a DvR Controller, the configuration is automatically pushed to the Leaf nodes within the domain.

- **`ip igmp version`**
- **`ip igmp query-interval`**
- **`ip igmp query-max-response`**
- **`ip igmp robust-value`**
- **`ip igmp last-member-query-interval`**
- **`ip igmp compatibility-mode`**

For more information on DvR, see Configuring IPv4 Routing for VOSS.

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 IGMP v2-v3 compatibility mode:

   ```
   ip igmp compatibility-mode
   ```

3. Configure the system to downgrade the version of IGMP:

```
ip igmp dynamic-downgrade-version
```

4. Configure message intervals and response times:

   ```
   ip igmp last-member-query-interval <0–255> [query-interval <1–65535>] [query-max-response <0–255>]
   ```

5. Configure expected packet loss and IGMP version:

   ```
   ip igmp robust-value <2–255> [version <1–3>]
   ```

6. Configure IGMP for a specific port:

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

7. Enable router alert:

   ```
   ip igmp router-alert
   ```

Example

Configure message intervals and response times:

```
Switch:1(config-if)#ip igmp last-member-query-interval 30 query-interval 60 query-max-response 90
```

Configure expected packet loss and IGMP version:

```
Switch:1(config-if)#ip igmp robust-value 2 version 3
```

Configure IGMP for a specific port:

```
Switch:1(config-if)#ip igmp port 1/4
```

Enable router alert:

```
Switch:1config-if)#ip igmp router-alert
```

Variable definitions

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

Variable	Value
access-list *WORD<1–64>* {A.B.C.D/X} <eny-tx|deny-rx|deny-both|allow-only-tx| allow-only-rx|allow-only-both>	Specifies the name of the access list from 1 to 64 characters. Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask that is used to determine the host or hosts that are covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host. Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
compatibility-mode	Activates v2-v3 compatibility mode. The default value is Disabled, which means IGMPv3 is not compatible with IGMPv2.

Table continues…

Variable	Value

	To use the default configuration, use the default option in the command:
	`default ip igmp compatibility-mode`
	or use the no option to disable compatibility mode:
	`no ip igmp compatibility-mode`

dynamic-downgrade-version	Configures whether the system downgrades the version of IGMP to handle older query messages. If the system downgrades, the host with IGMP-v3 only capability does not work. If you do not configure the system to downgrade the version of IGMP, the system logs a warning. The system downgrades to the oldest version of IGMP on the network by default. To use the default configuration, use the default option in the command:
	`default ip igmp dynamic-downgrade-version`
	or use the no option to disable downgrade:
	`no ip igmp dynamic-downgrade-version`

igmpv3-explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is Disabled.
immediate-leave	Enables fast leave on a port.
last-member-query-interval <0–255>	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The default is 10 tenths of a second. It is recommended that you configure this value between 3–10 (equal to 0.3 – 1.0 seconds).
port {slot/port[-slot/port][,...]}	Configures IGMP for a specific port.
query-interval <1–65535>	Configures the frequency (in seconds) at which the VLAN transmits host query packets. The default value is 125 seconds.
query-max-response <0–255>	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value to a number that is lower than the query-interval.

Table continues…

Variable	Value
robust-value <2–255>	Configures the expected packet loss of a network. The default value is 2 seconds. Increase the value if you expect the network to experience packet loss.
router-alert	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. **Important:** To maximize network performance, configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
stream-limit stream-limit-max-streams *<0-65535>*	Configure multicast stream limitation on a port to limit the number of concurrent multicast streams on the port. The default is 4.
version <1–3>	Configures the version of IGMP that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default value is 2 (IGMPv2).

Configuring IGMP on a VRF

You configure IGMP on a VRF instance the same way you configure IGMP for the Global Router, except that you must use VRF Router Configuration mode.

Procedure

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

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

2. Enable SSM dynamic learning:

   ```
   ip igmp ssm dynamic-learning
   ```

3. Configure the range group:

   ```
   ip igmp ssm group-range {A.B.C.D/X}
   ```

   The following message appears:.

```
Warning: Changing the SSM range will cause all spb-multicast and
spb-pim-gw enabled interfaces to be internally bounced. Do you wish
to continue? (y/n) ? (y/n)?
```

Enter `y` to continue.

4. Enable the SSM map table for all static entries:

   ```
   ip igmp ssm-map all
   ```

5. Create a static entry for a specific group:

   ```
   ip igmp ssm-map {A.B.C.D} {A.B.C.D} enable
   ```

6. Enable the generation of IGMP traps:

   ```
   ip igmp generate-trap
   ```

7. Enable the generation of IGMP log messages:

   ```
   ip igmp generate-log
   ```

8. Configure the fast leave mode:

   ```
   ip igmp immediate-leave-mode {multiple-user|one-user}
   ```

Example

For the VRF Red context, configure a new IP multicast group address and create an SSM map table entry for the multicast group and the source at 192.32.99.151. Configure the administrative state to enable all the static SSM map table entries.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#router vrf red
Switch:1(router-vrf)#ip igmp ssm group-range 232.1.1.10/32

WARNING: Changing the SSM range will cause all spb-multicast and spb-pim-gw enabled
interfaces to be internally bounced. Do you wish to continue? (y/n) ? (y/n)? y
Switch:1(router-vrf)#ip igmp ssm-map 232.1.1.10 192.32.99.151
Switch:1(router-vrf)#ip igmp ssm-map all
```

Variable definitions

Use the definitions in the following table to use the **`ip igmp`** command on a VRF.

Variable	Value
generate-log	Enables the generation of IGMP log messages. The default is disabled.
generate-trap	Enables the generation of IGMP traps. The default is disabled.

Table continues…

Variable	Value
immediate-leave-mode {multiple-user|one-user}	• multiple-user: Removes (from the group) the IGMP member who sent the leave message. The default is multiple-user. • one-user: Removes all group members on a fast leave enabled interface port after receiving the first leave message from a member.
ssm dynamic-learning	Enables dynamic learning from IGMPv3 reports. The default is enabled.
ssm group-range {A.B.C.D/X}	Changes the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.
ssm-map <all | {A.B.C.D} {A.B.C.D} enable	Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that sends traffic to the group. Enables the administrative state for a specific entry (group). This variable does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.

Chapter 5: IP multicast basic configuration using EDM

To provide multicasting services, you need a host membership protocol and a multicast routing protocol. Hosts use a host membership protocol to subscribe to multicast services. The Internet Group Management Protocol (IGMP) is an example of an IPv4 host membership protocol.

A multicast routing protocol optimizes the routing of multicast information to avoid loops and restrict multicast traffic to networks that use host membership. Examples of multicast routing protocols include Protocol Independent Multicast–Sparse Mode (PIM–SM) and Protocol Independent Multicast–Source Specific Multicast (PIM–SSM).

Configuring multicast on the switch

This procedure shows how to configure PIM and IGMP Snooping in an SMLT environment. The configuration steps show how to enable multicast, and then configure the usual PIM and IGMP Snooping related VLANs and global attributes. It includes steps to configure the following:

- Setting the boot config flag
- Configuring the vIST peer
- Enabling Simplified vIST

Before you begin

SPBM must not be enabled on the vIST peers or any router participating in the PIM network.

Procedure

1. In the navigation pane, expand the **Configuration** > **Edit** > **Chassis** folders.
2. Click the **Boot Config** tab.
3. Clear the **EnableSpbmConfigMode** to disable the boot flag.

   The system responds with these messages:

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

   Warning: Please carefully save your configuration file before
   rebooting the switch. Saving configuration file when spbm-config-
   ```

```
mode is changed to disable, removes SPBM configurations from the
configuration file.
```

4. Click **Apply**.
5. Save the configuration, and then reboot the switch.

   Important:

   Any change to the **EnableSpbmConfigMode** boot flag requires a reboot for the change to take effect.

6. Configure the *SMLT* MLT:
   a. Expand the following folders: **Configuration** > **VLAN** > **MLT/LACP**.
   b. Click the **MultiLink/LACP Trunks** tab.
   c. Click **Insert**.
   d. In the **Id** box, type the ID number of the MLT.
   e. In the **PortMembers** box, click the (**...**) button.
   f. In the **Port Editor: PortMembers** dialog box, select the desired ports.
   g. Click **Ok**
   h. Click **Insert**.

      The switch adds the SMLT MLT to the MultiLink/LACP Trunks tab in the MLT_LACP box.

7. Configure the *vIST* MLT:
   a. Repeat steps 6a to 6g to configure the MLT.
   b. Click **MltVistEnable** to enable Simplified vIST.

      Note:

      The **MltVistEnable** field enables Simplified vIST and is only available when the **EnableSpbmConfigMode** boot flag is disabled.

   c. Click **Insert**.
8. Create the *vIST* VLAN:
   a. Expand the following folders: **Configuration** > **VLAN** > **VLANs**
   b. In the **Basic** tab, click **Insert**.
   c. In the **Id** box, enter an unused VLAN ID, or use the ID provided.
   d. In the **MstpInstance** box, click the down arrow, and then choose an MSTI instance from the list.
   e. In the **Type** box, select **byPort**.
   f. Click **OK**.

g. Click **Insert**.

h. Select the vIST VLAN from the list of VLANs, and then click **IP**.

i. Click **Insert**.

j. Configure the IP address for the vIST VLAN.

9. Repeat Step 8 to create an *SMLT* VLAN and assign the SMLT MLT ID to it. Do not use the vIST MLT ID.
10. Configure PIM or IGMP Snooping on the *SMLT* VLAN:
    a. To enable PIM, select the SMLT VLAN from the list of VLANs and click **IP** > **PIM**. Select **Enable** and click **Apply**.
    b. To enable IGMP Snooping, select the SMLT VLAN from the list of VLANs and click **IP** > **IGMP**. Select **SnoopEnable** and click **Apply**.
11. Configure PIM on the *SMLT* VLAN:

    To enable PIM, select the SMLT VLAN from the list of VLANs and click **IP** > **PIM**. Select **Enable** and click **Apply**.
12. Click **IP** > **PIM** > **Globals** to enable PIM globally.
13. Select the **Enable** check box, and then click **Apply**.

Selecting and launching a VRF context view

About this task

Use this procedure to switch to another VRF context view when you use the embedded EDM. GlobalRouter is the default view at log in. You can configure both Global Router (GRT) and Virtual Routing and Forwarding (VRF) instances when you launch a VRF context view. You can open only five tabs for each EDM session.

Important:

If you log out from the GRT view, the system generates a warning: all tabs close and your session terminates. If you close a VRF view tab, you close only that view.

Note:

The Set VRF Context view function is not available to users in a service provider deployment where only a tenant VRF view is assigned.If you use a tenant VRF view, it is recommended to use the applicable EDM plugin with COM to access EDM. COM provides VRF mapping and Role-Based Access Control.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VRF Context View**.
2. Click **Set VRF Context View**.

3. Click the **VRF** tab.
4. Select a context to view.
5. Click **Launch VRF Context view**.

   A new browser tab opens containing the selected VRF view

VRF field descriptions

Use the descriptions in the following table to use the **VRF** tab.

Name	Description
Id	Shows the unique VRF ID.
Name	Shows the name of the virtual router.
ContextName	Shows the SNMPv3 context name that denotes the VRF context and logically separates the MIB port management.

Enabling PIM-SM globally

Enable PIM-SM to offer multicasting services. After you enable PIM-SM globally and on a particular interface, the IGMP parameters take effect.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Globals** tab.
4. Click **sm** (sparse mode).
5. Select the **Enable** check box.
6. Click **Apply**.

Globals field descriptions

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

Name	Description
Mode	Configures the mode on the routing switch: sm (Sparse Mode) or ssm (Source Specific Multicast).

Table continues…

Name	Description
Enable	Enables or disables PIM.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The range is from 1–18724 and the default is 60 seconds.
RegisterSuppTimer	Specifies how long (in seconds) the designated router suppresses sending registers to the rendezvous point (RP). The timer starts after the designated router receives a register-stop message from the RP. The range is from 6–65535 and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager for unicast routing information updates for PIM. The range is from 2–65535 and the default is 5 seconds. **Important:** If you lower this value, it increases how often the switch polls the routing table manager. This value can affect the performance of the switch, especially if a high volume of traffic flows through the switch.
DiscardDataTimeOut	Specifies how long (in seconds) to discard data until the switch receives a join message from the RP. An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The range is from 5–65535 and the default is 60 seconds.
CRPADVTimeOut	Specifies how often (in seconds) a router configured as a candidate rendezvous point router (C-RP) sends advertisement messages. After this timer expires, the C-RP router sends an advertisement message to the elected bootstrap router (BSR). The range is from 5–26214 and the default is 60 seconds.
BootStrapPeriod	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. The range is from 5–32757 and the default is 60 seconds.
StaticRP	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism.
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value ages PIM mroutes in seconds. The range is from 10–86400 and the default value is 210. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.
FastJoinPrune	Enables or disables the PIM fast join prune feature.

Enabling IPv6 PIM-SM globally

Enable IPv6 PIM-SM to offer multicasting services. After you enable IPv6 PIM-SM globally and on a particular interface, the MLD parameters take effect.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Globals** tab.
4. Select the **Enable** check box.
5. Click **sm** (sparse mode).
6. Click **Apply**.

Globals field descriptions

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

Name	Description
Enable	Enables or disables PIM.
Mode	Configures the mode on the routing switch: sm (Sparse Mode) or ssm (Source Specific Multicast).
RegisterSuppTimer	Specifies how long (in seconds) the designated router suppresses sending registers to the rendezvous point (RP). The timer starts after the designated router receives a register-stop message from the RP. The range is from 10–65535 and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager for unicast routing information updates for PIM. The range is from 2–65535 and the default is 5 seconds. **Important:** If you lower this value, it increases how often the switch polls the routing table manager. This value can affect the performance of the switch, especially if a high volume of traffic flows through the switch.
DiscardDataTimeOut	Specifies how long (in seconds) to discard data until the switch receives a join message from the RP. An IP multicast discard record is created after a register packet is sent, until the timer expires or the switch receives a join message. The range is from 5–65535 and the default is 60 seconds.

Table continues…

Name	Description
StaticRP	Enables or disables the static RP feature. You can use static RP to configure a static entry for an RP. A static RP permits communication with switches from other vendors that do not use the BSR mechanism.
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value ages PIM mroutes in seconds. The range is from 10–86400 and the default value is 210. Topology and hardware conditions can affect the polling interval and cause an inactive route to remain for up to 12-15 minutes.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The range is from 1–18724 and the default is 60 seconds.

Enabling PIM on a port

Enable PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you enable it on an interface.
- The interface uses a valid IP address.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **PIM** tab.
5. Select the **Enable** check box.
6. Click **Apply**.

PIM field descriptions

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

Name	Description
Enable	Enables (true) or disables (false) PIM for the specified port.
Mode	Displays the mode currently running on the routing switch.
IntfType	Indicates the interface type as active or passive.

Table continues…

Name	Description
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724 seconds.
CBSRPreference	Configures the preference for this local interface to become a candidate BSR (C-BSR). The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. The range is -1-255.

Enabling IPv6 PIM on a port

Enable IPv6 PIM for each interface to enable the interface to perform multicasting operations.

About this task

You can also right-click the port and use the **Edit IPv6** shortcut menu to reach this same tab.

Before you begin

- You must enable IPv6 interface before you enable PIM on a port.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the **Configuration** > **Edit** > **Port** folders.
3. Click **IPv6**.
4. Click the **PIM** tab.
5. Select **Enable**.
6. Click **Apply**.

PIM field descriptions

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

Name	Description
Address	Specifies the IPv6 address of the PIM interface.
NetMask	Specifies the network mask for the IPv6 address of the PIM interface.
Enable	Enables (true) or disables (false) PIM for the specified port.

Table continues…

Name	Description
Mode	Displays the mode currently running on the routing switch.
DR	Specifies the designated router on this PIM interface.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724 seconds.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724 seconds.
OperState	Specifies the current operational state of this PIM interface.
Type	Specifies the type of interface.

Enabling SSM globally

Enable Source Specific Multicast (SSM) to optimize PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

Before you begin

- Configure a unicast protocol, such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM. For more information about RIP and OSPF, see Configuring OSPF and RIP for VOSS.
- Enable PIM globally.

Important:

After you enable PIM in SSM mode, the IGMP parameters take effect. To take full advantage of SSM, enable IGMPv3 if hosts that attach to the switch run IGMPv3 or configure the SSM table.

About this task

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On an SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. In the navigation pane, expand the **Configuration** > **IP** folders.
2. Click **PIM**.
3. Click the **Globals** tab.
4. Click **ssm** (source specific multicast).
5. Select the **Enable** check box.

6. Click **Apply**.

   The following message appears:

   ```
   Are you sure you want to change the PIM mode? The traffic will not
   be stopped immediately. All Static Source Group entries in the SSM
   range will be deleted. Do you wish to continue?
   ```

7. Click **Yes**.

Enabling IPv6 SSM globally

Enable Source Specific Multicast (SSM) to optimize IPv6 PIM-SM by simplifying the many-to-many model (servers-to-receivers). Because most multicast applications distribute content to a group in one direction, SSM uses a one-to-many model that uses only a subset of the IPv6 PIM-SM features. This model is more efficient and reduces the load on multicast routing devices.

Before you begin

- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM.

  For more information about RIPng and OSPFv3, see Configuring IPv6 Routing for VOSS.

- Enable PIM globally.

  Important:

  After you enable IPv6 PIM in SSM mode, the MLD parameters take effect. To take full advantage of SSM, enable MLDv2 if hosts that attach to the switch run MLDv2.

About this task

SSM is a global configuration. After you enable SSM on a switch, it is enabled on all interfaces that run PIM. On a SSM-enabled switch, SSM behavior is limited to the SSM group range. For non-SSM groups, the protocol behavior is PIM-SM.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Globals** tab.
4. Select the **Enable** check box.
5. Click **ssm** (source specific multicast).
6. Click **Apply**.

   The following message appears:

   ```
   Warning: RP entries in the SSM range will be deleted
   ```

```
Do you wish to continue? (y/n)?
```

7. Click **Yes**.

Enabling PIM on a VLAN interface

Configure PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable PIM globally before you enable it on an interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Click **IP**.
6. Click the **PIM** tab.
7. Select the **Enable** check box.
8. Click **Apply**.

PIM field descriptions

Use the descriptions in the following table to use the **PIM** tab.

Name	Description
Enable	Enables (true) or disables (false) PIM.
Mode	Displays the mode that currently runs on the switch. The valid modes are SSM and Sparse. This variable is a read-only field.
IntfType	Specifies the type of interface: active or passive.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724.
CBSRPreference	Configures the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. The range is -1-255.

Enabling IPv6 PIM on a VLAN interface

Configure IPv6 PIM for each interface to enable the interface to perform multicasting operations.

Before you begin

- You must enable IPv6 PIM globally before you enable it on an interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Click **IPv6**.
6. Click the **PIM** tab.
7. Select the **Enable** check box.
8. Click **Apply**.

PIM field descriptions

Use the descriptions in the following table to use the **PIM** tab.

Name	Description
IfIndex	Specifies the interface index for PIM.
Address	Specifies the IPv6 address of the PIM interface.
Netmask	Specifies the network mask for the IPv6 address of the PIM interface.
Enable	Enables (true) or disables (false) PIM.
Mode	Displays the mode that currently runs on the switch. The valid modes are SSM and Sparse. This variable is a read-only field.
DR	Specifies the designated router on this PIM interface.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring routers. The default is 30 seconds. The range is 0-18724.
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds. The range is 1-18724.
OperState	Specifies the current operational state of this PIM interface.
Type	Specifies the type of interface.

Configuring IGMP parameters on a port

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **IGMP** tab.
5. Edit the appropriate values.

   Note:

   When you configure the following IGMP parameters on the DvR Controllers in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

   - Version
   - QueryInterval
   - QueryMaxResponseTime
   - Robustness
   - LastMembQueryIntvl
   - CompatibilityModeEnable

   For information on DvR, see Configuring IPv4 Routing for VOSS.

   Note:

   To use the fast leave feature on IGMP, enable explicit-host-tracking.
6. Click **Apply**.

IGMP field descriptions

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

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1.

Table continues…

Name	Description
	Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds). **Important:** You must configure this value lower than the QueryInterval.
Robustness	Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in 1/10 seconds) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Configure this parameter to values greater than 3. If you do not require a fast leave process, Use values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
SnoopEnable	Enables snoop on the interface. The default is disabled.
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
ProxySnoopEnable	Enables proxy snoop on the interface. The default is disabled.
Version	Configures the version of IGMP (1, 2 or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this port.
Maximum Number Of Stream	Configures the maximum number of streams this port permits. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This variable is a read-only value.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router. **Important:** Configure this variable only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN

Table continues…

Name	Description
	(technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.
RouterAlertEnable	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. To maximize network performance, configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
DynamicDowngradeEnable	Configures if the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.

Configuring IGMP parameters on a VLAN

Configure IGMP for each interface to enable the interface to perform multicasting operations.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Select **IGMP**.
7. Configure the relevant variables.

Note:

When you configure the following IGMP parameters on the DvR Controllers in a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

- Version
- QueryInterval
- QueryMaxResponseTime
- Robustness
- LastMembQueryIntvl
- CompatibilityModeEnable

Configuration of IGMP snooping is not supported on DvR enabled Layer 2 VSNs.

For information on DvR, see Configuring IPv4 Routing for VOSS.

8. Click **Apply**.

IGMP field descriptions

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

Name	Description
QueryInterval	Configures the frequency (in seconds) at which the IGMP host query packets transmit on the interface. The range is from 1–65535 and the default is 125.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The range is from 0–255 and the default is 100 tenths of a second (equal to 10 seconds.) **Important:** You must configure this value lower than the QueryInterval.
Robustness	Configure this parameter to tune for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect the network to lose query packets, increase the robustness value. The range is from 2–255 and the default is 2. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group

Table continues…

Name	Description
	messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of a second. Configure this parameter to values greater than 3. If you do not require a fast leave process, use values greater than 10. (The value 3 is equal to 0.3 seconds, and 10 is equal to 1 second.)
SnoopEnable	Enables snoop on the interface. The default is disabled.
SsmSnoopEnable	Enables SSM snoop. The default is disabled.
ProxySnoopEnable	Enables proxy snoop on the interface. The default is disabled.
Version	Configures the version of IGMP (1, 2, or 3) that you want to use on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
FastLeaveEnable	Enables or disables fast leave on the interface.
StreamLimitEnable	Enables or disables stream limitation on this VLAN.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this VLAN. The range is from 0–65535 and the default is 4.
Current Number Of Stream	Displays the current number of streams. This value is a read-only value.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router. **Important:** Configure this field only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.
RouterAlertEnable	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. To maximize network performance, configure this parameter according to the version of IGMP currently in use: • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
DynamicDowngradeEnable	Configures if the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with

Table continues…

Name	Description
	IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning. The default value is selected (enabled), which means the switch downgrades to the oldest version of IGMP on the network.
CompatibilityModeEnable	Enables or disables v2-v3 compatibility mode. The default value is clear (disabled), which means IGMPv3 is not compatible with IGMPv2.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts per channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.
SnoopQuerierEnable	Enables snoop querier. The default is disabled. When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups. Enable Layer 2 Querier on only one node in the VLAN.
SnoopQuerierAddr	Specifies the pseudo IP address of the IGMP snoop querier. The default IP address is 0.0.0.0.

Chapter 6: Multicast Listener Discovery

This chapter provides an overview of Multicast Listener Discovery (MLD) snooping for IPv6 multicast traffic.

MLD Fundamentals

MLD is an asymmetric protocol. It specifies separate behaviors for multicast address listeners (that is, hosts or routers that listen to multicast packets) and multicast routers. Each multicast router learns, for each directly attached link, which multicast addresses and which sources have listeners on that link. The information that MLD gathers is provided to the multicast routing protocols that the router uses. This information ensures that multicast packets arrive at all links where listeners require such packets.

A multicast router can itself be a listener of one or more multicast addresses; that is, the router performs both the multicast router role and the multicast address listener part of the protocol. The router collects the multicast listener information needed by the multicast routing protocol and informs itself and other neighboring multicast routers of the listening state.

IPv6 routers use MLD to discover:

- The presence of multicast listeners on directly attached links
- Multicast addresses required by neighboring nodes

MLD versions

The purpose of the MLD protocol in the IPv6 multicast architecture is to allow an IPv6 router to discover the presence of multicast listeners on directly-attached links and to discover which multicast addresses are of interest to neighboring nodes. MLD is the direct IPv6 replacement for the IGMP protocol used in IPv4. The MLD implementation described in this document is based on the MLDv2 standard, which is a backward-compatible update to the MLDv1 standard.

There are three versions of IGMP, and two versions of MLD. IGMPv2 is equivalent in function to MLDv1 and IGMPv3 is equivalent to MLDv2.

MLD Querier

MLD Querier is similar to IGMP querier. A multicast query router communicates with hosts on a local network by sending MLD queries. This router periodically sends a general query message to each local network of the router. This is standard multicast behavior.

Note:

Queries are sent only if PIM is enabled globally and on the interface. PIM and snooping cannot be enabled at the same time.

Each VLAN using MLD multicast must have a router performing multicast queries. Networks with no stand-alone devices currently have no capability for implementing the pruning of multicast traffic. A dedicated querier must be available on the network.

There are several behavioral differences between a traditional query router and a switch or stack using the MLD Querier functionality. The following are the differences:

- There is no election process. When a switch or stack restarts, queries are sent as part of MLD startup. This process stops other devices from sending queries while they detect the new device starting up. The last active device sending queries on the network is the active one. This is not the case with Layer 3 MLD behavior.
- If the current active device stops sending queries, a timeout period must elapse before another device takes over. This can result in an ageout of groups, and subsequent flooding, before a new query is sent and the pruning process restarts. This occurs only during the transition between active query devices. Once the new device is established, queries are sent as configured in the Query Interval and Robust Values fields.
- Multiple active query devices are not supported. Enabling multiple devices establishes one active device and other devices listening to take over should the active device fail.

The querier version is determined by the received query version and establishes the interface operational version. By default, the interface operational version is MLDv1. If the interface operational version is downgraded from MLDv2 to MLDv1 (when operational version is MLDv2 and a MLDv1 query is received), then all MLDv2 listeners (registered by MLDv2 reports) are removed and all incoming MLDv2 reports are dropped.

MLD snooping

MLD snooping is an IPv6 multicast constraining mechanism running on Layer 2 devices. When MLD snooping is enabled on a VLAN, the switch examines the MLD messages between hosts and multicast routers and learns which hosts are interested in receiving traffic for a multicast group. Based on the learning, the switch forwards multicast traffic only to those interfaces in the VLAN that are connected to the interested receivers instead of flooding traffic to all the interfaces.

The following figure shows an example of this scenario. On the left side of the figure, IPv6 multicast packets are transmitted when MLD snooping is not enabled. All the hosts that are interested and not interested receive the IP Multicast traffic consuming bandwidth. Whereas, on the right side of the

figure, when MLD snooping is enabled and IPv6 multicast packets are transmitted, only the interested hosts receive the IP multicast packets.

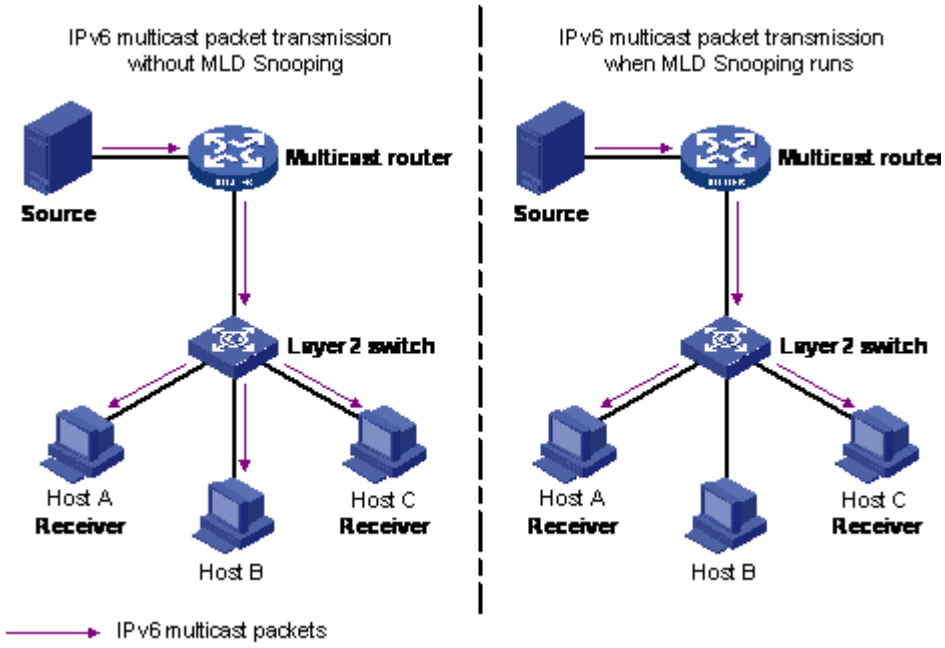

Figure 18: IPv6 multicast packet transmission when MLD snooping is enabled and not enabled

The following figure shows IPv6 multicast packets transmitted when MLD v2 snooping is enabled and not enabled.

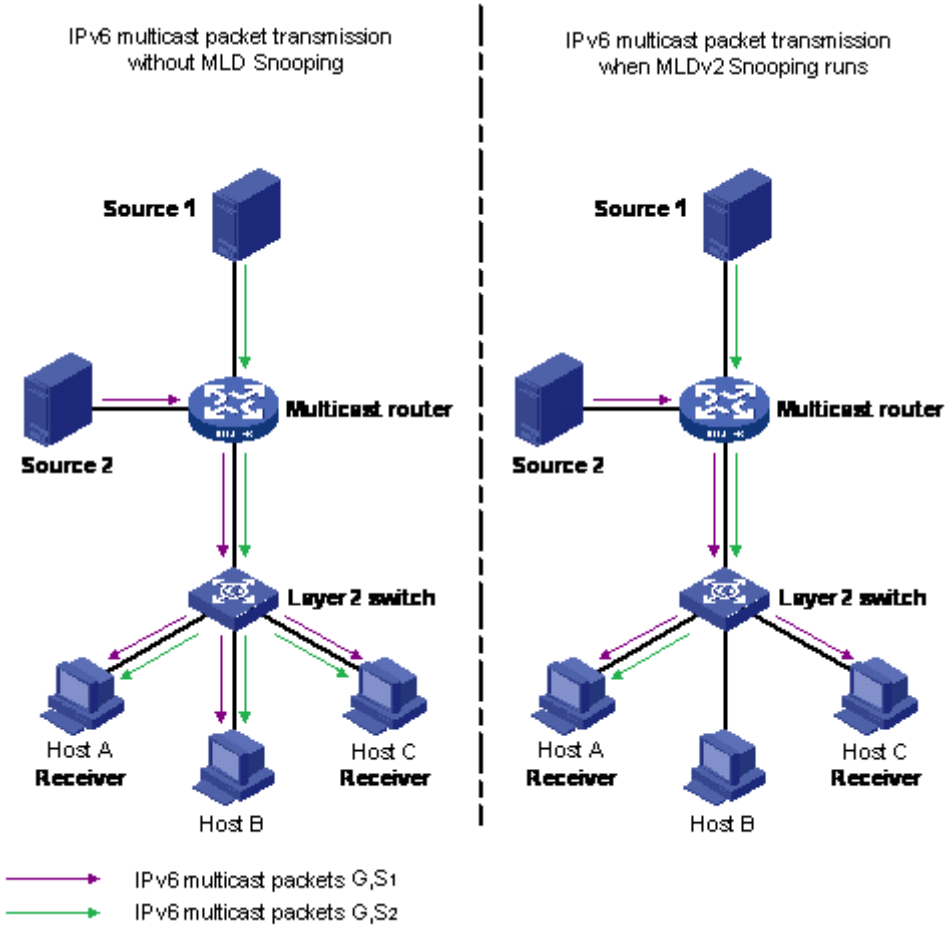

Figure 19: IPv6 multicast packet transmission when MLD v2 snooping is enabled and not enabled

MLD snooping configuration guidelines and restrictions

You can perform the following configurations to manage and control IPv6 multicast groups using the MLD snooping feature:

- Enable or disable MLD snooping on each VLAN. MLD snooping can be enabled on a maximum of 512 VLANs.
- Enable IGMP snooping and MLD snooping on the same VLAN.

Limitations

Following are the limitations for MLD snooping configuration:

- The maximum (S,G,V) entries supported in the IPv6 multicast routing table (L3_ENTRY_IPV6_MULTICAST) is 512.

  MLD snooping shares the (S,G,V) entries with IGMP snooping, where the (S,G,V) entries number = (G,V) MLD_V1 type entries number + (S,G,V) MLD_V2 type entries number + (*,G,V) MLD_V2 type entries number + number of groups without (*,G,V) registered listeners.

- IPv6 MLD proxy functionality is not supported.
- Multicast Flood Control (MFC) is not supported.
- Static mrouter ports cannot be configured.
- IPv6 MLD send query functionality is not supported.
- Configure static router ports is not supported.

MLD configuration using the CLI

This chapter describes the procedures you can use to configure and display Multicast Listener Discovery (MLD) parameters using the CLI.

Configuring MLD trap generation

About this task

Use this procedure to enable MLD traps.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable MLD trap generation:

   ```
   ipv6 mld generate-trap
   ```

3. Disable MLD trap generation:

   ```
   no ipv6 mld generate-trap
   ```

4. Set MLD trap enable status to default:

   ```
   default ipv6 mld generate-trap
   ```

Configuring MLD log status

About this task

Use this procedure to enable MLD traps.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

```
configure terminal
```

2. Enable MLD log status:

```
ipv6 mld generate-log
```

3. Disable MLD log status:

```
no ipv6 mld generate-log
```

4. Set MLD log enable status to default:

```
default ipv6 mld generate-log
```

Configuring MLD version

About this task

Use this procedure to configure MLD version.

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. Configure MLD version:

```
ipv6 mld version <1-2>
```

Note:

For MLD to function correctly, the MLD version must be the same on all routers in the network.

3. Set MLD version to default:

```
default ipv6 mld version
```

Variable definitions

The following table describes the variables for the `ipv6 mld version` command.

Variable	Description
<1–2>	Indicates the version of MLD that runs on this interface.

Configuring the MLD last listener query interval

About this task

Use this procedure to configure the last listener query interval in seconds for the MLD interface.

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. Configure the last listener query interval:

   ```
   ipv6 mld last-listener-query-interval <0-60>
   ```

3. Set the last listener query interval to its default value:

   ```
   default ipv6 mld last-listener-query-interval
   ```

Variable definitions

The following table describes the variables for the `ipv6 mld last-listener-query-interval` command.

Variable	Description
<0–60>	Indicates the last listener query interval in seconds.

Configuring the MLD query interval

About this task

Use this procedure to configure the query interval for the MLD interface.

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. Configure the query interval for the MLD interface:

```
ipv6 mld query-interval <1-65535>
```

3. Set the query interval to its default value:

```
default ipv6 mld query-interval
```

Variable definitions

The following table describes the variables for the `ipv6 mld query-interval` command.

Variable	Description
<1-65535>	Indicates the frequency at which MLD host query packets transmit on this interface.

Configuring the MLD query maximum response time

About this task

Use this procedure to configure the query maximum response time for mld interface.

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. Configure the query maximum response time for mld interface:

```
ipv6 mld query-max-response-time <0-60>
```

3. Set the query maximum response time to its default value:

```
default ipv6 mld query-max-response-time
```

Variable definitions

The following table describes the variables for the `ipv6 mld query-max-response-time` command.

Variable	Description
<0–60>	Indicates the query maximum response interval time in seconds.

Configuring the MLD robustness

About this task

The robustness value allows the tuning for the expected packet loss on a link. If a link expects packet loss, increase the robustness variable value.

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. Configure the MLD robustness:

   ```
   ipv6 mld robust-value <2-255>
   ```

3. Set the MLD robustness to its default value:

   ```
   default ipv6 mld robust-value
   ```

Variable definitions

The following table describes the variables for the `ipv6 mld robust-value` command.

Variable	Description
<2–255>	Specifies a numerical value for MLD snooping robustness.

Enabling MLD snooping on a VLAN

About this task

Use this procedure to enable MLD snooping on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable MLD snooping:

   ```
   ipv6 mld snooping
   ```

3. Set the MLD snooping to its default value:

   ```
   default ipv6 mld snooping
   ```

Enabling MLD ssm-snooping on a VLAN

About this task

Use this procedure to enable IPv6 MLD ssm-snooping on a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable MLD snooping:

   ```
   ipv6 mld ssm-snoop
   ```

3. Set the MLD snooping to its default value:

   ```
   default ipv6 mld ssm-snoop
   ```

Displaying MLD snooping configuration status

About this task

Use this procedure to display information about the MLD snooping configuration for the switch.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the switch MLD snooping configuration status:

   ```
   show ipv6 mld snooping
   ```

Example

```
Switch:1#show ipv6 mld snooping
================================================================================

              Mld Snooping - GlobalRouter
================================================================================

IFINDEX SNOOP   SSM     ACTIVE           MROUTER
        ENABLE  SNOOP   MROUTER          EXPIRATION
                ENABLE  PORTS            TIME
--------------------------------------------------------------------------------
V666        False    False    NONE                        0
1 out of 1 entries displayed
```

Job Aid

The following table describes the column headings in the command output for `show ipv6 mld snooping`.

Variable	Description
IFINDEX	Identifies the index of the physical interface.
SNOOP ENABLE	Identifies whether snoop is enabled (true) or disabled (false).
SSM SNOOP ENABLE	Identifies whether SSM snoop is enabled (true) or disabled (false).
ACTIVE MROUTER PORTS	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
MROUTER EXPIRATION TIME	Specifies the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Displaying MLD snooping tracing information

About this task

Use this procedure to display MLD snooping tracing information.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the MLD snooping tracing information:

   ```
   show ipv6 mld snoop-trace
   ```

Example

```
Switch:1#show ipv6 mld snoop-trace
================================================================
               Mld Snoop Trace - GlobalRouter

================================================================
GROUP/
SOURCE                  IN      IN    OUT     OUT       TYPE
ADDRESS                 VLAN    PORT VLAN     PORT
----------------------------------------------------------------
ff10:0:0:0:0:0:0:1/     10      2/15 10       3/16      ACCESS
5051:0:0:0:0:1:84:51
```

Job Aid

The following table describes the column headings in the command output for `show ipv6 mld snoop-trace`.

Variable	Description
GROUP ADDRESS	Specifies the IP multicast address of the group traversing the router
SOURCE ADDRESS	Specifies the IP source address of the multicast group address.
IN VLAN	Specfies the ingress VLAN ID for the multicast source.
IN PORT	Specifies the ingress port for the multicast group.
OUT VLAN	Specifies the egress VLAN ID for the multicast group.
OUT PORT	Specifies the egress port of the multicast group.
TYPE	Specifies the port type on which the snoop entry is learnt.

Displaying MLD interface information

About this task

Use this procedure to display MLD snooping interface parameters.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display MLD interface information:

   ```
   show ipv6 mld interface [gigabitethernet {slot/port[/sub-port]}]
   [vlan <1-4059>]
   ```

Example

```
Switch:1#show ipv6 mld interface
================================================================================================
                         Mld Interface - GlobalRouter
================================================================================================
IF   STATUS VERS  OPER VERS QUERIER                                  Wrong Query JOINS MODE
------------------------------------------------------------------------------------------------
P6/3 inact  2     2         2001:0db8:3c4d:0015:0000:0000:1a2f:1aaa   0              0     pim
V666 inact  2     2         2001:0db8:3c4d:0015:0000:0000:1a2f:1bbb   0              0     pim

Switch:1#show ipv6 mld interface vlan 10
================================================================================
                                  Vlan IPv6 Mld
================================================================================
VLAN QUERY QUERY ROBUST VERSION LAST  SNOOP  SSM     DYNAMIC
ID   INTVL MAX                  LIST  ENABLE SNOOP   DOWNGRADE
           RESP                 QUERY        ENABLE
--------------------------------------------------------------------------------
10   125   10    2      1       1     false  false   enabled

Switch(config)#show ipv6 mld interface gigabitethernet 1/11
================================================================================
                                 Port IPv6 MLD

================================================================================
PORT QUERY QUERY ROBUST VERSION LAST  DYNAMIC
NUM  INTVL MAX                  LIST  DOWNGRADE
           RESP                 QUERY
--------------------------------------------------------------------------------
1/11 125   10    2      1       1     enabled

1 out of 1 entries displayed
```

Variable definitions

The following table describes the variables for the `show ipv6 mld interface` command.

Variable	Description
vlan *<1-4059>*	Displays MLD snooping information for the configured VLANs.
gigabitEthernet *{slot/port[/sub-port]}*	Displays MLD snooping information on a specific interface.

Job Aid

The following table describes the column headings in the command output for `show ipv6 mld interface`.

Variable	Description
VLAN ID	Indicates the VLAN ID of the physical interface.
PORT NUM	Indicates the port number of the physical interface.

Table continues…

Variable	Description
QUERY INTVL	Indicates the query interval, the frequency at which IPv6 MLD snooping host-query packets are transmitted on this interface.
QUERY MAX RESP	Indicates the maximum query response time advertised in IPv6 MLD snooping queries on this interface.
ROBUST	Indicates the robustness value.
VERSION	Indicates the version.
LAST LIST QUERY	Indicates the last listener query interval. The last listener query interval is the maximum response delay inserted into group-specific queries sent in response to leave group messages, and it is also the amount of time between group-specific query messages.
SNOOP ENABLE	Indicates if snooping is enabled.
SSM SNOOP ENABLE	Indicates if ssm-snooping is enabled.
DYNAMIC DOWNGRADE ENABLE	Enables dynamic downgrade of the MLD version when older version query message is received.

Displaying MLD system parameters

About this task

Use this procedure to display information about the MLD traps and logs.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the system parameters:

   ```
   show ipv6 mld sys
   ```

Example

```
Switch:1#show ipv6 mld sys
================================================================================
                      Mld System Parameters - GlobalRouter
================================================================================
                     generate-trap : disable
                      generate-log : disable
```

Displaying MLD cache information

About this task

Use this procedure to display the learned multicast groups in the cache.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the learned multicast groups in the cache:

   ```
   show ipv6 mld cache
   ```

Example

```
Switch:1#show ipv6 mld cache
================================================================================================
                              MLD Cache Information

================================================================================================
================================================================================================
 GRPADDRESS/LASTREPORTER                  INTERFACE EXPIRATION
------------------------------------------------------------------------------------------------
 ff03:0:0:0:0:0:0:0/                      Vlan10    0 day(s), 00h:04m:12s
 fe80:0:0:0:200:9aff:fe68:3dd5

1 out of 1 entries displayed
```

Job Aid

The following table describes the column headings in the command output for `show ipv6 mld cache`.

Variable	Description
GRPADDR	Indicates the IPv6 address of the multicast address of interest.
LASTREPORTER	Indicates the IPv6 address of the last reporter.
INTERFACE	Indicates the ingress interface for MLDv2.

Displaying the MLD group information

About this task

Use this procedure to display the MLD group information to show the learned multicast groups and the attached ports.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display the MLD group information:

   ```
   show ipv6 mld group [count] [group] [member-subnet]
   ```

Example

```
Switch:1#show ipv6 mld group
================================================================================
                          Mld Group - GlobalRouter
================================================================================
GRPADDR/MEMBER                                  INPORT             EXPIRATION
--------------------------------------------------------------------------------
ff1e:0000:0000:0000:0000:0000:0002:4444/        V666-6/41          0
2001:0db8:3c4d:0015:0000:0000:1a2f:1a2c

1 out of 1 group Receivers displayed

Total number of unique groups 1

Switch:1#show ipv6 mld group group ff1e:0000:0000:0000:0000:0000:0002:4444 detail

================================================================================
                       Mld Group Detail - GlobalRouter
================================================================================

Interface:                    Vlan666-6/41
MLDv2 Group:                  ff1e:0000:0000:0000:0000:0000:0002:4444
Interface Group Mode:         EXCLUDE
Interface Compatibility Mode: MLD_V2
Interface Group Timer:        258
V1 Host Timer:                Not Running
Interface Group Include Source List:
        Source Address                                  Expires
        2001:0db8:3c4d:0015:0000:0000:1a2f:1aaa         258
Interface Group Exclude Source List :
        Source Address                                                  Expires
        2001:0db8:3c4d:0015:0000:0000:1a2f:1bbb         N/A
```

Job Aid

The following table describes the column headings in the command output for `show ipv6 mld group`.

Variable	Description
GRPADDR	Specifies the multicast group address that others want to join to. A group address can be the same for many incoming ports.
MEMBER	Specifies the IP address of a source that has sent group report whishing to join this group.
INPORT	Identifies a physical interface or a logical interface which has received group reports from various sources.
EXPIRATION	Specifies the time left before group report expires on this port. This is updated upon receiving a group report.

Viewing IPv6 MLD host cache

View the learned multicast group addresses in the host cache.

Procedure

1. Log on to the switch to enter User EXEC mode.
2. Show IPv6 MLD host cache:

   ```
   show ipv6 mld-host-cache
   ```

3. Show IPv6 MLD host cache for a management interface:

   ```
   show ipv6 mld-host-cache mgmtEthernet [mgmt]
   ```

Example

```
Switch:1#show ipv6 mld-host-cache
=======================================================
                   MLD Cache Information
=======================================================
PORT/VID  GRPADDRESS                                SELF
-------------------------------------------------------
mgmt      ff02::1:ff00:3                          enabled
mgmt      ff02::1:ff4c:9400                       enabled
mgmt      ff02::1                                 enabled
```

Job aid

The following table describes the fields in the output for the **`show ipv6 mld-host-cache`** command.

Parameter	Description
GRPADDRESS	Shows the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
SELF	Indicates whether or not the group is locally registered into MLD Host Cache. • enabled — group is locally registered into MLD Host Cache • disabled — group is not locally registered into MLD Host Cache
PORT/VID	Shows the port or VLAN that learns the multicast group.

MLD configuration using EDM

This chapter describes the procedures you can use to configure and display Multicast Listener Discovery (MLD) snooping parameters using Enterprise Device Manager (EDM).

Configuring MLD globally

About this task

Use the following procedure to configure MLD parameters for the switch.

Procedure

1. In the navigation pane, expand **Configuration** > **IPv6** folders.
2. Click **IPv6 MLD**.
3. Click the **Globals** tab.
4. Configure the MLD global parameters as required.
5. On the toolbar, click **Apply** to save the changes.
6. On the toolbar, click **Refresh** to update the changes.

Globals field description

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

Field	Description
GenerateTrap	Enables MLD to generate traps.
GenerateLog	Enables MLD to generate logs.

Viewing the MLD SSM global information

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click the **Ssm Globals** tab.

Ssm Globals field description

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

Field	Description
RangeGroup	Specifies the ssm range.
RangeMask	Specifies the ssm range mask.

MLD interface configuration

Configure the interfaces so that the switch forwards multicast traffic only to those interfaces in the VLAN that are connected to the interested receivers instead of flooding traffic to all the interfaces.

Configuring an MLD interface

Perform this procedure to change the configuration of existing MLD interfaces.

Procedure

1. In the navigation pane, expand the **Configuration** > **IPv6** folders.
2. Click **IPv6 MLD**.
3. Click the **Interfaces** tab.
4. On the toolbar, click **Insert**.
5. Configure the MLD interface parameters.
6. Click **Insert**.
7. On the toolbar, click **Apply** to save the changes.
8. On the toolbar, click **Refresh** to update the changes.

MLD interfaces field description

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

Field	Description
IfIndex	Specifies the internetwork layer interface value of the interface for which MLD is enabled.
QueryInterval	Specifies the frequency at which MLD host-query packets are transmitted on this interface. Values range from 1 to 65535.
Version	Indicates the MLD version.
Querier	Specifies the address of the MLD Querier on the IPv6 subnet to which this interface is attached.
QueryMaxResponseDelay	Specifies the maximum query response time advertised in MLD queries on this interface. Values range from 0 to 60.
Joins	Specifies the number of times a group membership has been added on this interface.
Groups	Specifies the current number of entries for this interface in the cache table.
Robustness	Specifies the robustness variable tuning for the expected packet loss on a subnet. If a subnet is

Table continues…

Field	Description
	expected to experience loss, the robustness variable can be increased. Values range from 2 to 255.
LastListenQueryIntvl	Specifies the maximum response delay inserted into the group-specific queries sent in response to the leave group messages. It also indicates the amount of time between group-specific query messages. Values range from 0 to 60. This value can be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
SnoopEnable	Indicates if snooping is enabled.
FlushAction	Specifies the MLD flush action as one of the following: • flushGrpMember • flushMrouter • flushSender
SsmEnable	Indicates if ssm is enabled.
NewQuerier	Specifies the IPv6 address of the new MLD querier.
DynamicDowngradeEnable	Enables dynamic downgrade of the MLD version when older version query message is received.
OperVersion	Specifies the operational version of the MLD running on this interface.
McastMode	Specifies the MLD interface mode as one of the following: • snoop • pim • snoopSpb • routerSpb • dvmrp • none

Configuring MLD on a port

Configure the MLD on a port.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the **Configuration** > **Edit** > **Port** folders.
3. Click **IPv6**.
4. Click the **MLD** tab.

5. Configure the MLD interface parameters.
6. On the toolbar, click **Apply** to save the changes.
7. On the toolbar, click **Refresh** to update the changes.

MLD field description

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

Field	Description
QueryInterval	Specifies the frequency at which MLD host-query packets are transmitted on this interface. Values range from 1 to 65535.
Version	Indicates the MLD version.
Querier	Specifies the address of the MLD Querier on the IPv6 subnet to which this interface is attached.
QueryMaxResponseDelay	Specifies the maximum query response time advertised in MLD queries on this interface. Values range from 0 to 60.
Joins	Specifies the number of times a group membership has been added on this interface.
Groups	Specifies the current number of entries for this interface in the cache table.
Robustness	Specifies the robustness variable tuning for the expected packet loss on a subnet. If a subnet is expected to experience loss, the robustness variable can be increased. Values range from 2 to 255.
LastListenQueryIntvl	Specifies the maximum response delay inserted into the group-specific queries sent in response to the leave group messages. It also indicates the amount of time between group-specific query messages. Values range from 0 to 60. This value can be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
SnoopEnable	Indicates if snooping is enabled.
FlushAction	Specifies the MLD flush action as one of the following: • flushGrpMember • flushMrouter • flushSender
SsmEnable	Indicates if ssm is enabled.
NewQuerier	Specifies the IPv6 address of the new MLD querier.

Table continues…

Field	Description
DynamicDowngradeEnable	Enables dynamic downgrade of the MLD version when older version query message is received.
OperVersion	Specifies the operational version of the MLD running on this interface.
McastMode	Specifies the MLD interface mode as one of the following: • snoop • pim • snoopSpb • routerSpb • dvmrp • none

Configuring MLD on a VLAN

About this task

Configure MLD on a VLAN.

Procedure

1. In the navigation pane, expand the **Configuration** > **VLAN** folders.
2. Click **VLANs**.
3. Select a VLAN from the list.
4. Click the **IPv6** tab.
5. Click the **MLD** tab.
6. Configure the MLD interface parameters.
7. On the toolbar, click **Apply** to save the changes.
8. On the toolbar, click **Refresh** to update the changes.

MLD field description

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

Field	Description
QueryInterval	Specifies the frequency at which MLD host-query packets are transmitted on this interface. Values range from 1 to 65535.
Version	Indicates the MLD version.
Querier	Specifies the address of the MLD Querier on the IPv6 subnet to which this interface is attached.

Table continues…

Field	Description
QueryMaxResponseDelay	Specifies the maximum query response time advertised in MLD queries on this interface. Values range from 0 to 60.
Joins	Specifies the number of times a group membership has been added on this interface.
Groups	Specifies the current number of entries for this interface in the cache table.
Robustness	Specifies the robustness variable tuning for the expected packet loss on a subnet. If a subnet is expected to experience loss, the robustness variable can be increased. Values range from 2 to 255.
LastListenQueryIntvl	Specifies the maximum response delay inserted into the group-specific queries sent in response to the leave group messages. It also indicates the amount of time between group-specific query messages. Values range from 0 to 60. This value can be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
SnoopEnable	Indicates if snooping is enabled.
FlushAction	Specifies the MLD flush action as one of the following: • flushGrpMember • flushMrouter • flushSender
SsmEnable	Indicates if ssm is enabled.
NewQuerier	Specifies the IPv6 address of the new MLD querier.
DynamicDowngradeEnable	Enables dynamic downgrade of the MLD version when older version query message is received.
OperVersion	Specifies the operational version of the MLD running on this interface.
McastMode	Specifies the MLD interface mode as one of the following: • snoop • pim • snoopSpb • routerSpb • dvmrp • none

Configuring MLD snooping

About this task

Use the following procedure to enable MLD snooping on the switch.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click **Snooping** tab.
4. Select a value, double-click the cell in **SnoopEnable** column, select **True** or **False**.
5. Select a value, double-click the cell in **SsmEnable** column, select **True** or **False**.
6. Click **Apply**.

Snooping field description

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

Field	Description
IfIndex	Specifies the interface on which you enabled MLD snooping. It specifies the port number if the interface is a brouter port or the VLAN number if the interface is a VLAN.
SnoopEnable	Indicates the status of MLD snooping on the specified interface: • True – MLD snooping is enabled • False – MLD snooping is disabled
SsmEnable	Indicates the status of SSM on the specified interface: • True – SSM is enabled • False – SSM is disabled

Viewing the MLD snoop trace information

About this task

Use this procedure to display information about the multicast groups traversing the snoop enabled router.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.

2. Click **IPv6 MLD**.
3. Click the **Snoop Trace** tab.

Snoop Trace field description

Use the data in the following table to use the **Snoop Trace** tab.

Field	Description
GrpAddr	Specifies the IP multicast address of the group traversing the router.
SrcAddr	Specifies the IP source address of the multicast group address.
OutVlan	Specifies the egress VLAN ID for the multicast group.
OutPort	Specifies the egress port of the multicast group.
InVlan	Specfies the ingress VLAN ID for the multicast source.
InPort	Specifies the ingress port for the multicast group.
Type	Specifies the port type on which the snoop entry is learnt.

Viewing the MLD cache information

About this task

Use this procedure to display information about the learned multicast groups in the cache.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click the **Cache** tab.

MLD cache field description

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

Field	Description
Address	The IPv6 multicast group address for which this entry contains information.
IfIndex	Indicates the internetwork-layer interface for which this entry contains information for an IPv6 multicast group address.

Table continues…

Field	Description
LastReporter	Indicates the source IPv6 address of the last membership report received for this IPv6 Multicast group address on this interface. If membership report is not received, the value is 0::0
ExpiryTime	Indicates the minimum amount of time remaining before the entry ages out.

Viewing the MLD V2 cache information

About this task

Use this procedure to display information about the MLDv2 corresponding to each interface, port and multicast group paired on a router.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click the **V2 Cache** tab.

V2 Cache field description

Use the data in the following table to use the **V2 Cache** tab.

Field	Description
GroupAddress	Specifies the multicast group address that others want to join. A group address can be the same for many incoming ports.
Ifindex	Identifies a physical interface or a logical interface (VLAN), which has received group reports from various sources.
InPort	Identifies a physical interface or a logical interface (VLAN), which has received group reports from various sources.
Version1HostTimer	Specifies the time remaining until the local router assumes that there are no more MLDv1 members on the IP subnet attached to the interface. This is applicable only for MLDv1 hosts. Upon receiving an MLDv1 report, this value is reset to the group membership timer.
SourceFilterMode	Specifies the current group state applicable on MLDv2 compatible nodes.

Viewing IPv6 MLD host cache

View the learned multicast group addresses in the host cache.

Note:

Not all hardware platforms include a dedicated, physical management interface. For more information about supported interfaces, see your hardware documentation.

Procedure

1. In the navigation tree, expand the **Configuration** > **IPv6** folders.
2. Click **IPv6 MLD**.
3. Click the **Host Cache** tab.

MLD host cache field descriptions

Use the data in the following table to use the Host Cache tab.

Name	Description
IfIndex	Shows the index value that uniquely identifies the interface to which this entry applies.
GrpAddress	Shows the IP address for the multicast group.
GrpLocallyRegistered	Shows the Group Locally Registered for an IPv6 MLD host-cache entry.
GrpLastReporter	Shows the Group Last Reporter address for an IPv6 MLD host-cache entry.
GrpUpTime	Shows the Group Uptime for an IPv6 MLD host-cache entry.
GrpExpiryTime	Shows the Group Expiry Time for an IPv6 MLD host-cache entry.
GrpFilterMode	Shows the Group Filter Mode for an IPv6 MLD host-cache entry.

Viewing the MLD source information

About this task

Use this procedure to display information about the MLD source.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click the **Source** tab.

Source field description

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

Field	Description
GroupAddress	Specifies the IPv6 multicast group address for which this entry contains information.
Ifindex	Specifies the interface for which this entry contains information for an IP multicast group address.
InPort	Identifies a physical interface or logical interface (VLAN), which has received group reports for this source.
HostAddress	Specifies the host address to which this entry corresponds.
MemberAddress	Specifies the IPv6 address of a member that has sent source specific report wishing to join this source.
Expire	Specifies the state of this entry.
Mode	Specifies the current member state. This is applicable to MLDv2 compatible nodes.
MemberExpire	Specifies the time until the member for this source expires.

Viewing the MLD sender information

About this task

Use this procedure to display information about the multicast senders.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click the **Sender** tab.

Source field description

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

Field	Description
GrpAddr	Specifies the IPv6 multicast group address.
Ifindex	Specifies the interface index of the sender.
MemberAddr	Specifies the IPv6 host address.

Table continues…

Field	Description
Action	Specifies the MLD action as one of the following: • none • flushEntry • flushGrp
Port	Specifies the MLD sender port.

Viewing the MLD group information

About this task

Use this procedure to display information about the groups configured in this device.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 MLD**.
3. Click the **Group** tab.

Group field description

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

Field	Description
IPv6Address	Specifies the multicast group address that others want to join to. A group address can be the same for many incoming ports.
Members	Specifies the IP address of a source that has sent group report whishing to join this group.
InPort	Identifies a physical interface or a logical interface which has received group reports from various sources.
Expiration	Specifies the time left before group report expires on this port. This is updated upon receiving a group report.
IfIndex	Identifies a physical interface or a logical interface which has received group reports from various sources.

Chapter 7: PIM configuration using the CLI

The switch supports two modes of Protocol Independent Multicast (PIM): Sparse Mode (SM) and Source Specific Multicast (SSM).

- PIM-SM supports multicast groups spread out across large areas of a company or the Internet.
- PIM-SSM optimizes PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Important:

The **spbm-config-mode** boot flag must be disabled before you can configure PIM or IGMP. To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

Before you begin

For an IPv4 PIM configuration using the CLI:

- Configure an IPv4 interface.

  For more information, see Configuring IPv4 Routing for VOSS.
- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM.

  For more information about RIP and OSPF, see Configuring OSPF and RIP for VOSS.
- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- You must first configure and enable PIM on an IP interface, which can be circuitless, before you can utilize that interface as a candidate rendezvous point (RP). To configure PIM-SM RP for an IP interface, see Configuring a candidate rendezvous point on page 165.
- Configure one or more bootstrap routers (BSR) to propagate RP information to all switches in the network.

For an IPv6 PIM configuration using the CLI:

- Configure an IPv6 interface.

  For more information, see Configuring IPv6 Routing for VOSS.
- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where you want to configure PIM.

  For more information about RIPng and OSPFv3, see Configuring OSPF and RIP for VOSS.
- Enable IPv6 PIM-SM globally

- Enable IPv6 PIM-SM on individual interfaces.

Changing the interface status to passive

Change the PIM interface status to passive to deny PIM control traffic on the interface.

Before you begin

- The PIM interface is disabled.

About this task

The command you use depends on the required administrative state of the interface (enable or disable).

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. Create a passive interface and enable it simultaneously:

   ```
   ip pim passive
   ```

3. Create a passive interface in the disabled state:

   ```
   ip pim interface-type passive
   ```

   You must manually enable the interface.

4. Enable a disabled interface:

   ```
   ip pim enable
   ```

Variable definitions

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

Variable	Value
active	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active. To configure this option to the default value, use the default operator with the command.
passive	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active. To configure this option to the default value, use the default operator with the command.

Changing the interface status to active

Change the PIM interface status to active to allow PIM control traffic on the interface.

Before you begin

- The PIM interface is disabled.

About this task

The command you use depends on the required administrative state of the interface (enable or disable).

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. Create an active interface in the disabled state:

   ```
   ip pim interface-type active
   ```

   You must manually enable the interface.

3. Create an active interface and enable it simultaneously:

   ```
   ip pim active
   ```

   OR

   ```
   ip pim enable
   ```

   The second command enables an active interface only if this is the first PIM interface you create on the port or VLAN or you created an active interface in the disabled state. If you already created a passive interface in the disabled state, the second command enables that passive interface.

Variable definitions

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

Variable	Value
active	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to other switches. The default is active. To configure this option to the default value, use the default operator with the command.
passive	Configures the selected interface. You can change the state of a PIM interface after you create the interface but only if you first disable PIM on the interface. An active interface permits PIM control transmitted and received traffic. A passive interface prevents PIM control traffic from transmitting or receiving, thereby reducing the load on a system. This feature is useful if a high number of PIM interfaces exist and connect to end users, not to

Table continues…

Variable	Value
	other switches. The default is active. To configure this option to the default value, use the default operator with the command.

Configuring the PIM virtual neighbor

Configure a PIM virtual neighbor if the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the PIM virtual neighbor:

   ```
   ip pim virtual-neighbor <A.B.C.D> <A.B.C.D>
   ```

Example

Configure the PIM virtual neighbor:

```
Switch:1(config)#ip pim virtual-neighbor 192.0.2.0 198.51.100.0
```

Variable definitions

Use the definitions in the following table to use the **`ip pim virtual-neighbor`** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	The first IP address indicates the IP address of the selected interface. The second IP address indicates the IP address of the neighbor.

Configuring a candidate rendezvous point

Configure a candidate rendezvous point (C-RP) to serve as backup to the RP router.

About this task

You can configure only one interface on the switch for multiple groups. You cannot configure multiple interfaces for multiple groups.

With the mask value, you can configure a C-RP router for several groups in one configuration.

For example, if you use a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Add a candidate rendezvous point:

   ```
   ip pim rp-candidate group <A.B.C.D> <A.B.C.D> rp <A.B.C.D>
   ```

3. Remove a candidate rendezvous point:

   ```
   no ip pim rp-candidate group <A.B.C.D> <A.B.C.D>
   ```

4. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ip pim rp-candidate
   ```

Example

Add a candidate rendezvous point:

```
Switch:1(config)#ip pim rp-candidate group 224.1.1.0 255.255.255.0 rp 198.51.100.0
```

Variable definitions

Use the definitions in the following table to use the **`ip pim rp-candidate`** command.

Variable	Value
group {A.B.C.D} {A.B.C.D}	Specifies the IP address and the address mask of the multicast group. After the IP address and group mask are combined, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
rp {A.B.C.D}	Specifies the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Job aid

The following table shows the field descriptions for the **`show ip pim rp-candidate`** command.

Table 5: show ip pim rp-candidate field descriptions

Field	Description
GRPADDR	Displays the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Displays the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
RPADDR	Displays the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Configuring static RP

Configure a static RP to ignore the bootstrap router (BSR) mechanism and use the statically configured RPs.

Before you begin

- Enable PIM-SM globally.

About this task

Static RP-enabled switches use this feature to communicate with switches from other vendors that do not use the BSR.

Important:

You cannot configure a static RP-enabled switch as a BSR or as a C-RP router.

All dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable static RP:

   ```
   ip pim static-rp
   ```

   The following message appears:

   ```
   WARNING: RP information learnt dynamically through BSR functionality will be lost.
   Do you wish to enable Static RP? (y/n) ?
   ```

3. Enter `y`.

4. Configure a static RP entry:

```
ip pim static-rp {A.B.C.D/X} {A.B.C.D}
```

5. Configure all the switches in the network (including switches from other vendors) to map to the same RP.
6. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ip pim static-rp
   ```

Example

Configure a static RP:

```
Switch:1(config)#  ip pim static-rp 239.255.0.0/255.255.0.0 198.51.100.0
```

Variable definitions

Use the definitions in the following table to use the **`ip pim static-rp`** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and address mask of the multicast group. When combined, the IP address and address mask identify the range of the multicast addresses that the RP handles.
{A.B.C.D}	Specifies the IP address of the static RP.

Configuring IPv6 PIM static RP

On IPv6 PIM BSR mechanism is not supported so static RP must be configured.

Before you begin

Enable IPv6 PIM-SM globally.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable static RP:

   ```
   ipv6 pim static-rp
   ```

3. Configure an IPv6 static RP entry:

   ```
   ipv6 pim static-rp WORD<0-255> WORD<0-255>
   ```

4. Configure all the switches in the network (including switches from other vendors) to map to the same RP.

5. Display information about the candidate rendezvous points for the PIM-SM domain:

   ```
   show ipv6 pim static-rp
   ```

Variable definitions

The following table describes the variables for the `ipv6 pim static-rp` command.

Variable	Description
WORD<0-255>	Specifies the IPv6 address and address mask of the multicast group. When combined, the IPv6 address and address mask identify the range of the multicast addresses that the RP handles.
WORD<0-255>	Specifies the IPv6 address of the static RP.

Configuring a candidate BSR on a port

Configure additional routers as candidate BSRs (C-BSR) to provide backup protection in the event that the primary BSR fails. PIM-SM cannot run without a BSR.

Before you begin

- Static RP is disabled.

About this task

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

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 candidate BSR:

   ```
   ip pim bsr-candidate preference <0-255>
   ```

Example

Configure a candidate BSR:

```
Switch:1(config-if)#ip pim bsr-candidate preference 2
```

Variable definitions

Use the definitions in the following table to use the **ip pim bsr-candidate** command.

Variable	Value
preference *<0–255>*	Activates the C-BSR on this interface and configures its preference value, from 0–255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. To set this option to the default value, use the default operator with the command.

Configuring a candidate BSR on a VLAN

Configure additional routers as candidate BSRs (C-BSR) to provide backup protection in the event that the primary BSR fails. PIM-SM cannot run without a BSR.

Before you begin

- Static RP is disabled.

About this task

The C-BSR with the highest configured preference becomes the BSR for the domain. If two C-BSRs use equal preference, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher preference to the domain, it automatically becomes the new BSR.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Configure a candidate BSR on a VLAN:

   ```
   ip pim bsr-candidate preference <0-255>
   ```

Example

Configure a candidate BSR on a VLAN:

```
Switch:1(config-if)#ip pim bsr-candidate preference 5
```

Variable definitions

Use the definitions in the following table to use the **`ip pim bsr-candidate`** command.

Variable	Value
preference *<0–255>*	Activates the C-BSR on this interface and configures its preference value, from 0–255, to become a BSR. The C-BSR with the highest BSR preference and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR. To configure this option to the default value, use the default operator with the command.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

About this task

Important:

The following command also activates full-mesh configurations.

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable square-SMLT:

   ```
   multicast smlt-square
   ```

Chapter 8: PIM configuration using EDM

The switch supports two modes of Protocol Independent Multicast (PIM): Sparse Mode (SM) and Source Specific Multicast (SSM).

- PIM-SM supports multicast groups spread out across large areas of a company or the Internet.
- PIM-SSM optimizes PIM-SM by simplifying the many-to-many model (servers-to-receivers).

Important:

The **`EnableSpbmConfigMode`** boot flag must be disabled before you can configure PIM or IGMP. To verify the setting, navigate to **Configuration** > **Edit** > **Chassis** and click on the **Boot Config** tab.

Before you begin

For an IPv4 PIM configuration using EDM:

- Configure an IP interface.

  For more information, see Configuring IPv4 Routing for VOSS.
- Configure a unicast protocol, for example, Routing Information Protocol (RIP) or Open Shortest Path First (OSPF), globally and on the interfaces where you want to configure PIM-SM.

  For more information about RIP and OSPF, see Configuring OSPF and RIP for VOSS.
- Enable PIM-SM globally.
- Enable PIM-SM on individual interfaces.
- Configure one or more rendezvous points (RP) for the groups that multicast applications use in the network.

  Important:

  If you configure the rendezvous point (RP) to be the address of a circuitless IP (CLIP) interface, then you must first configure and enable PIM on the CLIP interface before you can utilize that interface as a candidate RP. To configure a PIM-SM RP for a circuitless IP interface, see Configuring a candidate RP on page 185.
- Configure one or more bootstrap routers (BSR) to propagate RP information to all switches in the network.

For an IPv6 PIM configuration using EDM:

- Configure an IPv6 interface. For more information, see Configuring IPv6 Routing for VOSS.
- Configure an IPv6 unicast protocol, for example, Routing Information Protocol Next Generation (RIPng) or Open Shortest Path First Version 3 (OSPFv3), globally and on the interfaces where

you want to configure PIM. For more information about RIPng and OSPFv3, see Configuring OSPF and RIP for VOSS.

- Enable IPv6 PIM-SM globally.
- Enable IPv6 PIM-SM on individual interfaces.

Enabling static RP

Enable static RP to avoid the process of selecting an active RP from the list of candidate RPs and dynamically learning about RPs through the BSR mechanism.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Globals** tab.
4. Select **sm** (sparse mode).
5. Select **Enable**.
6. Select **Static RP**.
7. Click **Apply**.

   The following message appears:

   ```
   RP information learnt dynamically through BSR functionality will be
   lost. Do you wish to enable Static RP?
   ```

8. Click **Yes**.

Enabling IPv6 static RP

Use this procedure to enable IPv6 static RP.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Globals** tab.
4. Select **sm** (sparse mode).
5. Select **Enable**.
6. Select **Static RP**.

7. Click **Apply**.
8. Click **Yes**.

Configuring a static RP

Configure a static RP to ignore the BSR mechanism and use the statically configured RPs only. A static RP-enabled switch uses this feature to communicate with switches from other vendors that do not use the BSR mechanism.

Before you begin

- Before you can configure a static RP, you must enable the following:
  - PIM-SM
  - static RP

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Static RP** tab.
4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Static RP field descriptions

Use the descriptions in the following table to use the **Static RP** tab.

Name	Description
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, this value identifies the range of the multicast addresses that the RP handles.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the range of the multicast addresses that the RP handles.
Address	Configures the IP address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid if the switch uses a unicast route to the network for the static RP and is invalid otherwise.

Job aid

Keep in mind the following configuration considerations:

- Static RPs do not age; they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless you configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP for certain group range.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of vendor switches across the network, ensure that all switches or routers use the same active RP because vendors use different algorithms to elect the active RP. This switch uses the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.
- Static RP on the switch is active as long as the switch uses a unicast route to the network for the static RP. If the switch loses this route, the static RP is invalidated, and the hash algorithm is invoked to remap all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm is invoked to remap the affected groups.

Configuring an IPv6 static RP entry

Configure an IPv6 static RP to use the statically configured RPs. A static RP-enabled switch uses this feature to elect the active RP only from the statistically configured switches, without any relation to the RP information of other switches.

Before you begin

- Before you can configure a static RP, you must enable the following:
  - IPv6 PIM-SM
  - IPv6 static RP

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Static RP** tab.
4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Static RP field descriptions

Use the descriptions in the following table to use the **Static RP** tab.

Name	Description
GroupAddress	Configures the IPv6 address of the multicast group. When combined with the group mask, this value identifies the range of the multicast addresses that the RP handles.
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the range of the multicast addresses that the RP handles.
Address	Configures the global IPv6 address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid if the switch uses a unicast route to the network for the static RP.

Job aid

Keep in mind the following configuration considerations:

- Static RPs do not age; they cannot time out.
- Switches do not advertise static RPs, so, if a new PIM neighbor joins the network, it does not know about the static RP unless you configure it with that static RP.
- Configure all the switches in the network (including switches from other vendors) to map to the same RP for certain group range.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If you use a mix of vendor switches across the network, ensure that all switches or routers use the same active RP because vendors use different algorithms to elect the active RP. This switch uses the hash function defined in the PIM-SM standard to elect the active RP; other vendors can use the lowest IP address to elect the RP.
- Static RP on the switch is active as long as the switch uses a unicast route to the network for the static RP. If the switch loses this route, the static RP is invalidated, and the hash algorithm is invoked to remap all affected groups. If the switch regains this route, the static RP is validated and the hash algorithm is invoked to remap the affected groups.

Viewing the active RP

Perform this procedure to show information about the active RP for all the running multicast groups on the switch.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Active RP** tab.

Active RP field descriptions

Use the data in the following table to use the **Active RP** tab.

Name	Description
GroupAddress	Shows the IP address of the multicast group.
Address	Shows the IP address of the RP router. This address must be one of the local PIM-SM enabled interfaces.
Priority	Shows the priority of the RP.

Viewing the IPv6 active RP

Perform this procedure to show information about the IPv6 active RP for all the running multicast groups on the switch.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Active RP** tab.

Active RP field descriptions

Use the data in the following table to use the **Active RP** tab.

Name	Description
GroupAddress	Shows the IPv6 address of the multicast group.
Address	Shows the IPv6 address of the RP router. This address can be one of the local PIM-SM enabled interfaces or any reachable global IPv6 address configured using the **`static-rp`** CLI command.

Table continues…

Name	Description
	Note: IPv6 link local address is always used as the PIM interface address.
Priority	Shows the priority of the RP.

Configuring a candidate bootstrap router

Configure routers as candidate bootstrap routers (C-BSR) to provide backup protection in case the primary BSR fails. PIM-SM cannot operate without a BSR. A PIM-SM domain can use only one active BSR.

About this task

The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs use equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **PIM** tab.
5. Click **Enable**.
6. In the **CBSRPreference** box, type the preference.

   The C-BSR with the highest BSR-preference and address becomes the active BSR. The default is –1, which indicates that the current interface is not a C-BSR.
7. Click **Apply**.

Viewing current BSR information

View the current BSR information to review the configuration.

Before you begin

- You must disable static RP.

Procedure

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

2. Click **PIM**.
3. Click the **Current BSR** tab.

Current BSR field descriptions

Use the descriptions in the following table to use the **Current BSR** tab.

Name	Description
Address	Shows the IP address of the current BSR for the local PIM domain.
FragmentTag	Shows a randomly generated number that distinguishes fragments that belong to different bootstrap messages. Fragments that belong to the same bootstrap message carry the same fragment tag.
HashMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. The hashmask allows a small number of consecutive groups to always hash to the same RP.
Priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
BootStrapTimer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

Changing VLAN interface type

Change the state (active or passive) of PIM on a VLAN interface.

Before you begin

- Before you change the state of PIM on a VLAN interface, you must first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when the interface receives streams.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select the VLAN ID that you want to configure with PIM.
5. Click **IP**.
6. Click the **PIM** tab.
7. Clear the **Enable** check box.
8. Click **Apply**.

9. Select **active** or **passive**.
10. Reenable PIM on the VLAN interface.
11. Click **Apply**.

Editing PIM interface parameters

Edit PIM parameters for an interface to customize the PIM configuration.

Before you begin

- Before you change the state (active or passive) of a PIM interface, first disable PIM to prevent instability in the PIM operations, especially when neighbors exist or when the interface receives streams.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Interfaces** tab.
4. Edit the fields by double-clicking on them, and then select or type the new value.
5. Click **Apply**.

Interfaces field descriptions

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

Name	Description
IfIndex	Shows the interface Index. This variable is a read-only field.
Address	Shows the IP address of the PIM interface. This variable is a read-only field.
NetMask	Shows the network mask for the IP address of the PIM interface. This variable is a read-only field.
Mode	Shows the configured mode of this interface. The valid modes are SSM and sparse. This variable is a read-only field.
InterfaceType	Specifies if the interface is active or passive.
DR	Shows the router with the highest IP address on a LAN designated to perform these tasks.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds.

Table continues…

Name	Description
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds.
CBSRPreference	Configures the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OperState	Indicates the status of PIM on this interface: Up or Down.

Editing IPv6 PIM interface parameters

Edit the IPv6 PIM parameters for an interface to customize the IPv6 PIM configuration.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Interfaces** tab.
4. Edit the fields by double-clicking on them, and then select or type the new value.
5. Click **Apply**.

Interfaces field descriptions

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

Name	Description
IfIndex	Shows the interface Index. This variable is a read-only field.
Address	Shows the IPv6 address of the PIM interface. This variable is a read-only field.
NetMask	Shows the network mask for the IPv6 address of the PIM interface. This variable is a read-only field.
Enable	Shows the configured mode of this PIM interface. sparseDense mode is valid only for PIMv1.
Mode	Shows the configured mode of this interface. The valid modes are SSM and sparse. This variable is a read-only field.
DR	Shows the router with the highest IPv6 address on a LAN designated to perform these tasks.
HelloInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default is 30 seconds.

Table continues…

Name	Description
JoinPruneInterval	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default is 60 seconds.
OperState	Indicates the status of PIM on this interface: Up or Down.
Type	Specifies the interface type.

Configuring the PIM virtual neighbor

Configure a PIM virtual neighbor if the next hop for a static route cannot run PIM, such as the Virtual Router Redundancy Protocol (VRRP) address on an adjacent device.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Virtual Neighbors** tab.
4. Click **Insert**.
5. Specify the IP address of the virtual neighbor.
6. Specify the interface index for the PIM interface.
7. Click **Insert**.

Virtual Neighbors field descriptions

Use the descriptions in the following table to use the **Virtual Neighbors** tab.

Name	Description
Address	Specifies the IP address of the neighbor.
IfIndex	Specifies the IP address of the PIM interface.

Viewing PIM-SM neighbor parameters

View PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.

3. Click the **Neighbors** tab.

Neighbors field descriptions

Use the descriptions in the following table to use the **Neighbors** tab.

Name	Description
Address	Shows the IP address of the PIM neighbor.
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Shows the time since this neighbor became a neighbor of the local router.
ExpiryTime	Shows the time remaining before the neighbor expires.

Viewing IPv6 PIM-SM neighbor parameters

View IPv6 PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Neighbors** tab.

Neighbors field descriptions

Use the descriptions in the following table to use the **Neighbors** tab.

Name	Description
Address	Shows the IPv6 address of the PIM neighbor.
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Shows the time since this neighbor became a neighbor of the local router.
ExpiryTime	Shows the time remaining before the neighbor expires.

Viewing IPv6 Neighbor Secondary Address

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **Neighbor Secondary Address** tab.

Neighbor Secondary Address field descriptions

Use the descriptions in the following table to use the **Neighbor Secondary Address** tab.

Name	Description
IfIndex	Shows the slot and port number or VLAN ID of the interface used to reach this PIM neighbor.
Type	Shows the address type of this PIM neighbor.
Primary	The primary IPv6 address of this PIM neighbor.
SecAddress	The secondary IPv6 address of this PIM neighbor.

Viewing RP set parameters

View the RP set to see a list of rendezvous point addresses. The BSR constructs this list from C-RP advertisements, and then distributes it to all PIM routers in the PIM domain for the BSR. View the parameters for troubleshooting purposes.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **RP Set** tab.

RP Set field descriptions

Use the descriptions in the following table to use the **RP Set** tab.

Name	Description
GroupAddress	Shows the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.

Table continues…

Name	Description
GroupMask	Shows the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
Address	Shows the IP address of the C-RP router.
HoldTime	Shows the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
ExpiryTime	Shows the time remaining before this C-RP router times out.

Configuring a candidate RP

Configure a C-RP router to add it to the RP Set.

About this task

You can configure only one interface on a switch for multiple groups; that is, you cannot configure multiple interfaces for multiple groups.

Using the GroupMask value, you can configure a candidate RP for several groups in one configuration. For example, if you use a C-RP configuration with a GroupAddress value of 224.0.0.0 and a GroupMask of 240.0.0.0, you can configure the C-RP router for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **PIM**.
3. Click the **Candidate RP** tab.
4. Click **Insert**.
5. Type the required information in each box.
6. Click **Insert**.

Candidate RP field descriptions

Use the descriptions in the following table to use the **Candidate RP** tab.

Name	Description
GroupAddress	Configures the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.

Table continues…

Name	Description
GroupMask	Configures the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
InterfaceAddress	Configures the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

Enabling square-SMLT globally

Use square-Split MultiLink Trunking (SMLT) to form an SMLT aggregation group. In a square configuration, enable square-SMLT globally on each of the four switches.

About this task

Important:

The following configuration also activates full-mesh configurations.

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**
2. Click **Multicast**.
3. Click the **Globals** tab.
4. Select **MulticastSquareSmltEnable**.

   Clear this check box if you want to disable square-SMLT globally.
5. Click **Apply**.

Viewing IPv6 RP set parameters

View the IPv6 RP set to see a list of rendezvous point addresses. View the parameters for troubleshooting purposes.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 PIM**.
3. Click the **RP Set** tab.

RP Set field descriptions

Use the descriptions in the following table to use the **RP Set** tab.

Name	Description
GroupAddress	Specifies the IPv6 address of the multicast group. When combined with the group mask, this value identifies a group prefix for which the address is a static RP.
GroupMask	Specifies the address mask of the multicast group. When combined with the group address, this value identifies a group prefix for which the address is a static RP.
Address	Specifies the IPv6 address of the static RP.
HoldTime	Specifies the hold time of the static RP. The value is 0.
ExpiryTime	Specifies the minimum time remaining before the static RP is down. The value is 0.

Viewing IPv6 Mroute interface information

Use the following procedure to view IPv6 Mroute information for an interface.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 Mroute**.
3. Click the **Interfaces** tab.

Interfaces field descriptions

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

Name	Description
IfIndex	Displays the slot and port number or VLAN ID for this entry.
Ttl	Displays the datagram time-to-live (TTL) threshold for the interface. IPv6 multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means that all multicast packets are forwarded out of the interface.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured

Table continues…

Name	Description
	• netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

Viewing IPv6 Mroute next hop information

Use the following procedure to view IPv6 Mroute next hop information.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 Mroute**.
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
Group	Displays the IPv6 multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
IfIndex	Displays the slot and port number or VLAN ID for this entry.
Address	Displays the address of the next hop specific to this entry. For most interfaces, it is identical to the next-hop group. Non Broadcast Multiple Access (NBMA) interfaces, however, can use multiple next hop addresses out of a single outgoing interface.
State	Displays whether the outgoing interface and next hop represented by this entry currently forward IPv6 datagrams. A value of forwarding indicates the information is currently used; pruned indicates it is not used.

Table continues…

Name	Description
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	Displays the minimum number of hops between this router and members of the IPv6 multicast group reached through the next hop on this outgoing interface. IPv6 multicast datagrams for the group that use a time-to-live less than this number of hops are not forwarded to the next hop.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

Configuring resource usage counter for IPv6 Mroute

Configure the resource usage counters to query the number of ingress and egress IPv6 multicast streams traversing the switch. After you configure the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive notification through a trap on the console, a logged message, or both.

Important:

If you do not configure the thresholds, EDM displays only the ingress and egress records that are currently in use.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 Mroute**.
3. Click the **Resource Usage** tab.
4. Configure the ingress and egress thresholds.
5. Configure the notification methods.
6. Click **Apply**.

Resource Usage field descriptions

Use the data in the following table to use the **Resource Usage** tab.

Name	Description
Ingress Records In-Use	Displays the number of ingress records (source or group) traversing the switch.
Egress Records In-Use	Displays the number of egress records traversing the switch.
Ingress Threshold	Configures the ingress threshold level (0–32767).
Egress Threshold	Configures the egress threshold level (0–32767).
SendTrapAndLog	Sends both trap and log notification messages after the number of streams exceeds a threshold level.
SendTrapOnly	Sends only trap notification messages after the number of streams exceeds a threshold level. You can configure only one notification type.
LogMsgOnly	Sends only log notification messages after the number of streams exceeds a threshold level.

Viewing IPv6 multicast route information

Use the following procedure to view IPv6 Mroute route information.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 Mroute**.
3. Click the **Route** tab.

IPv6 Multicast Route field descriptions

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

Name	Description
Group	Displays the IPv6 multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.

Table continues…

Name	Description
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
UpStreamNeighbor	Shows the address of the upstream neighbor from which the IPv6 datagrams from these sources are received.
IfIndex	Displays the slot and port number or VLAN ID for this entry.
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

Chapter 9: IGMP configuration using the CLI

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.

Important:

The **spbm-config-mode** boot flag must be disabled before you can configure PIM or IGMP. To verify the setting, enter **show boot config flags** in Privileged EXEC mode.

Before you begin

- Complete one of the following tasks:
  - Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
  - Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Protocol Independent Multicast-Sparse Mode (PIM-SM) or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

  Important:

  To configure and use IGMP on a VRF instance you must first select and launch the VRF context.

  To select and launch the VRF context, see Configuring IGMP on a VRF on page 111.

Configuring multicast stream limitation on an Ethernet port

Configure multicast stream limitation on an Ethernet port to limit the number of concurrent multicast streams on the port. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the port drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Note:

Configuration of multicast stream limitation is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

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 multicast stream limitation and configure the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

3. If stream-limit is already enabled on the interface, change the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

4. Display multicast stream limitation information for the ports on a specific interface:

   ```
   show ip igmp stream-limit interface
   ```

Example

Enable multicast stream limitation on the Ethernet port and configure the maximum number of allowed streams to 8.

```
Switch:1(config-if)# ip igmp stream-limit
Switch:1(config-if)# ip igmp stream-limit stream-limit-max-streams 8
```

Variable definitions

Use the data in the following table to use the **`ip igmp stream-limit-max-streams`** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this port. The range is from 0–65535 and the default is 4.

Job aid

The following tables show the field descriptions for the **show ip igmp stream-limit interface** command.

Table 6: show ip igmp stream-limit interface field descriptions

Field	Description
INTERFACE	Indicates the interface IP address.
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

Configuring multicast stream limitation on a VLAN

Configure multicast stream limitation on a VLAN to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, providers can protect the bandwidth on a specific interface and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the VLAN drops joins to new streams. A service provider uses this feature to control the overall bandwidth usage in addition to restricting users from attaching more than the allowed television sets to a link.

Note:

Configuration of multicast stream limitation is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable multicast stream limitation and configure the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

3. If stream-limit is already enabled on the VLAN, change the maximum number of allowed streams:

   ```
   ip igmp stream-limit stream-limit-max-streams <0-65535>
   ```

4. Display multicast stream limitation information for the ports on a specific interface:

   ```
   show ip igmp stream-limit port
   ```

Example

Enable multicast stream limitation and configure the maximum number of allowed streams to 8.

```
Switch:1(config-if)# ip igmp stream-limit
Switch:1(config-if)# ip igmp stream-limit stream-limit-max-streams 8
```

Variable definitions

Use the data in the following table to use the **`ip igmp stream-limit`** command.

Variable	Value
<0-65535>	Configures the maximum number of allowed streams on this VLAN. The range is from 0–65535 and the default is 4.

Job aid

The following tables show the field descriptions for the **`show ip igmp stream-limit port`** command.

Table 7: show ip igmp stream-limit port field descriptions

Field	Description
INTERFACE	Indicates the interface IP address.
PORT	Indicates the port for the VLAN.
MAX STREAMS	Indicates the maximum number of streams.
NUM STREAMS	Indicates the current number of streams.

Configuring VLAN multicast stream limitation members

Configure multicast stream limitation members on ports of a specific VLAN to limit the number of multicast groups that can join a VLAN.

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable
   configure terminal
   ```

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

2. Configure multicast stream limitation members on a VLAN:

```
ip igmp stream-limit-group {slot/port[/sub-port][-slot/port[/sub-port]][,...]} enable max-streams <0-65535>
```

Example

Enable multicast stream limitation on ports 2/3 to 2/8 and configure the maximum allowed number of streams to 6 for this interface.

```
Switch:1(config-if)# ip igmp stream-limit-group 2/3-2/8 max-streams 6
```

Variable definitions

Use the data in the following table to use the **`ip igmp stream-limit-group`** command.

Variable	Value
<0–65535>	Configures the maximum number of allowed streams for the specified ports on this VLAN. The range is from 0–65535 and the default is 4.
{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 multicast router discovery options

Configure the multicast router discovery options to enable the automatic discovery of multicast-capable routers.

About this task

Important:

The switch does not support the Multicast Router Discovery (MRDISC) protocol on brouter ports.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable multicast router discovery:

```
ip igmp mrdisc
```

3. Configure the maximum advertisement intervals between successive advertisements:

   ```
   ip igmp mrdisc maxadvertinterval <2–180> maxinitadvertinterval <2–180>
   ```

4. Configure the maximum advertisements after initialization:

   ```
   ip igmp mrdisc maxinitadvertisements <2–15>
   ```

5. Configure the minimum advertisement interval between successive advertisements:

   ```
   ip igmp mrdisc minadvertinterval <3–180>
   ```

6. Configure the time allowed before a neighbor is declared dead:

   ```
   ip igmp mrdisc neighdeadinterval <2–180>
   ```

Example

Configure the maximum advertisement intervals between successive advertisements:

```
Switch:1(config-if)#ip igmp mrdisc maxadvertinterval 30 maxinitadvertinterval 5
```

Configure the maximum advertisements after initialization:

```
Switch:1(config-if)#ip igmp mrdisc maxinitadvertisements 8
```

Configure the minimum advertisement interval between successive advertisements:

```
Switch:1(config-if)#ip igmp mrdisc minadvertinterval 30
```

Configure the time allowed before a neighbor is declared dead:

```
Switch:1(config-if)#ip igmp mrdisc neighdeadinterval 60
```

Variable definitions

Use the data in the following table to use the **`ip igmp mrdisc`** command.

Variable	Value
maxadvertinterval <*2–180*>	Configures the maximum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 20.
maxinitadvertinterval <2–180>	Configures the maximum number (in seconds) between successive initial advertisements. For this change to take effect, you must save the configuration, and then reset the switch.

Table continues…

Variable	Value
	To configure this option to the default value, use the default operator with the command. The default is 2.
maxinitadvertisements <2–15>	Configures the maximum number of initial multicast advertisements after initialization. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 3.
minadvertinterval *<3–180>*	Configures the minimum number (in seconds) between successive advertisements. For this change to take effect, you must save the configuration, and then reset the switch. To configure this option to the default value, use the default operator with the command. The default is 15.
neighdeadinterval *<2–180>*	Configures the multicast router discovery dead interval—the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down. To configure this option to the default value, use the default operator with the command. The default is 60.

Configuring explicit host tracking

Configure explicit host tracking to track all the source and group members.

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. Configure explicit host tracking:

   ```
   ip igmp igmpv3-explicit-host-tracking
   ```

3. Display all the tracked members for a specific group:

```
show ip igmp group group <A.B.C.D> tracked-members [member-subnet
<A.B.C.D/X>] [source-subnet <A.B.C.D/X>] [port {slot/port[/sub-port]
[-slot/port[/sub-port]][,...]}] [vlan <1-4059>]
```

4. Display the IGMPv3 specific data:

```
show ip igmp group group <A.B.C.D> detail port {{slot/port[/sub-
port][-slot/port[/sub-port]][,...]}} vlan <1-4059>
```

Example

Configure explicit host tracking:

```
Switch:1(config-if)#ip igmp igmpv3-explicit-host-tracking
```

Display all the tracked members:

```
Switch:1(config-if)#show ip igmp group
================================================================================
                                Igmp Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
   225.1.1.1       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.2       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.3       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.4       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.5       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.6       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.7       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.8       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.9       V22-1/1         22.22.22.200    178        Dynamic
   225.1.1.10      V22-1/1         22.22.22.200    178        Dynamic
   225.12.12.1     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.2     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.3     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.4     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.5     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.6     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.7     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.8     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.9     V2222-2/16      22.2.2.200      172        Dynamic
   225.12.12.10    V2222-2/16      22.2.2.200      172        Dynamic
   226.1.1.1       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.2       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.3       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.4       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.5       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.6       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.7       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.8       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.9       V33-1/23        33.33.33.200    173        Dynamic
   226.1.1.10      V33-1/23        33.33.33.200    173        Dynamic
   226.22.22.1     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.2     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.3     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.4     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.5     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.6     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.7     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.8     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.9     V3333-2/22      33.3.3.200      173        Dynamic
   226.22.22.10    V3333-2/22      33.3.3.200      173        Dynamic
   228.45.45.45    V222-1/1        122.122.122.200 173        Dynamic
```

```
   228.56.56.56    V222-1/1        122.122.122.200 166        Dynamic
   229.1.1.1       V333-1/17       133.133.133.200 172        Dynamic
   229.32.32.32    V222-1/1        122.122.122.200 169        Dynamic
   232.1.1.1       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.2       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.3       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.4       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.5       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.6       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.7       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.8       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.9       V333-1/17       133.133.133.200 170        Dynamic
   232.1.1.10      V333-1/17       133.133.133.200 170        Dynamic
   232.32.32.1     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.2     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.3     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.4     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.5     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.6     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.7     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.8     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.9     V222-1/1        122.122.122.200 165        Dynamic
   232.32.32.10    V222-1/1        122.122.122.200 162        Dynamic
   232.42.42.1     V222-1/1        122.122.122.200 167        Dynamic

65 out of 65 group Receivers displayed

Total number of unique groups 65
```

Display all the tracked members for a specific group:

```
Switch:1(config-if)#show ip igmp group group 232.1.1.1 tracked-members
================================================================================
                          Members of Channels/Groups - GlobalRouter
================================================================================
INTERFACE       CHANNEL/GROUP         MEMBER           MEMBER_MODE EXP
--------------------------------------------------------------------------------
Vlan333-2/30    */232.1.1.1          133.133.133.200   IS_EXCLUDE  205

Note:
The "*" attached to the interface (if any) indicates that the interface has explicit host
tracking disabled.
```

Display IGMPv3 specific data:

```
Switch:1(config-if)#show ip igmp group group 232.32.32.10 detail
================================================================================
                       Igmp Group Detail - GlobalRouter
================================================================================
Interface:                    Vlan222-1/10
IGMPv3 Group:                 232.32.32.10
Interface Group Mode:         INCLUDE
Interface Compatibility Mode: IGMP_V3
V2 Host Timer:                Not Running
V1 Host Timer:                Not Running
Interface Group Include Source List:
        Source Address  Expires
        133.133.133.200 114
```

Variable definitions

Use the data in the following table to use the **ip igmp igmpv3-explicit-host-tracking** command.

Variable	Value
explicit-host-tracking	Enables explicit host tracking on IGMPv3. The default state is disable.
<A.B.C.D>	Specifies the IP address of the group of the tracked member.

Configuring IGMP static members

Configure IGMP static members to add members to a snoop group. You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams are always forwarded to the multicast router within the VLAN, in addition to the ports in this static entry.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Configure interface static members:

   ```
   ip igmp static-group {A.B.C.D} {A.B.C.D} {port {slot/port[/sub-port]
   [-slot/port[/sub-port]][,...]} [static|blocked]
   ```

Example

Configure interface static members:

```
Switch:1(config-if)#ip igmp static-group 239.1.1.1 239.1.2.1 port 2/1 static
```

Variable definitions

Use the data in the following table to use the **ip igmp static-group** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Indicates the IP address range of the selected multicast group.

Table continues…

Variable	Value
port	Adds ports to a static group entry
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Creates a static group entry. Specifies the port or list of ports that is a member of the VLAN interface being configured to which you want to redirect the multicast stream for this multicast group. 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.
<static|blocked>	Configures the route to static or blocked.

Configuring SSM dynamic learning and range group

Configure SSM dynamic learning and a range group to enable the IGMPv3 dynamic learning feature and to extend the default SSM range of 232/8 to include an IP multicast address. As new SSM channels are learned, they appear in the SSM channel table.

Before you begin

- To define the range group, you must first disable PIM.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable SSM dynamic learning:

   ```
   ip igmp ssm dynamic-learning
   ```

3. Configure the range group:

   ```
   ip igmp ssm group-range <A.B.C.D/X>
   ```

   The following message appears:.

   ```
   Warning: Changing the SSM range will cause all spb-multicast and
   spb-pim-gw enabled interfaces to be internally bounced. Do you wish
   to continue? (y/n) ? (y/n)?
   ```

Enter `y` to continue.

Example

Define the SSM range group address (234.0.0.0) and mask (255.0.0.0). Enable dynamic learning from IGMPv3 reports.

```
Switch:1(config)#ip igmp ssm group-range 234.0.0.0/255.0.0.0

WARNING: Changing the SSM range will cause all spb-multicast and spb-pim-gw enabled
interfaces to be internally bounced. Do you wish to continue? (y/n) ? (y/n)? y
Switch:1(config)#ip igmp ssm dynamic-learning
```

Variable definitions

Use the data in the following table to use the **`ip igmp ssm`** command.

Variable	Value
{A.B.C.D/X}	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Changing the SSM range group

Change the SSM range group to define the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address.

About this task

Important:

This procedure reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), it also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Disable PIM:

   ```
   no ip pim enable
   ```

   If you forget to disable PIM, the following error message appears:

   ```
   Error: PIM is enabled in SSM mode, disable PIM
   ```

3. Delete each entry in the SSM channel table:

   ```
   no ip igmp ssm-map [all] [{A.B.C.D} enable]
   ```

   If you forget to delete the SSM channels, the following error message appears:

   ```
   Error: SSM source group table not empty
   ```

4. Configure the new IP multicast group address:

   ```
   ip igmp ssm group-range {A.B.C.D/X}
   ```

   The following message appears:.

   ```
   Warning: Changing the SSM range will cause all spb-multicast and
   spb-pim-gw enabled interfaces to be internally bounced. Do you wish
   to continue? (y/n) ? (y/n)?
   ```

   Enter `y` to continue.

5. Enable PIM:

   ```
   ip pim enable
   ```

Example

Configure the new IP multicast group address:

```
Switch:1(config)#ip igmp ssm group-range 232.0.0.0/16

WARNING: Changing the SSM range will cause all spb-multicast and spb-pim-gw enabled
interfaces to be internally bounced. Do you wish to continue? (y/n) ? (y/n)? y
```

Variable definitions

Use the data in the following table to use the **`ip igmp ssm group-range`** and **`ip igmp ssm`** commands.

Variable	Value
{A.B.C.D/X}	Defines the SSM range. The SSM range parameter extends the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without having to change their group configurations. This parameter specifies an IP multicast address within the range of 224.0.0.0 and 239.255.255.255. The default is 232.0.0.0. The address mask is the IP address mask of the multicast group. The default is 255.0.0.0.

Configuring the SSM map table

Configure the SSM map table to map groups to their sending source. SSM maps cannot conflict with static source groups. After you configure an SSM map or a static source group, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) to different sources or multiple sources to the same group for both static source group and an SSM map.

About this task

The consistency check applies to all SSM map entries, even if they are disabled. If you disable an entry, it becomes inactive. If you do not delete the entry, you can reenable it later.

After you disable an SSM map, the switch stops multicast traffic from the specified source to the specified group. You can use this static configuration as a security feature to block traffic from a certain source to a specific group.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable the SSM map table for all static entries:

   ```
   ip igmp ssm-map all
   ```

3. Create a static entry for a specific group:

   ```
   ip igmp ssm-map {A.B.C.D} {A.B.C.D} enable
   ```

Example

Create an SSM map table entry for the multicast group 234.0.1.0 and the source at 192.32.99.151. Configure the administrative state to enable all the static SSM map table entries.

```
Switch:1(config)#ip igmp ssm-map 234.0.1.0 192.32.99.151
Switch:1(config)#ip igmp ssm-map all
```

Variable definitions

Use the data in the following table to use the **`ip igmp ssm-map`** command.

Variable	Value
{A.B.C.D} {A.B.C.D}	Creates a static SSM channel table entry by specifying the group and source IP addresses. The IP address is an IP multicast address within the SSM range. The source IP address is an IP host address that sends traffic to the group.

Table continues…

Variable	Value
{A.B.C.D} enable	Enables the administrative state for a specific entry (group). This variable does not affect the dynamically learned entries. This state determines whether the switch uses the static entry or saves it for future use. The default is enable for each entry.

Configuring multicast access control for an IGMP Ethernet port

Configure multicast access control for an IGMP Ethernet port to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

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 multicast access control:

   ```
   ip igmp access-list WORD<1-64> {A.B.C.D/X} <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

3. Change an existing access list:

   ```
   ip igmp access-list WORD<1-64>> {A.B.C.D/X} mode <deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp access-list`** command

Variable	Value
{A.B.C.D/X}	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this

Table continues…

Variable	Value
	configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
mode	Changes the access control group configuration.
WORD<1–64>	Specifies the name of the access list from 1–64 characters.

Configuring multicast access control for a VLAN

Configure multicast access control for an IGMP VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

Procedure

1. Enter VLAN Interface Configuration mode:

   ```
   enable

   configure terminal

   interface vlan <1-4059>
   ```

2. Configure multicast access control:

   ```
   ip igmp access-list WORD<1-64> [A.B.C.D/X] <deny-tx|deny-rx|deny-
   both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

3. Change an existing access list:

   ```
   ip igmp access-list WORD<1-64> [A.B.C.D/X] mode <deny-tx|deny-rx|
   deny-both|allow-only-tx|allow-only-rx|allow-only-both>
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp access-list`** command.

Variable	Value
{A.B.C.D/X}	Creates an access control group entry for a specific IGMP interface. Specify the IP address of the host and the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.

Table continues…

Variable	Value
deny-tx|deny-rx|deny-both|allow-only-tx|allow-only-rx|allow-only-both	Indicates the action for the specified IGMP interface. For example, if you specify deny-both, the interface denies both transmitted and received traffic
mode	Changes the access control group configuration.
WORD<1–64>	Specifies the name of the access list from 1–64 characters.

Configuring fast leave mode

Configure fast (immediate) leave mode to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces. Normal IGMP behavior is skipped. Fast leave mode provides one command that controls all IGMP fast leave enabled interfaces.

Before you begin

- You must enable explicit-host-tracking before configuring fast-leave mode for IGMPv3. For more information on enabling explicit-host-tracking, see Configuring explicit host tracking on page 198.

About this task

If a single user connects to an interface, you do not need to track if other users exist on the interface to perform the fast leave. In cases like this, you must change the mode to one-user.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. View the current fast leave mode:

   ```
   show ip igmp sys
   ```

   Note:

   This command is not supported on a node configured as the DvR Leaf within a DvR domain.

   For information on DvR, see Configuring IPv4 Routing for VOSS.

3. Configure fast leave mode:

   ```
   ip igmp immediate-leave-mode <multiple-user|one-user>
   ```

Example

Change the mode to one-user.

```
Switch:1(config)#ip igmp immediate-leave-mode one-user
```

Variable definitions

Use the data in the following table to use the **`ip igmp immediate-leave-mode`** command.

Variable	Value
multiple-user|one-user	multiple-user removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This configuration is the default. one-user removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.

Enabling fast leave mode on a port

Enable fast (immediate) leave mode to specify if a port receives a leave message from a member of a group. If you enable fast leave mode on a port, it uses the global fast leave mode configuration.

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 fast leave:

   ```
   ip igmp immediate-leave
   ```

Configuring IGMP fast leave members on a VLAN

Configure IGMP fast leave members on a VLAN to specify fast leave capable ports.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable fast leave on the VLAN:

   ```
   ip igmp immediate-leave
   ```

3. Configure fast leave members on a VLAN:

   ```
   ip igmp immediate-leave-members {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```

Variable definitions

Use the data in the following table to use the **`ip igmp immediate-leave-members`** 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.

Enabling IGMP Layer 2 Querier

When no multicast router exists in your network, you can use IGMP Layer 2 Querier to allow the Layer 2 switch to act as a multicast router so that the system can participate in multicast environments where multicast routing is not required.

Before you begin

- You must enable IGMP snooping.

About this task

When you enable IGMP Layer 2 Querier, Layer 2 switches in your network can snoop IGMP control packets exchanged with downstream hosts and upstream routers. The Layer 2 switches then generate the Layer 2 MAC forwarding table, used for switching sessions and multicast traffic regulation, and provide the recurring queries required to maintain IGMP groups.

By default, IGMP Layer 2 Querier is disabled.

Enable Layer 2 Querier on only one node in the VLAN.

On Shortest Path Bridging (SPB) Customer VLANs (CVLAN), IGMP Querier is enabled automatically when you enable snooping on the VLAN. For more information about SPB, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable IGMP Layer 2 Querier:

   ```
   ip igmp snoop-querier
   ```

Next steps

You must enable the IGMP Layer 2 Querier address. See EnablingIGMPLayer2QuerierAddress on page 211

Enabling IGMP Layer 2 Querier address

To use the IGMP Layer 2 Querier feature you must designate the IGMP Layer 2 Querier source IP address, the address the system uses in the query message.

Before you begin

- Enable IGMP Layer 2 Querier.

About this task

You must configure the IGMP Layer 2 Querier address to an IP address in the IP subnet that IGMP hosts, and to which IGMP snoopers in the VLAN belong.

The default IP address is 0.0.0.0 when the IGMP Layer 2 Querier is disabled.

Procedure

1. Enter VLAN Interface Configuration mode:

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

2. Enable the IGMP Layer 2 Querier address:

   ```
   ip igmp snoop-querier-addr {A.B.C.D}
   ```

3. Verify the configuration:

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

Example

Enable the IGMP Layer 2 Querier feature for VLAN 4, and configure the querier address. Verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#interface vlan 4
Switch:1(config-if)#ip igmp snoop-querier
Switch:1(config-if)#ip igmp snoop-querier-addr 192.0.2.1
Switch:1(config-if)#show ip igmp snooping

================================================================================
                          Igmp Snooping - GlobalRouter
================================================================================
IFINDEX SNOOP   PROXY   SSM     STATIC           ACTIVE           MROUTER
        ENABLE  SNOOP   SNOOP   MROUTER          MROUTER          EXPIRATION
                ENABLE  ENABLE  PORTS            PORTS            TIME
--------------------------------------------------------------------------------
V2      false   false   false                                     0
V3      false   false   false                                     0
V4      true    false   false                                     0
V200    false   false   false                                     0

IFINDEX SNOOP   SNOOP           DYNAMIC   COMPATIBILITY
        QUERIER QUERIER         DOWNGRADE MODE
        ENABLE  ADDRESS         VERSION
--------------------------------------------------------------------------------
V2      false   0.0.0.0         enable    disable
V3      false   0.0.0.0         enable    disable
V4      true    192.0.2.1       enable    disable
V200    false   0.0.0.0         enable    disable

4 out of 4 entries displayed
```

Chapter 10: IGMP configuration using EDM

Hosts use the Internet Group Management Protocol (IGMP) to report their IP multicast group memberships to neighboring multicast routers. Configure IGMP on an individual interface basis.

Important:

The **EnableSpbmConfigMode** boot flag must be disabled before you can configure PIM or IGMP. To verify the setting, navigate to **Configuration** > **Edit** > **Chassis** and click on the **Boot Config** tab.

Before you begin

- Configure IGMP on a Layer 2 interface by enabling IGMP snoop.
- Configure IGMP on a Layer 3 interface by enabling multicast routing, for example, Protocol Independent Multicast-Sparse Mode (PIM-SM), or Protocol Independent Multicast-Source Specific Multicast (PIM-SSM).

Important:

To configure and use IGMP on a VRF instance you must first select and launch the VRF context.

To select and launch the VRF context, see Selecting and launching a VRF context view on page 116.

Enabling IGMP snoop on a VLAN

Enable IGMP snooping on a VLAN to optimize the multicast data flow for a group within a VLAN to only those that are members of the group that uses IGMP snoop.

About this task

The switch listens to group reports from each port and builds a database of multicast group members for each port. The switch suppresses the reports heard by not forwarding them to other hosts, forcing the members to continuously send their own reports.

The switch relays group membership from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN. The switch multicasts data only to the participating group members and to the multicast routers within the VLAN.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **VLAN**.

2. Click **VLANs**.
3. Click the **Basic** tab.
4. Select a VLAN.
5. Click **IP**.
6. Click the **IGMP** tab.
7. Select the **SnoopEnable** check box.
8. Select the **ProxySnoopEnable** check box.
9. For SteamLimtEnable, select **enable**.
10. Click **Apply**.

Configuring IGMP interface static members

Configure IGMP interface static members to add members to a snoop group.

About this task

You can create a static entry to forward multicast data streams to a particular set of ports within the VLAN. After you create the entry, multicast data streams always forward to the multicast router within the VLAN, in addition to the ports in this static entry.

Important:

IGMP snoop can optimize only local multicast data flow. IGMP snoop does not manage the forwarding state of the multicast tree. You cannot configure a port as a static receiver in an IGMP snoop-enabled VLAN that does not contain at least one dynamic receiver port and forward multicast data.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Static** tab.
4. Click **Insert**.
5. Type the appropriate information.
6. Click **Insert**.

Static field descriptions

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

Name	Description
IfIndex	Shows the interface where the IGMP entry is enabled.
GrpAddr	Indicates the start of the IP multicast address range of the multicast stream. Within the indicated valid range (224.0.0.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.
ToGrpAddr	Indicates the end of the IP multicast address range of the multicast stream. If an address is not entered, the IP address in the GrpAddr field is the single address.
MemberPorts	Specifies the ports to which you want to redirect the multicast stream for this multicast group. The ports must be member ports of the VLAN.
NotAllowedToJoin	Specifies the ports that do not receive the multicast stream for this multicast group.

Configuring the SSM map table

Configure the SSM map table to map groups to their sending source. SSM maps cannot conflict with static source groups. After you configure an SSM map or a static source group, the switch performs a consistency check to make sure no conflicts exist. You can map one group (G) or multiple groups to different sources for both static source group and an SSM channel.

About this task

The consistency check applies to all SSM channel entries, even if they are disabled. If you disable an entry, it becomes inactive. If you do not delete the entry, you can reenable it later.

After you disable an SSM map, the switch stops multicast traffic from the specified source to the specified group. You can use this static configuration as a security feature to block traffic from a certain source to a specific group.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Ssm Map** tab.
4. Click **Insert**.
5. Type the IP address for the multicast group and source.

6. Click **Insert**.

   You can change the default status of an SSM map from enable to disable by clicking in the AdminState field.

Ssm Map field descriptions

Use the data in the following table to use the **Ssm Map** tab.

Name	Description
IpMulticastGrp	Specifies an IP multicast address that is within the SSM range.
IpSource	Specifies the IP address of the source that sends traffic to the group.
LearningMode	Displays whether the entry is statically configured (Static) or dynamically-learned from IGMPv3 (Dynamic). This variable a read-only field.
Activity	Displays the current activity of the selected (S,G) entry. True indicates that traffic is flowing to the switch, otherwise, it appears false. This variable a read-only field.
AdminState	Configures the administrative state for the selected static entry. This state determines whether the switch uses the static entries. Configure this field to enable (default) to use the entry or disable to save for future use.

Configuring SSM range and global parameters

Configure the SSM range parameter to extend the default SSM range of 232/8 to include an IP multicast address. You can configure existing applications without changing their group configurations.

Before you begin

- To change the RangeGroup configuration, you must first disable PIM.
- To change the RangeGroup configuration, you must delete all entries in the SSM channel table before you configure the new IP multicast group address.

About this task

The other global parameters enable the IGMPv3 dynamic learning feature and configure the administrative state for all the entries in the SSM channel table.

Important:

If you change the RangeGroup configuration, the switch reinitializes PIM and temporarily stops all PIM traffic. For those multicast groups out of SSM range (for example, under PIM-SM behavior), this procedure also causes a rendezvous point (RP) relearn delay of up to 60 seconds. This delay can be longer if the bootstrap router (BSR) is local.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Ssm Global** tab.
4. Configure the appropriate fields.
5. Click **Apply**.

Ssm Global field descriptions

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

Name	Description
DynamicLearning	Activates the dynamic learning of SSM channel (S,G) pairs from IGMPv3 reports. As new SSM channels are learned, they appear in the SSM channel table.
RangeGroup	Configures the IP multicast group address. The lowest group address is 224.0.0.0 and the highest is 239.255.255.255. The default is 232.0.0.0.
RangeMask	Configures the address mask of the multicast group. The default is 255.0.0.0.
SsmMapAdminAction	Configures the administrative state, which determines whether the switch uses the table entries: • enableAll—Globally activates all the static entries in the SSM channel table. This value does not affect the dynamically learned entries. • disableAll—Globally inactivates all the static entries in the SSM channel table. This value does not affect the dynamically learned entries.

Configuring multicast stream limitation on an interface

Configure multicast stream limitation to limit the number of concurrent multicast streams on the interface. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific interface and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the interface drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a specific limit of multicast streams on an interface.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit** tab.
4. To change the status of an interface, double-click on the **StreamLimitEnable** field for the interface, and then select **enable** or **disable** from the menu. If the interface is enabled, you can edit the **Maximum Number of Stream** field.
5. Click **Apply**.

StreamLimit field descriptions

Use the data in the following tab to use the **StreamLimit** tab.

Name	Description
Interface	Displays the slot and port number or VLAN ID for this interface.
StreamLimitEnable	Enables or disables stream limitation on this interface.
Maximum Number Of Stream	Configures the maximum number of streams allowed on this interface. The range is from 0–65535, and the default is 4.
Current Number Of Stream	Displays the current number of streams received on this interface. This value is a read-only value.

Configuring multicast stream limitation on a VLAN

Configure multicast stream limitation to limit the number of concurrent multicast streams on the VLAN. By limiting the number of concurrent multicast streams, you can protect the bandwidth on a specific VLAN and control access to multicast streams.

About this task

You can configure the maximum number of streams independently. After the number of streams reaches the limit, the VLAN drops additional join reports for new streams. You can control the overall bandwidth usage in addition to restricting users from receiving more than a specific limit of multicast streams on an interface.

Procedure

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

4. Select a VLAN.
5. Click **IP**.
6. Click the **IGMP** tab.
7. For StreamLimitEnable, select **enable**.
8. Configure the maximum number of streams.
9. Click **Apply**.

Configuring multicast stream limitation on a port

Configure multicast stream limitation to limit the number of concurrent multicast streams on the port. Limit the number of streams to protect the bandwidth on a specific port and control access to multicast streams.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **IGMP** tab.
5. In the StreamLimitEnable field, select the **Enable** option button.
6. Configure the maximum number of streams.
7. Click **Apply**.

Configuring multicast stream limitation members

Configure multicast stream limitation members on ports of the specified interface to configure the maximum number of streams on the interface.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.

4. Click **Insert**.
5. Type the number of the VLAN to which you want to add a member or click **Vlan** to select an ID from the list.
6. Type the number of the slot and port that you want to add as a member or click **Port**, and then select one from the graphic display. 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.

   Important:

   You must select one of the ports in the VLAN that you selected in step 4.
7. Type a maximum number of streams or accept the default of 4.
8. Click **Insert**.

StreamLimit Members field descriptions

Use the data in the following table to use the **StreamLimit Members** tab.

Name	Description
IfIndex	Displays the ID of the VLAN.
Port	Lists each slot and port number for this interface with stream limitation enabled. 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.
MaxStreams	Configures the maximum number of allowed streams for this specific port. The number of allowed streams cannot exceed the maximum number for the interface. The range is from 0–65535 and the default is 4.
NumStreams	Displays the current number of streams received on this interface. This value is a read-only value.

Deleting multicast stream limitation member

Delete a multicast stream limitation member from an interface to remove it from the configuration.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **StreamLimit Members** tab.
4. Click on the row that lists the member you want to delete.
5. Click **Delete**.

Configuring the IGMP interface

Configure the IGMP interface to change global IGMP values for the interface. Use the Interface tab to view or edit the IGMP interface table.

About this task

If an interface does not use an IP address, it does not appear in the IGMP table. If an interface uses an IP address, but PIM-SM is not enabled, the interface appears as notInService in the Status field.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**
2. Click **IGMP**.
3. Click the **Interface** tab.
4. Edit the appropriate information.
5. Click **Apply**.

Interface Field Descriptions

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

Name	Description
IfIndex	Shows the interface where IGMP is enabled.
QueryInterval	Configures the frequency (in seconds) at which the interface transmits IGMP host query packets. The default is 125.
Status	Shows the IGMP row status. If an interface uses an IP address and PIM-SM is enabled, the status is active. Otherwise, it is notInService.

Table continues…

Name	Description
Version	Configures the version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	Shows the version of IGMP that currently runs on this interface.
Querier	Shows the address of the IGMP querier on the IP subnet to which this interface attaches.
QueryMaxResponseTime	Configures the maximum response time (in tenths of a second) advertised in IGMPv2 general queries on this interface. You cannot configure this value for IGMPv1. Smaller values allow a router to prune groups faster. The default is 100 tenths of a second (equal to 10 seconds.) **Important:** You must configure this value lower than the QueryInterval.
WrongVersionQueries	Shows the number of queries received with an IGMP version that does not match the interface. You must configure all routers on a LAN to run the same version of IGMP. If the interface receives queries with the wrong version, this value indicates a version mismatch.
Joins	Shows the number of times this interface added a group membership, which is the same as the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	Tunes for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, increase the robustness value. The default value of 2 means that the switch drops one query for each query interval without the querier aging out.
LastMembQueryIntvl	Configures the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages. This value is also the time between group-specific query messages. You cannot configure this value for IGMPv1. Decrease the value to reduce the time to detect the loss of the last member of a group. The range is from 0–255 and the default is 10 tenths of second. It is recommended that you configure this parameter to values greater than 3. If you do not need a fast leave process, you can configure values greater than 10. (The value 3 is equal to 0.3 seconds and 10 is equal to 1 second.)
OtherQuerierPresent Timeout	Shows the length of time that must pass before a multicast router determines that no other querier exists. If the local router is the querier, the value is 0.

Table continues…

Name	Description
FlushAction	Configures the flush action to one of the following: • none • flushGrpMem • flushMrouter • flushSender
RouterAlertEnable	Instructs the router to ignore IGMP packets that do not contain the router alert IP option. If you disable this variable (default configuration), the router processes IGMP packets regardless of the status of the router alert IP option. **Important:** To maximize network performance, configure this parameter according to the version of IGMP currently in use. • IGMPv1—Disable • IGMPv2—Enable • IGMPv3—Enable
SsmSnoopEnable	Enables SSM snoop.
SnoopQuerierEnable	Enables IGMP Layer 2 Querier.
SnoopQuerierAddr	Specifies the pseudo address of the IGMP snoop querier.
ExplicitHostTrackingEnable	Enables or disables IGMPv3 to track hosts for each channel or group. The default is disabled. You must select this field if you want to use fast leave for IGMPv3.
McastMode	Indicates the protocol configured on the VLAN. • snoop — Indicates IGMP snooping is enabled on a VLAN. • snoop-spb — Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over Fabric Connect for a Layer 2 VSN). • pim — Indicates PIM is enabled. • routed-spb — Indicates IP multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts.

Configuring IGMP sender entries

Configure IGMP sender entries to identify a source that sends multicast data to a multicast group.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Sender** tab.
4. Change the appropriate options.
5. Click **Apply**.

Sender field descriptions

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

Name	Description
IfIndex	Specifies the interface where you enabled the IGMP entry.
GrpAddr	Specifies the multicast group address of the multicast stream. Within the indicated valid range (224.0.0.0 to 239.255.255.255), the following are invalid addresses: 244.0.0.x and the corresponding 31 multicast addresses that map to the IP MAC addresses. If you try to select them, you receive an invalid message.
MemberAddr	Specifies the IP address of a host.
Action	Flushes an entry or a group.
TPort	Identifies the T port.
State	Indicates whether a sender exists because of an IGMP access filter. The options are filtered and not filtered.
L2Isid	Specifies the Layer 2 I-SID of the C-VLAN.

Configuring fast leave mode

Configure fast leave mode to control all IGMP fast leave enabled interfaces.

Before you begin

- You must enable explicit-host-tracking before configuring fast-leave mode. To enable explicit-host-tracking, see Configuring IGMP parameters on a port on page 126 and Configuring IGMP parameters on a VLAN on page 128.

About this task

Fast leave relies on an alternative leave process where the switch stops sending traffic for the group immediately after it receives a leave message, without issuing a query to check if other group members exist on the network. Use this global parameter to alter the leave processing on fast leave enabled IGMPv2, IGMPv3, and IGMP snoop interfaces.

Important:

Fast leave mode applies only to fast leave enabled IGMP interfaces.

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Global** tab.
4. Select the mode.
5. Click **Apply**.

Global field descriptions

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

Name	Description
FastLeaveMode	Configures the mode to one of the following values: • multipleUser: Removes from the group only the IGMP member who sent the leave message. Traffic does not stop if other receivers exist on the interface port. This value is the default. • oneUser: Removes all group members on a fast leave enabled interface port after receiving the first leave message from a member. This behavior is the same as the conventional fast leave process.
GenerateTrap	Generates a trap. The default is disable.
GenerateLog	Generates a log message. The default is disable.

Configuring multicast access control for an interface

Configure multicast access control for a selected IGMP interface or VLAN to restrict access to certain multicast streams and to protect multicast streams from spoofing (injecting data to the existing streams).

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

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

2. Click **IGMP**.
3. Click the **Access Control** tab.
4. Click **Insert**.
5. Type the number of the slot and port or VLAN ID that you want to add as a member or click the appropriate button, and then select one from the graphic display.
6. Click the ellipsis button (**...**) next to **PrefixListId**.
7. Select a prefix list ID.
8. Click **OK**.
9. Type the host address and host mask.
10. Select the action mode that you want for the specified host.
11. Click **Insert**.

Access Control field descriptions

Use the data in the following table to use the **Access Control** tab.

Name	Description
IfIndex	Specifies the interface where the IGMP entry is enabled.
PrefixListId	Specifies a numeric string that identifies the prefix list.
HostAddr	Specifies the IP address of the host.
HostMask	Specifies the subnet mask that determines the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the network for the host.
PrefixListName	Specifies the name of the prefix list.
ActionMode	Specifies the action for the host identified by HostAddr. The options include the following: • denied IP multicast transmitted traffic (deny-tx). • denied IP multicast received traffic (deny-rx). • denied both IP multicast transmitted and received traffic (deny-both). • allowed IP multicast transmitted traffic (allow-only-tx). • allowed IP multicast received traffic (allow-only-rx). • allowed both IP multicast transmitted and received traffic (allow-only-both).

Viewing IGMP cache information

View IGMP cache information to view the group for which members exist on a specific interface.

About this task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**
3. Click the **Cache** tab.

Cache field descriptions

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

Name	Description
Address	Shows the IP multicast group address for this entry that contains this information.
IfIndex	Shows the interface from which the corresponding multicast group address is heard.
LastReporter	Shows the IP address of the source of the last membership report received for this IP multicast group address on this interface. If no membership report is received, the object uses the value 0.0.0.0.
ExpiryTime	Shows the amount of time (in seconds) that remain before this entry ages out.
Version1HostTimer	Shows the time that remains until the local router assumes that no IGMPv1 members exist on the IP subnet attached to the interface. Upon hearing IGMPv1 membership report, this value resets to the group membership timer. When the time that remains is nonzero, the local router ignores IGMPv2 leave messages for this group that it receives on this interface.
Type	Shows the type of IGMP entry.
StaticPorts	Shows the static ports associated with the entry.

Viewing IGMPv3 cache

View the IGMPv3 specific data corresponding to each interface, port, and multicast group pair on a router.

About this task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **IGMPv3 Cache** tab to view the IGMPv3 cache information.

IGMPv3 Cache field descriptions

Use the data in the following table to use the **IGMPv3 Cache** tab.

Name	Description
GroupAddress	Specifies the Multicast group Address (Class D) that others want to join. A group address can be the same for many incoming ports.
IfIndex	An unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports from various sources.
InPort	An unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports from various sources.
ModeExpiryTimer	Represents the time remaining before the interface EXCLUDE state expires and the interface state transitions to INCLUDE mode. This value is applicable only to IGMPv3-compatible nodes.
Version1HostTimer	Specifies the time remaining until the local router assumes that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. This entry only applies to IGMPv1 hosts. Upon hearing any IGMPv1 report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 Leave messages for this group that it receives on this interface.
Version2HostTimer	Specifies the time remaining until the local router assumes that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. Assuming no IGMPv1 hosts have been detected, the local router does not ignore any IGMPv2 Leave messages for this group that it receives on this interface.
SourceFilterMode	Specifies the current group state, applicable to IGMPv3-compatible nodes. The value indicates whether the state is INCLUDE or EXCLUDE.

Viewing and editing multicast router discovery information

View multicast router discovery information to view the current configuration.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Multicast Router Discovery** tab.
4. To edit the current configuration, double-click the value, make the change, and then click **Apply**.

Multicast Router Discovery field descriptions

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

Name	Description
Interface	Shows the interface where IGMP is enabled.
MrdiscEnable	Enables (true) or disables (false) the router interface to listen for multicast router discovery messages to determine where to send multicast source data and IGMPv2 reports. If you enable snoop, you automatically enable multicast router discovery.
DiscoveredRouterPorts	Lists ports that the Multicast Router Discovery (MRDISC) protocol discovers. **Important:** The switch does not support the MRDISC protocol on brouter ports.
MaxAdvertiseInterval	Shows the maximum time allowed between sending router advertisements from the interface, in seconds. The range is from 2–180 seconds. The default is 20 seconds.
MinAdvertiseInterval	Shows the minimum time allowed between sending unsolicited router advertisements from the interface, in seconds. This value must be more than 3 seconds but no greater than the value assigned to the MaxAdvertiseInterval value.
MaxInitialAdvertiseInterval	Configures the maximum number (in seconds) of multicast advertisement intervals that you can configure on the switch.
MaxInitialAdvertisements	Configures the maximum number of initial multicast advertisements that you can configure on the switch.

Table continues…

Name	Description
NeighborDeadInterval	Shows the time interval (in seconds) before the router interface drops traffic after a user leaves the multicast group.

Viewing the IGMP router source list

View the source list entries corresponding to each interface and multicast group pair on a router.

About this task

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Igmp Router Source List** tab to view the IGMPv3 cache information.

Igmp router source list field descriptions

Use the data in the following table to use the **Igmp Router Source List** tab.

Name	Description
GroupAddress	Specifies the IP multicast group address for which this entry contains information.
IfIndex	Specifies the interface for which this entry contains information for an IP multicast group address.
InPort	Specifies a unique value to identify a physical interface or a logical interface (VLAN), which has received Group reports for this source.
HostAddress	Specifies the host address to which this entry corresponds.
MemberAddress	Specifies the IP Address of a member that has sent source specific report wishing to join this source.
Expire	This value indicates the relevance of the source list entry, where a non-zero value indicates this is an INCLUDE state value, and a zero value indicates this to be an EXCLUDE state value.
Mode	Specifies the current member state, applicable to IGMPv3-compatible nodes. The value indicates whether the state is INCLUDE or EXCLUDE.
MemberExpire	This value indicates the time until the member for this source expires.

Viewing IGMP snoop information

View information about IGMP snoop to see the current configuration.

About this task

You can configure IGMP on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Snoop** tab.

Snoop field descriptions

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

Name	Description
Interface	Shows the VLAN ID for the VLAN.
SnoopEnable	Shows the status of IGMP snoop. IGMP snoop works only if a multicast router exists in the VLAN.
SsmSnoopEnable	Shows the status of SSM snoop.
ProxySnoopEnable	Indicates whether the IGMP report proxy feature is enabled. If you enable this feature, the switch forwards reports from hosts to the multicast router once for each group for each query interval, or after new group information is available. If you disable this feature, the switch forwards all reports from different hosts to multicast routers, and can forward more than one group report for the same multicast group for each query interval. The default is enabled.
FastLeaveEnable	Shows the status of fast leave for this port.
FastLeavePortMembers	Lists ports that are enabled for fast leave.
SnoopMRouterPorts	Shows the configuration of ports as multicast router ports. Such ports attach to a multicast router, and forward multicast data and group reports to the router. **Important:** Configure this variable only if you use multiple multicast routers that do not attach to one another, but attach to the VLAN (technically, an invalid configuration). If multicast routers use a route between them (the valid configuration) and you configure this variable, a multicast loop forms.

Table continues…

Name	Description
SnoopActiveMRouterPorts	Shows the active multicast router ports. Active multicast router ports are ports that directly attach to a multicast router. These ports include the querier port and all ports in the forwarding state that you configure as well as those that were dynamically learned through receiving queries.
SnoopMRouterExpiration	Indicates the time that remains before the multicast router ages out. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The query maximum response interval (obtained from the queries received) is used as the timer resolution.

Viewing IGMP Snoop Trace Information

View the multicast group trace to track the data flow path of multicast streams.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Snoop Trace** tab.

Snoop Trace Field Descriptions

Use the data in the following table to use the **Snoop Trace** tab.

Name	Description
GrpAddr	Displays the IP multicast address of the group traversing the router.
SrcAddr	Displays the IP source address of the multicast group.
OutVlan	Displays the egress VLAN ID for the multicast group.
InPort	Displays the ingress port for the multicast group.
InVlan	Displays the ingress VLAN ID for the multicast group.
OutPort	Displays the egress port of the multicast group.
Type	Displays the port type on which the snoop entry is learned.

Viewing IGMP Group Information

View information about IGMP groups to see the current group operation on the switch.

About this task

Note:

The following procedure displays the dynamically learned IGMP groups. **IP** > **IGMP** > **Static** displays statically configured IGMP groups. This is in contrast to the CLI command **`show ip igmp group`**, which displays both dynamically learned and statically configured IGMP groups, and the CLI command **`show ip igmp static`**, which displays only the statically configured groups.

You can view IGMP information on a VRF instance the same way you view the Global Router except that you must first launch the appropriate VRF context.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **IGMP**.
3. Click the **Groups** tab.

Groups Field Descriptions

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

Name	Description
IpAddress	Shows the multicast group address (Class D). A group address can be the same for many incoming ports.
Members	Shows the IP address of the host that issues the membership report to this group.
InPort	Shows the port that receives the group membership report.
IfIndex	Shows a unique value that identifies a physical interface or a logical interface (VLAN) that receives the membership report.
Expiration	Shows the time left before the group report expires on this port. This variable is updated after the port receives a group report.

Chapter 11: Route management using the CLI

With multicast route commands, you can configure and view IP multicast routing parameters on the switch.

Configuring multicast stream limits

Limit the number of multicast streams to protect the CPU from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

About this task

You can enable or disable the mroute stream limit for the entire device or for individual ports when the switch is operating. If you enable the mroute stream limit for the device and for an individual port, only the periodic check is performed for that port.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Enable stream limitation globally:

   ```
   ip mroute stream-limit
   ```

3. Enter GigabitEthernet Interface Configuration mode.

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

4. Enable stream limits:

   ```
   ip mroute stream-limit
   ```

5. For Gigabit Ethernet interfaces, configure the maximum number of streams and the interval at which to sample:

```
ip mroute max-allowed-streams <1–32768> max-allowed-streams-timer-
check <1–3600>
```

6. Show the mroute stream limit configuration:

```
show ip mroute interface gigabitethernet [{slot/port[/sub-port][-
slot/port[/sub-port]][,...]}]
```

Example

```
Switch:1(config)#ip mroute stream-limit
Switch:1(config)#interface gigabitethernet 3/6
Switch:1(config-if)#ip mroute stream-limit
Switch:1(config-if)#ip mroute max-allowed streams 1000 max-allowed-streams-timer-check 20
```

Variable definitions

Use the data in the following table to use the **`interface`** command.

Variable	Value
<1-4059>	Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1.
{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.

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

Variable	Value
max-allowed-streams *<1–32768>*	Configures the maximum number of streams on the specified port. The port is shut down if the number of streams exceeds this limit. The value is a number between 1–32768. The default value is 1984 streams. To configure this option to the default value, use the default operator with the command.
max-allowed-streams-timer-check *<1–3600>*	Configures the sampling interval, which checks if the number of ingress multicast streams to the CPU is under a configured limit or if the port needs to shut down. The range is between 1–3600. The default value is 10 seconds. To configure this option to the default value, use the default operator with the command.

Job aid

The following message appears if the system shuts down the port due to excessive multicast streams:

```
Shutdown port <port> due to excessive multicast streams <# of streams
ingressed>; Configured limit max streams <configured limit> in
<configured sampling interval> sec. Please disable and re-enable the
port.
```

The following table shows the field descriptions for the **`show ip mroute interface`** command.

Table 8: show ip mroute interface field descriptions

Field	Description
PORT	Indicates the slot and port number.
MROUTE STR LIMIT	Indicates the maximum number of multicast streams that can enter the CPU through this port.
MROUTE STR LIMIT TIMER	Indicates the sampling period (in seconds) to check the number of multicast streams that enter the CPU through this port.
ENABLE	Indicates the status of the mroute stream limit on the port.

Configuring multicast static source groups

Configure static source group entries in the Protocol Independent Multicast (PIM) multicast routing table. The PIM cannot prune these entries from the distribution tree.

Before you begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - PIM-Sparse Mode (SM)
  - PIM-Source Specific Multicast (SSM)

About this task

Even if no receivers exist in the group, the multicast stream for a static source group entry remains active.

The maximum number of static source groups must not exceed 1024.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a static source group entry:

   ```
   ip mroute static-source-group <A.B.C.D> <A.B.C.D/X>
   ```

Example

Create a static source group for two multicast groups: 224.32.2.1 and 226.50.2.2. The static source group for group 224.32.2.1 is for a source subnet 10.10.10.0/24. The static source group for group 226.50.2.2 is for the host 20.20.20.100/32.

```
Switch:1(config)#  ip mroute static-source-group 224.32.2.1 10.10.10.0/24
Switch:1(config)#  ip mroute static-source-group 226.50.2.2 20.20.20.100/32
```

Variable definitions

Use the definitions in the following table to use the **`ip mroute static-source-group`** command.

Variable	Value
A.B.C.D	Specifies the IP address of the multicast group. Use the no operator to later remove this configuration.
A.B.C.D/X	Specifies the multicast source IP address and subnet mask for the static source group entry. You cannot create duplicate groups. How you configure the source address depends on the protocol and mode you use. Use the no operator to later remove this configuration.

Configuring IP multicast software forwarding

When you use the IP multicast software forwarding feature you can avoid initial data loss experienced by multicast applications; this is suitable for low bandwidth conditions.

When you configure the IP multicast software forwarding feature the system forwards the initial packets of an IP multicast data stream it receives and creates a corresponding hardware record for subsequent packets.

By default, multicast software forwarding is disabled.

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

IP multicast software forwarding is a global system configuration feature that is only applicable to traditional PIM protocol and IGMP Snooping protocols, not SPB-PIM Gateway or Layer 3 VSN SPB

Multicast. If you enable IP multicast software forwarding, the hardware continues to forward IP multicast traffic. The software only forwards initial data traffic.

After a new data stream arrives, the first data packet is sent to the CPU, which programs the multicast route in hardware, and all packets that arrive subsequent to this programming are forwarded by hardware only.

If you enable software forwarding, all initial packets received before hardware programming is complete are sent to the CPU for forwarding and packet suppression by the hardware is disabled.

If you do not enable software forwarding, only the first data packet is sent to the CPU and subsequent packets are suppressed by the hardware so that the CPU is not overwhelmed with traffic. During this time, packets suppressed by the hardware are dropped.

Important:

To avoid overloading the CPU, ensure that you do not use the IP multicast software forwarding feature for video multicast applications.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable software forwarding:

   ```
   multicast software-forwarding
   ```

3. Show the software forwarding configuration:

   ```
   show multicast software-forwarding
   ```

Example

```
Switch:1#show multicast software-forwarding

================================================================================
                    Mcast Software Forwarding - GlobalRouter
================================================================================
McastSoftwareForwarding        :enabled
```

Configuring the resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing the switch.

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

After you configure the counter thresholds for ingress and egress records, if the record usage exceeds the threshold, you receive notification by a trap on the console, a logged message, or both.

If you do not configure the thresholds, the system displays only the ingress and egress records currently in use.

You can configure the resource usage counter on a VRF instance the same way you configure the Global Router except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure the thresholds:

   ```
   ip mroute resource-usage egress-threshold <0-32767> ingress-threshold <0-32767>
   ```

3. Configure one of the following notification methods:
   - Configure a log-only notification method:

     ```
     ip mroute resource-usage log-msg
     ```
   - Configure a trap-only notification method:

     ```
     ip mroute resource-usage trap-msg
     ```
   - Configure both notification methods:

     ```
     ip mroute resource-usage log-msg trap-msg
     ```

Example

Configure the egress threshold to 200.

```
Switch:1(config)# ip mroute resource-usage egress-threshold 200
```

Configure the ingress threshold to 100.

```
Switch:1(config)# ip mroute resource-usage ingress-threshold 100
```

Enable the log message notification method.

```
Switch:1(config)# ip mroute resource-usage log-msg
```

Variable definitions

Use the data in the following table to use the **`ip mroute resource-usage`** command.

Variable	Value
egress-threshold *<0–32767>*	Configures the egress record threshold (S,G). The system sends a notification message after the number of streams exceeds a threshold level. To configure this option to the default value, use the default operator with the command. The default is 0.
ingress-threshold *<0–32767>*	Configures the ingress record threshold. The system sends a notification message after the number of streams exceeds a threshold level. To configure this option to the default value, use the default operator with the command. The default is 0.

Configuring prefix lists

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or noncontiguous routes. Reference prefix lists by name from within a routing policy.

About this task

Important:

When you configure a prefix list for a route policy, add the prefix as a.b.c.d/32. You must enter the full 32-bit mask to exact a full match of a specific IP address.

You configure prefix lists on a VRF instance the same way you configure the GlobalRouter, except that you must use VRF Router Configuration mode.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a prefix list:

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

3. **(Optional)** Rename an existing prefix list:

   ```
   ip prefix-list WORD<1-64> name WORD<1-64>
   ```

4. Display the prefix list:

   ```
   show ip prefix-list [prefix {A.B.C.D}] [vrf WORD<0-16>] [vrfids
   WORD<0-512>] [WORD <1-64>]
   ```

Example

Configure a prefix-list. Display the prefix list.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list LIST1 47.17.121.50/255.255.255.0
Switch:1(config)#show ip prefix-list LIST1
================================================================================
                          Prefix List - GlobalRouter
================================================================================

        PREFIX           MASKLEN FROM TO
--------------------------------------------------------------------------------

List 1    LIST1:
        47.17.121.50     24      24   24
1 Total Prefix List entries configured
--------------------------------------------------------------------------------
Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
```

Variable definitions

Use the data in the following table to use the **`ip prefix-list`** command.

Variable	Value
{A.B.C.D/X}	Specifies the IP address and the mask in one of the following formats: • a.b.c.d/x • a.b.c.d/x.x.x.x • default
ge <0–32>	Specifies the minimum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
le <0–32>	Specifies the maximum length to match. Lower bound and higher bound mask lengths together can define a range of networks.
name WORD<1-64>	Renames the specified prefix list. The name length is 1–64 characters.
WORD<1-64>	Specifies the name for a new prefix list.

Use the data in the following table to use the **`show ip prefix-list`** command.

Variable	Value
{A.B.C.D}	Specifies the prefix to include in the command output.

Table continues…

Variable	Value
vrf WORD<0-16>	Specifies the name of the VRF.
vrfids WORD<0-512>	Specifies the ID of the VRF and is an integer in the range of 0–512.
WORD<1-64>	Specifies a prefix list, by name, to use for the command output.

Use the following table to use the **`show ip prefix-list`** command output.

Variable	Value
PREFIX	Indicates the member of a specific prefix list.
MASKLEN	Indicates the prefix mask length in bits.
FROM	Indicates the prefix mask starting point in bits.
TO	Indicates the prefix mask endpoint in bits.

Chapter 12: Route management using EDM

View or edit interface configuration information for Layer 3 IP multicast protocols on the switch.

Viewing multicast route information

View multicast route information for troubleshooting purposes.

This tab shows multicast routing information for IP datagrams from a particular source and addressed to a particular IP multicast group address.

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

You can view the multicast routes for a Layer 3 Virtual Services Network (VSN) the same way you view the Global Router except that you must first launch the appropriate VRF context. For more information about Layer 3 VSNs, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP** > **Multicast**.
2. Click the **Routes** tab.

Routes field descriptions

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

Name	Description
Group	Displays the IP multicast group address for this entry that contains multicast routing information.
Source	Displays the network address that, when combined with the corresponding route SourceMask value, identifies the source that contains multicast routing information.

Table continues…

Name	Description
SourceMask	Displays the network mask that, when combined with the corresponding route Source value, identifies the multicast source.
UpstreamNeighbor	Shows the address of the upstream neighbor from which the IP datagrams from these sources are received. The address is 0.0.0.0 if the network is local.
Interface	Displays the interface, slot and portnumber, or VLAN ID where IP datagrams sent by these multicast sourcesto this multicast address are received.
ExpiryTime	Displays the amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb (12) • spbpimgw(13)

Viewing multicast next-hop information

View all multicast next-hop information.

This tab shows information about the next hops used by outgoing interfaces to route IP multicast datagrams. Each entry is one in a list of next hops on outgoing interfaces for particular sources that send to a particular multicast group address.

About this task

You can view the multicast routes for a Layer 3 Virtual Services Network (VSN) the same way you view the Global Router except that you must first launch the appropriate VRF context. For more information about Layer 3 VSNs, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Procedure

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

Next Hops field descriptions

Use the data in the following table to use the **Next Hops** tab.

Name	Description
Group	Displays the IP multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
SourceMask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
ReceiverPort	Displays the receiver port for this next hop.
OutInterface	Displays the interface slot and portnumber or VLAN ID for the outgoing interface for this next hop.
Address	Displays the address of the next hop specific to this entry. For most interfaces, it is identical to the next-hop group. Non Broadcast Multiple Access (NBMA) interfaces, however, can use multiple next hop addresses out of a single outgoing interface.
State	Displays whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. A value of forwarding indicates the information is currently used; pruned indicates it is not used.
UpTime	Displays the up time for this entry.
ExpiryTime	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	Displays the minimum number of hops between this router and members of the IP multicast group reached through the next hop on this outgoing interface. IP multicast datagrams for the group that use a time-to-live less than this number of hops are not forwarded to the next hop.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb
Pkts	Displays the number of next hop packets.

Viewing multicast interface information

View multicast interface information to verify the multicast configuration.

This tab shows multicast routing information specific to interfaces.

About this task

You can view multicast interface information for a Layer 3 VSN the same way you view the Global Router except that you must first launch the appropriate VRF context. For more information about Layer 3 VSNs, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Interfaces** tab.

Interfaces field descriptions

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

Name	Description
Interface	Displays the slot and port number or VLAN ID for this entry.
Ttl	Displays the datagram time-to-live (TTL) threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means that all multicast packets are forwarded out of the interface.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

Adding new static source groups

Add a new static source group to create an entry that the switch cannot prune from the distribution tree. An attempt to add a duplicate of an existing source-group entry results in an error message.

Before you begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - PIM-SM
  - PIM-SSM

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

The switch supports PIM only in the Global Router. You cannot configure static source groups for specific VRF contexts.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Static Source Group** tab.
4. Click **Insert**.
5. Complete the information in the dialog box.
6. Click **Insert**.

Editing static source groups

Configure static source-group entries in the PIM multicast routing table. PIM cannot prune these entries from the distribution tree. In other words, even if no receivers exist in the group, the multicast stream for a static source-group entry stays active.

Before you begin

- Before you can configure a static source group, you must globally enable one of the following protocols:
  - PIM-Sparse Mode (SM)
  - PIM-Source Specific Multicast (SSM)

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

The maximum number of static source groups must not exceed 1024.

The switch supports PIM only in the Global Router. You cannot configure static source groups for specific VRF contexts.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Static Source Group** tab.
4. Edit the required information.
5. Click **Apply**.

Static Source Group field descriptions

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

Name	Description
GroupAddress	Configures the multicast group IP address for this static source-group entry.
SourceSubnet	Configures the multicast source address for this static source-group entry. How you configure the source address depends on the protocol and mode you use.
SrcSubnetMask	Configures the subnet mask of the source for this static source-group entry.

Configuring IP multicast software forwarding

Configure IP multicast software forwarding to enable the system to initially forward IP multicast data until a hardware record is created. The system forwards the initial packets of a stream it receives and creates a corresponding hardware record for subsequent packets. The advantage of this feature is that it avoids initial data loss experienced by multicast applications and is most suited for low bandwidth.

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

IP multicast software forwarding is a global system configuration feature that is only applicable to traditional PIM protocol and IGMP Snooping protocols, not SPB-PIM Gateway or Layer 3 VSN SPB Multicast. If you enable IP multicast software forwarding, the hardware still forwards IP multicast traffic. The software forwards only initial data traffic.

After a new data stream arrives, the first data packet is sent to the CPU, which programs the multicast route in hardware, and all packets that arrive subsequent to this programming are forwarded by hardware only. If you enable software forwarding, all initial packets received before hardware programming is complete are sent to the CPU for forwarding. If you enable software forwarding, packet suppression by the hardware is disabled. If you do not enable software forwarding, only the first data packet is sent to the CPU and subsequent packets are suppressed by the hardware so that the CPU is not overwhelmed with traffic. During this time, packets suppressed by the hardware are dropped.

By default, the feature is disabled.

Important:

To avoid overloading the CPU, do not use the IP multicast software forwarding feature for video multicast applications.

If you configure multicast software forwarding from within a VRF context, the configuration applies to the Global Router and all VRF contexts. You cannot change the multicast software forwarding configuration for individual VRF contexts.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Globals** tab.
4. Select the **SWForwardingEnable** check box.
5. Click **Apply**.

Globals field descriptions

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

Name	Description
SWForwardingEnable	Enables the system to initially forward IP multicast data until a hardware record is created. The default is disabled.
StatsEnabled	Enables or disables multicast route statistics. The default is disabled.
StatsClear	Clears multicast route statistics.

Configuring mroute stream limit

Limit the number of multicast streams to protect a CPU from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

Procedure

1. On the Device Physical View tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.
4. Select the **Mroute Stream Limit** tab.
5. Select the **StreamLimitEnable** box.
6. Edit other fields as required.
7. Click **Apply**.

Mroute Stream Limit field descriptions

Use the data in the following table to use the **Mroute Stream Limit** tab.

Name	Description
StreamLimitEnable	Enables or disables mroute stream limit on the port.
StreamLimit	Specifies the maximum number of multicast streams allowed to enter the CPU through this port.
StreamTimerCheck	Specifies the sampling period, in seconds, to check the number of multicast streams that enter the CPU through this port.

Configuring Mroute Stream Limit on an Insight Port

About this task

Perform this procedure to limit the number of multicast streams to protect a Central Processing Unit (CPU) from multicast data packet bursts generated by malicious applications, such as viruses that cause the CPU to reach 100 percent utilization, or that prevent the CPU from processing protocol packets or management requests. If more than a certain number of multicast streams ingress to a CPU through a port during a sampling interval, the port shuts down until you take appropriate action.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Insight Port**.

2. Click the Insight port you want to configure.
3. Click the **Mroute Stream Limit** tab.
4. Select **StreamLimitEnable**.
5. Configure other fields as required.
6. Click **Apply**.

Mroute Stream Limit Field Descriptions

Use data in the following table to configure the Mroute Stream Limit tab.

Name	Description
StreamLimitEnable	Enables or disables mroute stream limit on the Insight port. The default is disabled.
StreamLimit	Specifies the maximum number of multicast streams allowed to enter the CPU through the Insight port. The default value is 1984.
StreamTimerCheck	Specifies the sampling period, in seconds, to check the number of multicast streams that enter the CPU through the Insight port. The default is 10 seconds.

Configuring resource usage counter for multicast streams

Configure the resource usage counters to query the number of ingress and egress IP multicast streams traversing the switch. After you configure the counter thresholds for ingress and egress records, if the record usage goes beyond the threshold, you receive notification through a trap on the console, a logged message, or both.

About this task

Note:

This procedure is supported only on a DvR Controller. It is not supported on a DvR Leaf node.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Important:

If you do not configure the thresholds, EDM displays only the ingress and egress records that are currently in use.

You can configure the resource usage counter on a VRF instance the same way you configure the Global Router except that you must first launch the appropriate VRF context.

Procedure

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

2. Click **Multicast**.
3. Select the **Resource Usage** tab.
4. Configure the ingress and egress thresholds.
5. Configure the notification methods.
6. Click **Apply**.

Resource Usage field descriptions

Use the data in the following table to use the **Resource Usage** tab.

Name	Description
Egress Records In-Use	Displays the number of egress records traversing the switch.
Ingress Records In-Use	Displays the number of ingress records (source or group) traversing the switch.
Egress Threshold	Configures the egress threshold level (0–32767).
Ingress Threshold	Configures the ingress threshold level (0–32767).
SendTrapOnly	Sends only trap notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type. You can configure only one notification type.
SendTrapAndLog	Sends both trap and log notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type.
LogMsgOnly	Sends only log notification messages after the number of streams exceeds a threshold level. Select disable if you select a different notification type.

Configuring a prefix list

Configure a prefix list to allow or deny specific route updates. A prefix list policy specifies route prefixes to match. After a match occurs, the system uses the route.

The prefix list contains a set of contiguous or non-contiguous routes. Reference prefix lists by name from within a routing policy.

Before you begin

- Change the VRF instance as required to configure a prefix list on a specific VRF instance.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Policy**.

3. Click the **Prefix List** tab.
4. Click **Insert**.
5. In the **Id** box, type an ID for the prefix list.
6. In the **Prefix** box, type an IP address for the route.
7. In the **PrefixMaskLength** box, type the length of the prefix mask.
8. Configure the remaining parameters as required.
9. Click **Insert**.

Prefix List field descriptions

Use the data in the following table to use the **Prefix List** tab.

Name	Description
Id	Configures the list identifier.
Prefix	Configures the IP address of the route.
PrefixMaskLen	Configures the specified length of the prefix mask. You must enter the full 32-bit mask to exact a full match of a specific IP address, for example, if you create a policy to match on the next hop.
Name	Names a specified prefix list during the creation process or renames the specified prefix list. The name length can use from 1 to 64 characters.
MaskLenFrom	Configures the lower bound of the mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	Configures the higher bound mask length. The default is the mask length. Lower bound and higher bound mask lengths together can define a range of networks.

Chapter 13: Multicast route statistics configuration using the CLI

The following sections provide procedural information you can use to configure multicast route statistics using the Command Line Interface (CLI).

Important:

This feature does not apply to all hardware platforms. For more information about feature support, see Release Notes for VSP 8600.

Enabling IP multicast route statistics

Enable the collection and display of IP multicast route statistics.

These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on multicast group classification. By default, collection of multicast route statistics is disabled.

Note:

When you enable IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable the collection of IP multicast route statistics.

   ```
   ip mroute stats enable
   ```

3. **(Optional)** Set the IP multicast route statistics to default.

   ```
   default ip mroute stats enable
   ```

4. **(Optional)** Disable the collection of IP multicast route statistics.

   ```
   no ip mroute stats enable
   ```

5. View the IP multicast route statistics.

   `show ip mroute stats [WORD<3-160> {A.B.C.D[,E.F.G.H][,...]}]`

   Note:

   The maximum number of multicast group IP addresses is 10.

Example

Enable the collection of IP multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ip mroute stats enable
```

View the IP multicast route statistics:

```
Switch:1#show ip mroute stats

================================================================================
                     Multicast Stats

================================================================================
--------------------------------------------------------------------------------
                            Statistics : Enabled
```

View the statistics for the multicast group IP address `225.0.0.1`:

```
Switch:1#show ip mroute stats 225.0.0.1

================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress      SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1         1                  30452198             3897881344
   128
```

View the statistics for multiple (up to a maximum of 10) group IP addresses.

```
Switch:1#show ip mroute stats
225.0.0.1,225.0.0.2,225.0.0.3,225.0.0.4,225.0.0.5,225.0.0.6,225.0.0.7,225.0.0.8,225.0.0.9,
225.0.0.10

================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress      SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1         1                  32446194             4153112832
   128
225.0.0.2         1                  32446196             4153112960
   127
225.0.0.3         1                  32446197             4153113088
   127
225.0.0.4         1                  32446198             4153113216
   127
225.0.0.5         1                  32446199             4153113472
   128
```

```
225.0.0.6        1                  32446200            4153113600
   128
225.0.0.7        1                  32446201            4153113728
   128
225.0.0.8        1                  32446203            4153113856
   127
225.0.0.9        1                  32446203            4153113856
   127
225.0.0.10       1                  32446203            4153113984
   128
```

Variable definitions

Use the data in the following table to use the **`show ip mroute stats`** command.

Variable	Definition
WORD<*3-160*> *{A.B.C.D[,E.F.G.H][,...]}*	Specifies the multicast group IP address for which to display statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses. You can specify a maximum of 10 groups.

Job aid

The following table shows the field descriptions for viewing multicast route statistics.

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.
SourceCounter	Specifies the number of sources associated with the multicast route record.
IngressPackets	Specifies the number of packets received for the associated IP address.
IngressBytes	Specifies the number of bytes received for the associated IP address.
AverageSize	Specifies the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the formula: ingress packet/ ingress byte.

Clearing IP multicast route statistics

Use this procedure to clear the IP multicast route statistics. This resets the IP multicast statistics counters.

Note:

When you clear IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Clear the IP multicast route statistics:

   ```
   clear ip mroute stats
   ```

Example:

Clear the IP multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clear ip mroute stats
```

Monitoring IP multicast route statistics

Use this procedure to monitor the IP multicast route statistics at regular intervals.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Monitor the IP multicast route statistics:

   ```
   monitor ip mroute stats WORD<7-160> {A.B.C.D[,E.F.G.H][,...]}
   ```

   Note:

   You can monitor a maximum of 10 group IP addresses.

Example:

Monitor the IP multicast route statistics for the group IP address `225.0.0.1`. In this example, the statistics are monitored at intervals of 5 seconds for a duration of 300 seconds.

The output from monitoring three consecutive intervals is displayed below.

```
Switch:1>en
Switch:1#monitor ip mroute stats 225.0.0.1


                          MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Mon Dec 21 16:12:07 2015


================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress     SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1        1                  4716624              603727872
   128
                          MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Mon Dec 21 16:12:13 2015


================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress     SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1        1                  4767325              610217600
   128
                          MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Mon Dec 21 16:12:19 2015

...
...
Switch:1#
```

Monitor the IP multicast route statistics for a maximum of 10 group IP addresses. The statistics are monitored at intervals of 5 seconds for a duration of 300 seconds.

The output from monitoring three consecutive intervals is displayed below.

```
Switch:1#monitor ip mroute stats
225.0.0.1,225.0.0.2,225.0.0.3,225.0.0.4,225.0.0.5,225.0.0.6,225.0.0.7,225.0.0.8,225.0.0.9,
225.0.0.10
                          MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Mon Dec 21 16:22:07 2015


================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress     SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1        1                  9532039              1220100992
   128
225.0.0.2        1                  9532041              1220101120
   127
225.0.0.3        1                  9532042              1220101248
   127
225.0.0.4        1                  9532043              1220101376
```

```
   127
225.0.0.5        1                  9532044              1220101632
   128
225.0.0.6        1                  9532045              1220101760
   128
225.0.0.7        1                  9532046              1220101888
   128
225.0.0.8        1                  9532047              1220101888
   127
225.0.0.9        1                  9532048              1220102016
   127
225.0.0.10       1                  9532048              1220102144
   128
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Mon Dec 21 16:22:13 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress     SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1        1                  9582672              1226582016
   128
225.0.0.2        1                  9582674              1226582144
   127
225.0.0.3        1                  9582675              1226582272
   127
225.0.0.4        1                  9582676              1226582400
   127
225.0.0.5        1                  9582677              1226582656
   128
225.0.0.6        1                  9582678              1226582784
   128
225.0.0.7        1                  9582679              1226582912
   128
225.0.0.8        1                  9582681              1226583040
   127
225.0.0.9        1                  9582681              1226583040
   127
225.0.0.10       1                  9582681              1226583168
   128
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Mon Dec 21 16:22:19 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress     SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
225.0.0.1        1                  9625009              1232001152
   128
225.0.0.2        1                  9625011              1232001280
   127
225.0.0.3        1                  9625012              1232001408
   127
225.0.0.4        1                  9625013              1232001536
   127
225.0.0.5        1                  9625014              1232001792
   128
```

```
225.0.0.6        1                   9625015             1232001920
   128
225.0.0.7        1                   9625016             1232002048
   128
225.0.0.8        1                   9625018             1232002176
   127
225.0.0.9        1                   9625019             1232002304
   127
225.0.0.10       1                   9625018             1232002304
   128

...
...
Switch:1#
```

Variable definitions

Use the data in the following table to use the **`monitor ip mroute stats`** command.

Variable	Definition
WORD<*7-160*> *{A.B.C.D[,E.F.G.H][,...]}*	Specifies the multicast group IP address for which to monitor statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses, up to a maximum of 10.

Enabling IPv6 multicast route statistics

Enable the collection of IPv6 multicast route statistics.

These statistics are not related to the interface (port) statistics. Rather, the statistics are displayed based on multicast group classification. By default, collection of multicast route statistics is disabled.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable the collection of IPv6 multicast route statistics.

   ```
   ipv6 mroute stats enable
   ```

3. **(Optional)** Set the IPv6 multicast route statistics to default:

   ```
   default ipv6 mroute stats
   ```

4. **(Optional)** Disable the collection of IPv6 multicast route statistics.

   ```
   no ipv6 mroute stats
   ```

5. View the IPv6 multicast route statistics.

   ```
   show ipv6 mroute stats [WORD<7-400> {Ipv6address[,Ipv6address]
   [,...]}]
   ```

   Note:

   The maximum number of multicast group IP addresses is 10.

Example:

Enable collection of IPv6 multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#ipv6 mroute stats enable
```

View the IPv6 multicast route statistics:

```
Switch:1#show ipv6 mroute stats

================================================================================
                     Multicast Stats

================================================================================
--------------------------------------------------------------------------------
                     Statistics : Enabled
```

View the statistics for the multicast group IP address `FF05::1`:

```
Switch#show ipv6 mroute stats FF05::1

================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress                SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1          1                  1962750              2355300000
   1200
```

View the statistics for multiple group IP addresses (up to a maximum of 10).

```
Switch#show ipv6 mroute stats
FF05::1,FF05::2,FF05::3,FF05::4,FF05::5,FF05::6,FF05::7,FF05::8,FF05::9,FF05::a

================================================================================
                          Multicast Stats - GlobalRouter

================================================================================
GroupAddress                SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1          1                  2027508              2433009600
   1200
ff05:0:0:0:0:0:0:2          1                  2027507              2433008400
   1200
ff05:0:0:0:0:0:0:3          1                  2027507              2433008400
   1200
ff05:0:0:0:0:0:0:4          1                  2027507              2433008400
   1200
ff05:0:0:0:0:0:0:5          1                  2027507              2433008400
```

```
   1200
ff05:0:0:0:0:0:0:6         1                  2027505              2433006000
   1200
ff05:0:0:0:0:0:0:7         1                  2027505              2433006000
   1200
ff05:0:0:0:0:0:0:8         1                  2027505              2433006000
   1200
ff05:0:0:0:0:0:0:9         1                  2027505              2433006000
   1200
ff05:0:0:0:0:0:0:a         1                  2027505              2433006000
   1200
```

Variable definitions

Use the data in the following table to use the **`show ipv6 mroute stats`** command

Variable	Definition
WORD<*7-400*> *{Ipv6address[,Ipv6address][,...]}*	Specifies the multicast group IP address for which to display statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses. You can specify a maximum of 10 groups.

Job aid

The following table shows the field descriptions for viewing multicast route statistics.

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.
SourceCounter	Specifies the number of sources associated with the multicast route record.
IngressPackets	Specifies the number of packets received for the associated IP address.
IngressBytes	Specifies the number of bytes received for the associated IP address.
AverageSize	Specifies the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the formula: ingress packet/ ingress byte.

Clearing IPv6 multicast route statistics

Use this procedure to clear the IPv6 multicast route statistics. This resets the IP multicast statistics counters.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Clear the IPv6 multicast route statistics:

   ```
   clear ipv6 mroute stats
   ```

Example:

Clear the IPv6 multicast route statistics:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#clear ipv6 mroute stats
```

Monitoring IPv6 multicast route statistics

Use this procedure to monitor IPv6 multicast route statistics at regular intervals.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Monitor IPv6 multicast route statistics:

   ```
   monitor ipv6 mroute stats WORD<7-400> {Ipv6address[,Ipv6address]
   [,...]}
   ```

   Note:

   You can monitor a maximum of 10 group IP addresses.

Example:

Monitor the IPv6 multicast route statistics for the group IPv6 address `FF05::1`. In this example, the statistics are monitored at intervals of 5 seconds for a duration of 300 seconds.

The output from monitoring three consecutive intervals is displayed below.

```
Switch:1>en
Switch:1#monitor IPv6 mroute stats FF05::1

                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Tue Dec 22 16:54:25 2015
```

```
================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress               SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1         1                  2446250              2935500000
   1200
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Tue Dec 22 16:54:31 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress               SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1         1                  2448947              2938736400
   1200
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Tue Dec 22 16:54:37 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress               SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1         1                  2452185              2942622000
   1200

...
...
Switch:1#
```

Monitor the IPv6 multicast route statistics for a maximum of 10 group IPv6 addresses. The statistics are monitored at intervals of 5 seconds for a duration of 300 seconds.

The output from monitoring three consecutive intervals is displayed below.

```
Switch:1#$stats
FF05::1,FF05::2,FF05::3,FF05::4,FF05::5,FF05::6,FF05::7,FF05::8,FF05::9,FF05::a
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Tue Dec 22 17:04:55 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress               SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1         1                  2768926              3322711200
   1200
ff05:0:0:0:0:0:0:2         1                  2768925              3322710000
   1200
ff05:0:0:0:0:0:0:3         1                  2768925              3322710000
   1200
```

```
ff05:0:0:0:0:0:0:4       1                  2768925              3322710000
   1200
ff05:0:0:0:0:0:0:5       1                  2768925              3322710000
   1200
ff05:0:0:0:0:0:0:6       1                  2768923              3322707600
   1200
ff05:0:0:0:0:0:0:7       1                  2768923              3322707600
   1200
ff05:0:0:0:0:0:0:8       1                  2768923              3322707600
   1200
ff05:0:0:0:0:0:0:9       1                  2768923              3322707600
   1200
ff05:0:0:0:0:0:0:a       1                  2768923              3322707600
   1200
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Tue Dec 22 17:05:01 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress             SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1       1                  2771625              3325950000
   1200
ff05:0:0:0:0:0:0:2       1                  2771625              3325950000
   1200
ff05:0:0:0:0:0:0:3       1                  2771625              3325950000
   1200
ff05:0:0:0:0:0:0:4       1                  2771624              3325948800
   1200
ff05:0:0:0:0:0:0:5       1                  2771624              3325948800
   1200
ff05:0:0:0:0:0:0:6       1                  2771622              3325946400
   1200
ff05:0:0:0:0:0:0:7       1                  2771622              3325946400
   1200
ff05:0:0:0:0:0:0:8       1                  2771622              3325946400
   1200
ff05:0:0:0:0:0:0:9       1                  2771622              3325946400
   1200
ff05:0:0:0:0:0:0:a       1                  2771622              3325946400
   1200
                         MULTICAST STATISTIC
Monitor Interval: 5sec | Monitor Duration: 300sec       Tue Dec 22 17:05:07 2015


================================================================================
                         Multicast Stats - GlobalRouter

================================================================================
GroupAddress             SourceCounter      IngressPackets       IngressBytes
   AverageSize
--------------------------------------------------------------------------------
ff05:0:0:0:0:0:0:1       1                  2774864              3329836800
   1200
ff05:0:0:0:0:0:0:2       1                  2774863              3329835600
   1200
ff05:0:0:0:0:0:0:3       1                  2774863              3329835600
   1200
ff05:0:0:0:0:0:0:4       1                  2774863              3329835600
   1200
ff05:0:0:0:0:0:0:5       1                  2774863              3329835600
```

```
   1200
ff05:0:0:0:0:0:0:6       1                   2774861              3329833200
   1200
ff05:0:0:0:0:0:0:7       1                   2774861              3329833200
   1200
ff05:0:0:0:0:0:0:8       1                   2774861              3329833200
   1200
ff05:0:0:0:0:0:0:9       1                   2774861              3329833200
   1200
ff05:0:0:0:0:0:0:a       1                   2774861              3329833200
   1200

...
...

Switch:1#
```

Variable definitions

Use the data in the following table to use the **`monitor ipv6 mroute stats`** command:

Variable	Definition
WORD<*7-400*> *{Ipv6address[,Ipv6address][,...]}*	Specifies the multicast group IP address for which to monitor statistics. The group IP address is in one of the following formats: a single IP address or a series of IP addresses, up to a maximum of 10.

Chapter 14: Multicast route statistics configuration using EDM

The following sections provide procedural information you can use to configure multicast route statistics using the Enterprise Device Manager (EDM).

Important:

This feature does not apply to all hardware platforms. For more information about feature support, see Release Notes for VSP 8600.

Enabling IP multicast route statistics

Use this procedure to enable IP multicast route statistics.

Note:

When you enable or clear IP multicast route statistics on the Controller node of a DvR domain, the configuration is automatically pushed to the Leaf nodes within the domain.

For more information on DvR, see Configuring IPv4 Routing for VOSS.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**
2. Click **Multicast**.
3. Click the **Globals** tab.
4. In the **StatsEnabled** field, select the option to enable or disable the collection of statistics.
5. **(Optional)** To clear the statistics, click **StatsClear**.
6. Click **Apply**.

Field Definitions

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

Field	Description
StatsEnabled	Displays whether the multicast route statistics is enabled.
StatsClear	Clears the multicast route statistics.

Viewing IP multicast route statistics

Use this procedure to view IP multicast route statistics.

Before you begin

- You must enable the collection of multicast statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IP**.
2. Click **Multicast**.
3. Click the **Stats** tab to view the statistics.

Field Definitions

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

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.
SourceCounter	Specifies the number of sources associated with the multicast route record.
Pkts	Specifies the number of packets received for the associated IP address.
Bytes	Specifies the number of bytes received for the associated IP address.
AverageSizePerPkt	Specifies the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the following formula: ingress packet/ingress byte.

Enabling IPv6 multicast route statistics

Enable the collection of IPv6 multicast route statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 Mroute**.
3. Click the **Globals** tab.
4. In the **StatsEnabled** field, select the option to enable or disable the collection of statistics.
5. **(Optional)** To clear the statistics, click **StatsClear**.
6. Click **Apply**.

Field Definitions

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

Field	Description
StatsEnabled	Displays whether the multicast route statistics is enabled.
StatsClear	Clears the multicast route statistics.

Viewing IPv6 multicast route statistics

Use this procedure to view IPv6 multicast route statistics.

Before you begin

- You must enable the collection of multicast statistics.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **IPv6**.
2. Click **IPv6 Mroute**.
3. Click the **Stats** tab to view the statistics.

Field Definitions

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

Field	Description
GroupAddress	Specifies the multicast group IP address for which to show statistics.

Table continues…

Field	Description
SourceCounter	Specifies the number of sources associated with the multicast route record.
Pkts	Specifies the number of packets received for the associated IP address.
Bytes	Specifies the number of bytes received for the associated IP address.
AverageSizePerPkt	Specifies the average packet length for the associated group IP address. This information indicates only the ingress packet length and is calculated using the following formula: ingress packet/ingress byte.

Chapter 15: CLI show command reference

This reference information provides show commands to view the operational status of multicast routing on the switch.

Note:

If you do not specify a VRF name or range of VRF IDs in the show command, the results display information for the Global Router.

If you do specify a VRF name or range of VRF IDs, the results display information only for the VRFs you specify.

General show commands

This section explains the show commands for general multicast routing operations.

Multicast route information

Use the **show ip mroute route** command to display information about the multicast routes on the switch. The syntax for this command is as follows.

```
show ip mroute route [vrf WORD <0-32>] [vrfids <0-255>]
```

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following section shows a sample output for the **show ip mroute route** command.

In this table, every stream uses one (*,G) entry and x (S,G) entries, depending on how many servers forward traffic to the same group.

The 0.0.0.0 mask is always tied to a (*,G) entry.

Every time a new stream comes in, Protocol Independent Multicast (PIM) creates two entries in the table; one is a (*,G) entry that points toward the rendezvous point (RP) router, and the other is an (S,G) entry that points toward the source.

```
Switch:1#show ip mroute route
```

```
================================================================================
                          Mroute Route - GlobalRouter
================================================================================
GROUP           SOURCE          SRCMASK         UPSTREAM_NBR    IF     EXPIR  PROT
--------------------------------------------------------------------------------
233.252.0.1       0.0.0.0         0.0.0.0         0.0.0.0        V3      30  pimsm
233.252.0.1       198.51.100.99   255.255.255.0   0.0.0.0        -        0  pimsm
233.252.0.2       0.0.0.0         0.0.0.0         0.0.0.0        V2      30  pimsm
225.1.1.1         198.51.100.99   255.255.255.0   0.0.0.0        V3      173 spb-pim-gw

Total 4
```

The following table shows the field descriptions for this command.

Table 9: show ip mroute route command

Field	Description
GROUP	Indicates the IP multicast group for this multicast route.
SOURCE	Indicates the network address that, when combined with the corresponding value of SRCMASK, identifies the sources for this multicast route.
SRCMASK	Indicates the network mask that, when combined with the corresponding value of SOURCE, identifies the sources for this multicast route.
UPSTREAM_NBR	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the (S,G) source is local or if the RP for this the (*,G) group is an address on this router.
IF	Indicates the value of ifIndex for the interface that receives IP datagrams sent by these sources to this multicast address. A value of 0 in a (*,G) route indicates that datagrams are not subject to an incoming interface check, but datagrams can be received on any interface.
EXPIR	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging. **Note:** The value you configure for fwd-cache-timeout applies only to the locally learned sender; it does not apply to SMLT synchronized sender records.
PROT	Indicates the multicast protocol through which the switch learned this route.

Multicast route next hop

Use the **show ip mroute next-hop** command to show information about the next hop for the multicast routes on the switch. The syntax for this command is as follows.

show ip mroute next-hop [vrf WORD <0-16>] [vrfids <0-512>]

The following table shows the field descriptions for this command.

Table 10: show ip mroute next-hop command

Field	Description
INTERFACE	Indicates the interface identity.
GROUP	Indicates the IP multicast group for which this entry specifies a next-hop PIM neighbor toward receivers for a specific outgoing interface.
SOURCE	Indicates the network address, which when combined with the corresponding value of SRCMASK, identifies the sources for which this entry specifies a next-hop PIM neighbor toward receivers for a specific outgoing interface.
SRCMASK	Indicates the network mask, which when combined with the corresponding value of SOURCE, identifies the sources for which this entry specifies a next-hop PIM neighbor toward receivers for a specific outgoing interface.
ADDRESS	Indicates the address of the next hop specific to this entry. The next hop must be the address of a PIM neighbor. This table does not represent local receivers.
STATE	Indicates whether the outgoing interface and next hop represented by this entry currently forward IP datagrams. The value forwarding indicates the information is currently used; the value pruned indicates it is not used.
EXPTIME	Indicates the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
CLOSEHOP	Indicates the minimum number of hops between this router and members of this IP multicast group reached through this next hop on this outgoing interface. IP multicast datagrams for the group that use a time-to-live less than this number of hops are forwarded to the next hop
PROTOCOL	Indicates the routing mechanism through which the switch learned this next hop.
L2ISID	Indicates the I-SID associated with the Layer 2 interface.

Multicast routes on an interface

Use the **show ip mroute interface** command to display information about the multicast routes on the switch for a specific interface. The syntax for this command is as follows.

```
show ip mroute interface gigabitethernet [{slot/port[/sub-port][-slot/port[/sub-port]][,...]}]
```

```
show ip mroute interface [vrf WORD <1-16>][vrfids WORD <0-512>]
```

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command if you do not use optional command parameters.

Table 11: show ip mroute interface command without parameters

Field	Description
INTERFACE	Indicates the interface.
TTL	Indicates the datagram TTL threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means all multicast packets are forwarded out of the interface.
PROTOCOL	Indicates the routing protocol running on this interface.

The following table shows the field descriptions for this command if you use optional command parameters.

Table 12: show ip mroute interface command with parameters

Field	Description
PORT	Shows the slot and port location.
MROUTE STR LIMIT	Indicates the maximum number of multicast streams that can enter the CPU through this port.
MROUTE STR LIMIT TMR	Indicates the sampling period (in seconds) to check number of multicast streams that enter the CPU through this port.
ENABLE	Indicates the status of the mroute stream limit on the port.

Multicast hardware resource usage

Use the **`show ip mroute hw-resource-usage`** command to display information about the hardware resource use of an IP multicast route.

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The syntax for this command is as follows:

`show ip mroute hw-resource-usage [vrf WORD <0-32>] [vrfids WORD <0-255>]`

Field	Description
EGRESS REC IN-USE	Displays the number of egress records traversing the switch.
INGRESS REC IN-USE	Displays the number of ingress records (source or group) traversing the switch.
EGRESS THRESHOLD	Displays the configured egress threshold level (0–32767).

Table continues…

Field	Description
	A notification message is sent if this value is exceeded. The default is 0.
INGRESS THRESHOLD	Displays the configured ingress threshold level (0–32767). A notification message is sent if this value is exceeded. The default is 0.
LOG MSG ONLY	Displays whether only log notification messages are sent after the threshold level is exceeded. The default is false (disabled).
SEND TRAP ONLY	Displays whether only trap notification messages are sent after the threshold level is exceeded. The default is false (disabled).
SENT TRAP AND LOG	Displays whether both trap and log notification messages are sent after the threshold level is exceeded. The default is false (disabled).

Static source groups

Use the **show ip mroute static-source-group** command to display information about the static source groups. You can see all the valid entries that were created. If an entry is created with a x bit mask, it shows as a x bit in the output. The syntax for this command is as follows.

```
show ip mroute static-source-group [<A.B.C.D>][vrf WORD <0-32>][vrfids
WORD <0-255>]
```

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 13: show ip mroute static-source-group command

Field	Description
Group Address	Indicates the IP multicast group address.
Source Address	Indicates the network address.
Subnet Mask	Indicates the network mask.

VLAN port data

Use the **show vlan members** command to display VLAN port data. The syntax for this command is as follows.

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

The following table shows the field descriptions for this command.

Table 14: show vlan members command

Field	Description
VLAN ID	Indicates the VLAN ID.
PORT MEMBER	Indicates the set of ports that are members (static or dynamic) of this VLAN.
ACTIVE MEMBER	Indicates the set of ports that are currently active in this VLAN. Active ports include all static and dynamic ports that meet the VLAN policy.
STATIC MEMBER	Indicates the set of ports that are static members of this VLAN. A static member of a VLAN is always active and never ages.
NOT_ALLOW MEMBER	Indicates the set of ports that cannot become members of this VLAN.
VLAN PORT NUM	Indicates the VLAN port number for the passive OSPF interface.

IGMP show commands

This section explains the show commands for the Internet Group Management Protocol (IGMP).

IGMP access

Use the **show ip igmp access** command to display information about the IGMP multicast access control groups. The syntax for this command is as follows.

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

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 15: show ip igmp access field descriptions

Field	Description
INTERFACE	Identifies the interface where multicast access control is configured.
GRP PREFIX	Shows an alphanumeric string that identifies the name of the access policy.
HOSTADDR	Shows the IP address of the host.
HOSTMASK	Shows the subnet mask used to determine the host or hosts covered by this configuration. You can use the host subnet mask to restrict access to a portion of the host network.
ACCESSMODE	Specifies the action of the access policy. The actions are: • deny-tx—deny IP multicast transmitted traffic. • deny-rx—deny IP multicast received traffic. • deny-both—deny both IP multicast transmitted and received traffic. • allow-only-rx—allow IP multicast transmitted traffic. • allow-only-rx—allow IP multicast received traffic. • allow-only-both—allow both IP multicast transmitted and received traffic.

IGMP cache

Use the **show ip igmp cache** command to display information about the IGMP cache. The syntax for this command is as follows.

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

The following table shows the field descriptions for this command.

Table 16: show ip igmp cache command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INTERFACE	Indicates the physical interface or a logical interface (VLAN), which received group reports from various sources.
LASTREPORTER	Indicates the IP address of the source of the last membership report received for this IP multicast group address on this interface. If the interface does not receive a membership report, this object uses the value 0.0.0.0.
EXPIRATION	Indicates the minimum amount of time that remains before this entry ages out.
V1HOSTTIMER	Indicates the time that remains until the local router assumes that no IGMPv1 members exist on the IP subnet attached to this interface.

Table continues…

Field	Description
TYPE	Indicates whether the entry is learned dynamically or is added statically.
STATICPORTS	Indicates the list of statically-defined ports.
L2ISID	Indicates the I-SID associated with the Layer 2 interface.

IGMP group

Use the **show ip igmp group** command to display information about the IGMP group. The syntax for this command is as follows.

show ip igmp group [count] [member-subnet {default|A.B.C.D/X}] [group {A.B.C.D} <detail|tracked-members>][vrf WORD <0-16>] [vrfids <0-512>]

Note:

The CLI command **show ip igmp group** displays both static and dynamically learned IGMP groups, and the CLI command **show ip igmp static** command displays only the statically configured IGMP groups. In contrast, the EDM display command under **IP > IGMP > Groups** displays the dynamically learned groups, and the EDM command under **IP > IGMP > Static** displays the statically configured IGMP groups.

The following table shows the field descriptions for this command.

Table 17: show ip igmp group command

Field	Description
GRPADDR	Indicates the multicast group address (Class D) that others want to join. Many incoming ports can use the same group address.
INPORT	Indicates the physical interface or a logical interface (VLAN), which received group reports from various sources.
MEMBER	Indicates the IP address of a source that sent a group report to join this group.
EXPIRATION	Indicates the time left before the group report expires on this port. The port updates this variable after it receives a group report.
TYPE	Indicates the group type.
L2ISID	Indicates the I-SID associated with the Layer 2 interface.

Example

```
Switch:1(config)#show ip igmp group
```

```
Switch:1#show ip igmp group

================================================================================
                           Igmp Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE       L2ISID
--------------------------------------------------------------------------------
224.5.2.1       V701-1/4        62.0.1.1        214        Dynamic    40400
224.5.2.2       V702-1/4        62.0.2.1        221        Dynamic    40400
```

```
224.5.2.3       V703-1/4        62.0.3.1        217         Dynamic    40400
224.5.2.4       V704-1/4        62.0.4.1        223         Dynamic    40400

4 out of 4 group Receivers displayed

Total number of unique groups 2
```

IGMP interface

Use the **show ip igmp interface** command to display information about the interfaces where IGMP is enabled. This syntax for this command is as follows.

show ip igmp interface [*gigabitethernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}|vlan <1-4059>*] *[vrf WORD <1-16>] [vrfids WORD <0-512>]*

The following table shows the field descriptions for this command if you do not use optional parameters.

Table 18: show ip igmp interface command without parameters

Field	Description
IF	Indicates the interface where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
STATUS	Indicates the activation of a row, which activates IGMP on the interface. The destruction of a row disables IGMP on the interface.
VERS.	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
OPER VERS	Indicates the operational version of IGMP.
QUERIER	Indicates the address of the IGMP querier on the IP subnet to which this interface attaches.
QUERY MAXRSPT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
WRONG QUERY	Indicates the number of queries received where the IGMP version does not match the interface version. You must configure all routers on a LAN to run the same version of IGMP. If queries are received with the wrong version, a configuration error occurs.
JOINS	Indicates the number of times this interface added a group membership.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
LASTMEM QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is

Table continues…

Field	Description
	also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1.
MODE	Indicates the protocol configured on the VLAN added. • snoop—Indicates IGMP snooping is enabled on a VLAN. • snoop-spb—Indicates IGMP is enabled on a VLAN with an associated I-SID (IP multicast over Fabric Connect for a Layer 2 VSN). • pim—Indicates PIM is enabled. • routed-spb—Indicates IP multicast over Fabric Connect is enabled on the Layer 3 VSN or for IP Shortcuts. • pim-gw-spb—Indicates the SPB-PIM Gateway is enabled on the interface.
L2ISID	Indicates the I-SID associated with the Layer 2 interface.

The following table shows the field descriptions for this command if you use the interface parameters.

Table 19: show ip igmp interface command with interface parameters

Field	Description
VLAN ID or PORT NUM	Identifies the VLAN or port where IGMP is configured.
QUERY INTVL	Indicates the frequency at which IGMP host query packets transmit on this interface.
QUERY MAX RESP	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
ROBUST	Indicates the robustness variable, which you can configure for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.
VERSION	Indicates the version of IGMP that runs on this interface. This object configures a router capable of running either version. For IGMP to function correctly, you must configure all routers on a LAN to run the same version of IGMP.
LAST MEMB QUERY	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This variable does not apply to IGMPv1
PROXY SNOOP ENABLE	Indicates if proxy snoop is enabled on the interface.
SNOOP ENABLE	Indicates if snoop is enabled on the interface.
SSM SNOOP ENABLE	Indicates if SSM snoop is enabled on the interface.

Table continues…

Field	Description
FAST LEAVE ENABLE	Indicates if fast leave mode is enabled on the interface.
FAST LEAVE PORTS (VLAN parameter only)	Indicates the set of ports that are enabled for fast leave
SNOOP QUERIER ENABLE (VLAN parameter only)	Indicates if the IGMP Layer 2 Querier feature is enabled.
SNOOP QUERIER ADDRESS (VLAN parameter only)	Indicates the IP address of the IGMP Layer 2 querier.
DYNAMIC DOWNGRADE VERSION	Indicates if the dynamic downgrade feature is enabled.
COMPATIBILITY MODE	Indicates if compatibility mode is enabled.
EXPLICIT HOST TRACKING	Indicates if explicit host tracking is enabled to track all the source and group members.

Example

```
Switch:1#show ip igmp interface

================================================================================================
                          Igmp Interface - GlobalRouter
================================================================================================
       QUERY              OPER            QUERY   WRONG              LASTMEM
IF     INTVL STATUS VERS. VERS QUERIER    MAXRSPT QUERY JOINS ROBUST QUERY   MODE   L2ISID
------------------------------------------------------------------------------------------------
V100   125    activ  2     2   0.0.0.0    100     0     0     2      10      pim    1100

1 out of 1 entries displayed
```

IGMP multicast router discovery

Use the **show ip igmp mrdisc** command to display information about the IGMP multicast discovery routes. The syntax for this command is as follows.

show ip igmp mrdisc [vrf WORD <0-16>] [vrfids WORD <0-512>]

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 20: show ip igmp mrdisc command

Field	Description
VLAN ID	Indicates the VLAN ID.
MRDISC	Indicates the status of multicast router discovery.

Table continues…

Field	Description
MAX ADV INTERVAL	Indicates the maximum number (in seconds) between successive advertisements.
MIN ADV INTERVAL	Indicates the minimum number (in seconds) between successive advertisements.
MAX INIT ADV INTERVAL	Indicates the maximum number (in seconds) between successive initial advertisements.
MAX INIT ADV	Indicates the maximum number of initial multicast advertisements after initialization.
NBR DEAD INTERVAL	Indicates the multicast router discovery dead interval — the number of seconds the multicast route neighbors for the switch must wait before assuming that the multicast router is down.
DISCOVERED RTR PORTS	Indicates the ports discovered.

IGMP multicast router discovery neighbors

Use the **show ip igmp mrdisc neighbors** command to display information about the IGMP multicast router discovery neighbors. The syntax for this command is as follows.

```
show ip igmp mrdisc neighbors [vrf WORD <0-16>] [vrfids WORD <0-512>]
```

The following table shows the field descriptions for this command.

Table 21: show ip igmp mrdisc-neighbors command

Field	Description
VLAN ID	Indicates the VLAN ID.
SRC_PORT	Indicates the source port.
IP Addr	Indicates the IP address.
Advert-int	Indicates the advertisement interval in seconds.
QUERY-int	Indicates the query interval in seconds.
Robust-val	Indicates the tuning for the expected packet loss on a subnet. If you expect packet loss on a subnet, increase the robustness variable.

IGMP router-alert

Use the **show ip igmp router-alert** command to display the status of IGMP router alert. The syntax for this command is as follows.

```
show ip igmp router-alert [vrf WORD <0-16>] [vrfids WORD <0-512>]
```

Note:

This command is not supported on a node configured as the DvR Leaf within the DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 22: show ip igmp router-alert command

Field	Description
IFINDEX	Indicates the interface index number.
ROUTER ALERT ENABLE	Indicates the status of the router alert check.

IGMP sender

Use the **show ip igmp sender** command to display information about the IGMP senders. The syntax for this command is as follows.

show ip igmp sender [count] [member-subnet {default|A.B.C.D/X}] [group {A.B.C.D}] [vrf WORD <0-16>] [vrfids WORD <0-512>]

The following table shows the field descriptions for this command.

Table 23: show ip igmp sender command

Field	Description
GRPADDR	Indicates the IP multicast address.
IFINDEX	Indicates the interface index number.
MEMBER	Indicates the IP address of the host.
PORT/MLT	Indicates the IGMP sender ports.
STATE	Indicates if a sender exists because of an IGMP access filter. Options include filtered and nonfiltered.
L2ISID	Indicates the I-SID associated with the Layer 2 interface.

Example

Display information about IGMP senders:

```
Switch:1#show ip igmp sender
```

```
================================================================================
                           Igmp Sender - GlobalRouter
================================================================================
                                          PORT/
GRPADDR         IFINDEX    MEMBER         MLT        STATE           L2ISID
--------------------------------------------------------------------------------
239.0.0.1       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED     1100
239.0.0.2       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED     1100
239.0.0.3       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED     1100
239.0.0.4       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED     1100
239.0.0.5       Vlan 60    20.0.60.105    MLT-2      NOTFILTERED     1100

5 out of 5 entries displayed
```

IGMP snoop

Use the **show ip igmp snooping** command to display the status of IGMP snoop. The syntax of this command is as follows.

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

Note:

This command is not supported on a node configured as the DvR Leaf within the DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 24: show ip igmp snooping command

Field	Description
IFINDEX	Indicates the interface index number.
SNOOP ENABLE	Indicates the status of IGMP snoop.
PROXY SNOOP ENABLE	Indicates the status of IGMP proxy snoop.
SSM SNOOP ENABLE	Indicates the status of IGMP Source Specific Multicast (SSM) snoop.
STATIC MROUTER PORTS	Indicates the set of ports in this VLAN that provide connectivity to an IP multicast router.
ACTIVE MROUTER PORTS	Indicates the active ports.
MROUTER EXPIRATION TIME	Indicates the multicast querier router aging timeout in seconds.
SNOOP QUERIER ENABLE	Indicates if the IGMP Layer 2 Querier feature is enabled.
SNOOP QUERIER ADDRESS	Indicates the IP address of the IGMP Layer 2 querier.
DYNAMIC DOWNGRADE VERSION	Indicates if the switch downgrades the version of IGMP to handle older query messages.
COMPATIBILITY MODE	Indicates if IGMPv3 is compatible with IGMPv2

IGMP static and blocked ports

Use the **show ip igmp static** command to display information about the static and blocked ports for the IGMP-enabled interfaces. The syntax for this command is as follows.

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

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 25: show ip igmp static command

Field	Description
GRPADDR	Indicates the IP multicast address. The group address holds the starting range for the address range.
TO-GRPADDR	Indicates the end of the range for the group address.
INTERFACE	Indicates the interface IP address.
STATICPORTS	Indicates the egressing ports.
BLOCKEDPORTS	Indicates the ports not allowed to join.

Multicast group trace for IGMP snoop

Use the **show ip igmp snoop-trace** command to view multicast group trace information for IGMP snoop. Multicast group trace tracks the data flow path of the multicast streams. This command provides information such as the multicast group address, the source address, ingress VLAN and port, and egress VLAN and port. The syntax for the command is as follows.

show ip igmp snoop-trace [source {A.B.C.D}] [group {A.B.C.D}] [vrf WORD <0-16>] [vrfids WORD <0-512>]

Note:

This command is not supported on a node configured as the DvR Leaf within the DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table provides the field descriptions for this command.

Table 26: show ip igmp snoop-trace command

Field	Description
GROUP ADDRESS	Indicates the IP multicast group address for which this entry contains information.
SOURCE ADDRESS	Indicates the source of the multicast traffic.
IN VLAN	Indicates the incoming VLAN ID.
IN PORT	Indicates the incoming port number.
OUT VLAN	Indicates the outgoing VLAN ID.
OUT PORT	Indicates the outgoing port number.
TYPE	Indicates where the stream is learned. ACCESS indicates the stream is learned locally.

```
Switch:1# show ip igmp snoop-trace

================================================================================
                          Snoop Trace - GlobalRouter
================================================================================
GROUP              SOURCE             IN     IN         OUT    OUT        TYPE
ADDRESS            ADDRESS            VLAN   PORT       VLAN   PORT
```

```
--------------------------------------------------------------------------------
233.252.0.1        192.0.2.6         500    1/1        500    1/5         ACCESS
233.252.0.10        192.0.2.7        500    1/1        500    1/10       ACCESS
```

SSM map information

Use the **show ip igmp ssm-map** command to display the list of SSM maps. The syntax for this command is as follows.

show ip igmp ssm-map [vrf WORD <0-16>] [vrfids WORD <0-512>]

The following table shows the field descriptions for this command.

Table 27: show ip igmp ssm-map command

Field	Description
GROUP	Indicates the IP multicast group address that uses the default range of 232/8.
SOURCE	Indicates the IP address of the source that sends traffic to the group source.
MODE	Indicates that the entry is a statically configured entry (static) or a dynamically learned entry from IGMPv3 (dynamic).
ACTIVE	Indicates the activity on the corresponding source and group. If the source is active and traffic is flowing to the switch, this status is active; otherwise, it is nonactive.
STATUS	Indicates the administrative state and whether to use the entry. If the status is enabled (default), the entry is used. If the status is disabled, the entry is not used but is saved for future use.

Example

```
Switch:1(config)#show ip igmp ssm-map
================================================================================
                                Igmp Ssm Channel - GlobalRouter
================================================================================
GROUP           SOURCE           MODE       ACTIVE     STATUS
--------------------------------------------------------------------------------
232.1.1.1       122.122.122.200 dynamic    false      enabled
232.1.1.2       122.122.122.200 dynamic    false      enabled
232.1.1.3       122.122.122.200 dynamic    false      enabled
232.1.1.4       122.122.122.200 dynamic    false      enabled
232.1.1.5       122.122.122.200 dynamic    false      enabled
232.1.1.6       122.122.122.200 dynamic    false      enabled
232.1.1.7       122.122.122.200 dynamic    false      enabled
232.1.1.8       122.122.122.200 dynamic    false      enabled
232.1.1.9       122.122.122.200 dynamic    false      enabled
232.1.1.10      122.122.122.200 dynamic    false      enabled

10 out of 10 entries displayed
```

SSM group range and dynamic learning status

Use the **show ip igmp ssm** command to display the SSM group range and the status of dynamic learning. The syntax for this command is as follows.

show ip igmp ssm [vrf WORD <0-16>] [vrfids WORD <0-512>]

Note:

This command is not supported on a node configured as the DvR Leaf within a DvR domain.

For information on DvR, see Configuring IPv4 Routing for VOSS.

The following table shows the field descriptions for this command.

Table 28: show ip igmp ssm command

Field	Description
DYNAMIC LEARNING	Indicates whether dynamic learning is enabled at a global level.
SSM GROUP RANGE	Indicates the IP address range for the SSM group.

PIM show commands

This section explains the show commands for Protocol Independent Multicast (PIM).

PIM active RP

Use the **show ip pim active-rp** command to display information about the active rendezvous point (RP) for all groups or a specific group. If you do not specify an IP address, you receive information about the active RP for all the running multicast groups on the switch. The syntax for this command is as follows.

show ip pim active-rp [group {A.B.C.D}]

The following table shows the field descriptions for this command.

Table 29: show ip pim active-rp command

Field	Description
GRPADDR	Shows the IP address of the multicast group.
RP-ADDR	Shows the IP address of the RP router. This address must be one of the local PIM-SM enabled interfaces.
RP-PRIORITY	Shows the priority of the RP.

Example

Display information about the active rendezvous points:

```
Switch#show ip pim active-rp
================================================================================
                          Pim Grp->RP Active RP Table - GlobalRouter
================================================================================
GRPADDR          RP-ADDR          RP-PRIORITY
```

```
--------------------------------------------------------------------------------
239.0.0.1          20.0.0.90          0
239.0.0.2          20.0.0.90          0
239.0.0.3          20.0.0.90          0
239.0.0.4          20.0.0.90          0
239.0.0.5          20.0.0.90          0
239.255.255.250    20.0.0.90          0
```

PIM bootstrap router

Use the **show ip pim bsr** command to display information about the bootstrap router (BSR) for this PIM-SM domain. The syntax for this command is as follows.

show ip pim bsr

The following table shows the field descriptions for this command.

Table 30: show ip pim bsr command

Field	Description
Current BSR address	Shows the IP address of the current BSR for the local PIM domain.
Current BSR priority	Shows the priority of the current BSR. The C-BSR with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.
Current BSR HaskMask	Shows the mask used in the hash function to map a group to one of the C-RPs from the RP set. The hash-mask allows a small number of consecutive groups (for example, 4) to always hash to the same RP.
Current BSR Fragment	Shows a randomly generated number that distinguishes fragments that belong to different bootstrap messages. Fragments that belong to the same bootstrap message carry the same fragment tag.
Pim Boostrap Timer	Shows the bootstrap timer. After the bootstrap timer expires, the BSR sends out bootstrap messages.

PIM candidate rendezvous points

Use the **show ip pim rp-candidate** command to display information about the candidate rendezvous points for the PIM-SM domain.The syntax for this command is as follows.

show ip pim rp-candidate

The following table shows the field descriptions for this command.

Table 31: show ip pim rp-candidate command

Field	Description
GRPADDR	Displays the IP address of the multicast group. When combined with the group mask, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Displays the address mask of the multicast group. When combined with the group address, this value identifies the prefix that the local router uses to advertise itself as a C-RP router.
RPADDR	Displays the IP address of the C-RP router. This address must be one of the local PIM-SM enabled interfaces.

PIM interface

Use the **show ip pim interface** command to display information about the PIM-SM interface configuration on the switch. The syntax of this command is as follows.

show ip pim interface *[gigabitethernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}|vlan <1-4059>*]

The following table shows the field descriptions for this command if you do not use optional parameters.

Table 32: show ip pim interface command without parameters

Field	Description
IF	Indicates the slot and port number or VLAN ID of the interface where PIM is enabled.
ADDR	Shows the IP address of the PIM interface.
MASK	Shows the network mask for the IP address of the PIM interface.
MODE	Indicates the configured mode of this interface. The valid modes are SSM and Sparse.
DR	Shows the designated router (DR) for this interface.
HLINT	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join and prune interval is 60 seconds.
CBSPR	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
OPSTAT	Indicates the status of PIM on this interface: up or down.
INTF TYPE	Indicates whether the PIM interface is active or passive.

The following table shows the field descriptions for this command if you use optional parameters.

Table 33: show ip pim interface command with parameters

Field	Description
VLAN-ID or PORT-NUM	Indicates the slot and port number or VLAN ID of the interface where PIM is enabled.
PIM ENABLE	Indicates the administrative status of PIM.
MODE	Indicates the configured mode of this interface. The valid modes are SSM and Sparse.
HELLOINT	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join and prune interval is 60 seconds.
CBSRPREF	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
INTF TYPE	Indicates whether the PIM interface is active or passive.

Example

```
Switch:1(config)#show ip pim interface
```

```
================================================================================
                          Pim Interface - GlobalRouter
================================================================================
IF        ADDR            MASK            MODE       DR               HLINT
JPIN
  T  CBSPR          OPSTAT INTF TYPE
Port1/4   10.1.1.1        255.255.255.0   ssm        10.1.1.1         30
60
     -1 (disabled)  up     active
Clip1     111.10.10.10    255.255.255.255 ssm        111.10.10.10     30
60
     11 (enabled) up     passive
Vlan300   21.0.0.206      255.255.255.0   ssm        21.0.0.206       30
60
     -1  (disabled)  up     active
Vlan400   41.0.0.206      255.255.255.0   ssm        41.0.0.206       30
60
     -1  (disabled)  up     active
Vlan500   31.0.0.206      255.255.255.0   ssm        31.0.0.206       30
60
     -1  (disabled)  up     active
Vlan700   62.0.0.206      255.255.255.0   ssm        62.0.0.206       30
60
     -1  (disabled)  up     active
Vlan701   62.0.1.206      255.255.255.0   ssm        62.0.1.206       30
60
     -1  (disabled)  up     active
Vlan702   62.0.2.206      255.255.255.0   ssm        62.0.2.206       30
60
     -1  (disabled)  up     active
Vlan703   62.0.3.206      255.255.255.0   ssm        62.0.3.206       30
60
     -1  (disabled)  up     active
```

PIM mode

Use the **show ip pim mode** command to show the PIM mode (SM or SSM). The syntax for this command is as follows.

```
show ip pim mode
```

The following table shows the field description for this command.

Table 34: show ip pim mode command

Field	Description
Mode	Indicates the PIM mode as SM or SSM.

PIM neighbor

Use the **show ip pim neighbor** command to display information about the neighboring routers configured with PIM-SM. The syntax for this command is as follows.

```
show ip pim neighbor
```

The following table shows the field descriptions for this command.

Table 35: show ip pim neighbor command

Field	Description
INTERFACE	Indicates the interface number.
ADDRESS	Indicates the IP address of the PIM neighbor.
UPTIME	Indicates the elapsed time since this PIM neighbor last became a neighbor of the local router.
EXPIRE	Indicates the time that remains before this PIM neighbor times out.

PIM route

Use the **show ip pim mroute** command to display information from the route table. The syntax for this command is as follows.

```
show ip pim mroute [group <A.B.C.D>] [source <A.B.C.D>] [terse]
```

Note:

In a PIM-SM or PIM-SSM Layer 3 MLT/SMLT multicast environment, when an SMLT link down or SMLT link up event occurs, or when an individual port in an (S)MLT goes down or comes back up, traffic can be re-hashed (switched over) either to another port in the (S)MLT or to any of the IST's MLT ports. This is valid, as the nature of the (S)MLT environment is that traffic can ingress on any one of these ports and be successfully forwarded to receivers. However, the

Incoming Port record in the following table may not accurately reflect which port the data is arriving on at any given time. This does not cause traffic loss. Checking traffic statistics on the ports of the (S)MLT/ IST can be used to determine the ingress port.

The following table shows the field descriptions for this command.

Table 36: show ip pim mroute command

Field	Description
Src	Displays the IP address of the source that sends the multicast stream. A nonzero value indicates that a source sends multicast traffic. 0.0.0.0 indicates that this entry is created in response to a receiver that wants to receive this traffic.
Grp	Displays the IP multicast group address.
RP	Displays the IP address of the RP router.
Upstream	Displays the IP address of the next hop that a multicast packet takes when received on the correct port as listed on the incoming interface.
Flags	Displays the flags configured based on the condition of the receivers, the RP, and the senders. Use the legend at the bottom of the output to explain the flag values.
Incoming Port	Lists the port through which a multicast packet can ingress. If the port is a member of a Multi-Link Trunk (MLT), the packets can ingress on any port of the MLT.
Outgoing Ports	Lists all ports through which traffic that enters on incoming ports exit.
Joined Ports	Lists all ports that received PIM join messages.
Pruned Ports	List all ports that received PIM prune messages.
Leaf Ports	Lists multicast receivers that directly connect to the router.
Asserted Ports	Lists all ports that received assert messages. The router uses assert messages to help determine the best path to the source.
Prune Pending Ports	Lists all ports currently in the prune-pending state.
Assert Winner Ifs	Lists interfaces elected the assert winner. The winner continues to forward multicast traffic to the LAN.
Assert Loser Ifs	Lists interfaces not elected as the assert winner. The loser interface is pruned.
Timers	Displays the up time and expiration time for the entry in the routing table.
AssertVifTimer	Displays the time after which the assert winner state refreshes.

Example

```
Switch:1(config)#show ip pim mroute
```

```
================================================================================
                              Pim Multicast Route - GlobalRouter
================================================================================
Src: 10.1.1.3      Grp: 232.2.1.1    RP: 0.0.0.0     Upstream: 70.70.70.4
Flags: SPT CACHE SG
Incoming  Port: Vlan70-MLT-4(1/24),
Outgoing  Ports: Vlan2-1/8,1/40,
Joined    Ports: Vlan2-1/8,
Pruned    Ports:
Leaf      Ports: Vlan2-1/40,
Asserted  Ports:
Prune Pending Ports:
Assert Winner Ifs:
Assert Loser Ifs:
TIMERS:
Entry   JP   RS  Assert
207     9    0       0
VLAN-Id:     2     3     4    70
Join-P:   191     0     0     0
Assert:     0     0     0     0
--------------------------------------------------------------------------------
Src: 10.1.1.4      Grp: 232.2.1.1    RP: 0.0.0.0     Upstream: 70.70.70.4
Flags:
SPT CACHE SG
Incoming  Port: Vlan70-MLT-4(1/24),
Outgoing Ports: Vlan2-1/8,1/40,
Joined   Ports: Vlan2-1/8,
Pruned   Ports:
Leaf     Ports: Vlan2-1/40,
Asserted Ports:
Prune Pending Ports:
Assert Winner Ifs:
Assert Loser Ifs:
TIMERS:
Entry   JP   RS  Assert
230     19    0       0
VLAN-Id:     2     3     4    70
Join-P:   203     0     0     0
Assert:     0     0     0     0
--------------------------------------------------------------------------------
Total Num of Entries Displayed 2/2
```

PIM virtual neighbor

Use the **show ip pim virtual-neighbor** command to display the virtual neighbor. The syntax for this command is as follows.

show ip pim virtual-neighbor

The following table shows the field descriptions for this command.

Table 37: show ip virtual-neighbor command

Field	Description
INTERFACE	Indicates the interface.
ADDRESS	Indicates the IP address of the virtual neighbor.

Rendezvous points (for groups)

Use the **show ip pim rp-hash** command to display information about the RPs selected for a multicast group. The syntax for this command is as follows.

show ip pim rp-hash

The following table shows the field descriptions for this command.

Table 38: show ip pim rp-hash command

Field	Description
GRPADDRESS	Shows the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Shows the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
ADDRESS	Shows the IP address of the C-RP router.
HOLDTIME	Shows the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
EXPTIME	Shows the time that remains before this C-RP router times out.

Static RP table

Use the **show ip pim static-rp** command to display the static RP table. The syntax for this command is as follows.

show ip pim static-rp

The following table shows the field descriptions for this command.

Table 39: show ip pim static-rp command

Field	Description
GRPADDR	Indicates the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a static RP.

Table continues…

Field	Description
GRPMASK	Indicates the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a static RP.
RPADDR	Indicates the IP address of the static RP. This address must be one of the local PIM-SM enabled interfaces.
STATUS	Indicates the status of static RP.

Example

Display the static RP table:

```
Switch#show ip pim static-rp
================================================================================
                              Pim Static RP Table - GlobalRouter
================================================================================
GRPADDR          GRPMASK          RPADDR           STATUS
--------------------------------------------------------------------------------
239.0.0.0        255.0.0.0        20.0.0.90       valid
```

IPv6 PIM show commands

This section explains the show commands for IPv6 Protocol Independent Multicast (PIM).

IPv6 PIM mode

Use the **show ipv6 pim mode** command to show the IPv6 PIM mode (SM or SSM). The syntax for this command is as follows.

show ipv6 pim mode

Example

```
Switch:1(config)#show ipv6 pim mode

================================================================
                        Pim Global Mode - GlobalRouter
================================================================
Mode     : sparse
```

The following table shows the field description for this command.

Table 40: show ipv6 pim mode command

Field	Description
Mode	Indicates the PIM mode as SM or SSM.

IPv6 PIM neighbor

Use the **show ipv6 pim neighbor** command to display information about the neighboring routers configured with IPv6 PIM-SM. The syntax for this command is as follows.

show ipv6 pim neighbor

Example

```
Switch:1(config)#show ipv6 pim neighbor
================================================================================
                          Pim Neighbor - GlobalRouter

================================================================================
INTERFACE ADDRESS                          UPTIME               EXPIRE
--------------------------------------------------------------------------------
Vlan2     fe80:0:0:0:12cd:aeff:fe69:f900   0 day(s), 00:08:33   0 day(s), 00:01:43
Vlan7     fe80:0:0:0:b647:5eff:fe3a:8582   0 day(s), 00:08:30   0 day(s), 00:01:18

Total PIM  Neighbors Displayed:  2/2
--------------------------------------------------------------------------------
```

The following table shows the field descriptions for this command.

Table 41: show ipv6 pim neighbor command

Field	Description
INTERFACE	Indicates the interface number.
ADDRESS	Indicates the IPv6 address of the PIM neighbor.
UPTIME	Indicates the elapsed time since this PIM neighbor last became a neighbor of the local router.
EXPIRE	Indicates the time that remains before this PIM neighbor times out.

IPv6 PIM interface

Use the **show ipv6 pim interface** command to display information about the IPv6 PIM-SM interface configuration on the switch. The syntax of this command is as follows.

show ipv6 pim interface *[gigabitethernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}|vlan <1-4059>]*

Example

```
Switch:1(config)#show ipv6 pim interface
================================================================================
                          Pim Interface - GlobalRouter

================================================================================
IF               MODE        HLINT     JPINT     OPSTAT    INTF TYPE
--------------------------------------------------------------------------------
Vlan2            sparse      30         60        up       active
    ADDR/MASK :  fe80:0:0:0:12cd:aeff:fe6a:1902/64
    DR        :  fe80:0:0:0:12cd:aeff:fe6a:1902
```

```
Vlan7            sparse      30         60        up       active
    ADDR/MASK :  fe80:0:0:0:12cd:aeff:fe6a:1903/64
    DR        :  fe80:0:0:0:b647:5eff:fe3a:8582

Vlan10           sparse      30         60        up       active
    ADDR/MASK :  fe80:0:0:0:12cd:aeff:fe6a:1900/64
    DR        :  fe80:0:0:0:12cd:aeff:fe6a:1900

Vlan20           sparse      30         60        up       active
    ADDR/MASK :  fe80:0:0:0:12cd:aeff:fe6a:1901/64
    DR        :  fe80:0:0:0:12cd:aeff:fe6a:1901


Total PIM  Interface Entries:  4/0
------------------------------------------------------------------------
```

The following table shows the field descriptions for this command if you do not use optional parameters.

Table 42: show ipv6 pim interface command without parameters

Field	Description
IF	Indicates the slot and port number or VLAN ID of the interface where PIM is enabled.
MODE	Indicates the configured mode of this interface. The valid modes are SSM and Sparse.
HLINT	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join and prune interval is 60 seconds.
OPSTAT	Indicates the status of PIM on this interface: up or down.
INTF TYPE	Indicates the PIM interface type. The PIM interface type is active.

The following table shows the field descriptions for this command if you use optional parameters.

Table 43: show ip pim interface command with parameters

Field	Description
VLAN-ID or PORT-NUM	Indicates the slot and port number or VLAN ID of the interface where PIM is enabled.
PIM ENABLE	Indicates the administrative status of PIM.
MODE	Indicates the configured mode of this interface. The valid modes are SSM and Sparse.
HELLOINT	Specifies how long to wait (in seconds) before the PIM router sends out the next hello message to neighboring switches. The default hello interval is 30 seconds.
JPINT	Specifies how long to wait (in seconds) before the PIM router sends out the next join or prune message to its upstream neighbors. The default join and prune interval is 60 seconds.
INTF TYPE	Indicates the PIM interface type. The PIM interface type is active.

Show IPv6 PIM route

Use the **show ipv6 pim mroute** command to display information from the route table. The syntax for this command is as follows.

show ipv6 pim mroute [group WORD<0-255>] [source WORD<0-255>] [terse]

Example

```
Switch:1(config)#show ipv6 pim mroute
================================================================================================================
                        Pim Multicast Route - GlobalRouter

================================================================================================================
Src: 5010:0:0:0:0:1:82:10
Grp: ff30:0:0:0:0:0:0:1
RP: 5040:0:0:0:0:1:84:1
Upstream: NULL
 Flags: SPT CACHE SG
 Incoming  Port: Vlan10-1/9,
 Outgoing Ports: Vlan7-1/41/3-1/41/4,
 Joined   Ports: Vlan7-1/41/3(MLT- 7),
 Pruned   Ports: Vlan2-1/41/1(MLT- 2),
 Leaf     Ports:
 Asserted Ports:
 Prune Pending Ports:
 Assert Winner Ifs:
 Assert Loser Ifs:
TIMERS:
  Entry   JP   RS  Assert
    203    0   39       0
 VLAN-Id:     2     7    10    20
  Join-P:     0   160     0     0
  Assert:     0     0     0     0
 ----------------------------------------------------------------------------------------------------------------
Src: 5010:0:0:0:0:1:82:11
Grp: ff30:0:0:0:0:0:0:1
RP: 5040:0:0:0:0:1:84:1
Upstream: NULL
 Flags: SPT CACHE SG
 Incoming  Port: Vlan10-1/9,
 Outgoing Ports: Vlan7-1/41/3-1/41/4,
 Joined   Ports: Vlan7-1/41/3(MLT- 7),
 Pruned   Ports: Vlan2-1/41/1(MLT- 2),
 Leaf     Ports:
 Asserted Ports:
 Prune Pending Ports:
 Assert Winner Ifs:
 Assert Loser Ifs:
TIMERS:
  Entry   JP   RS  Assert
    176    0   34       0
 VLAN-Id:     2     7    10    20
  Join-P:     0   173     0     0
  Assert:     0     0     0     0
 ----------------------------------------------------------------------------------------------------------------

Total Num of Entries Displayed 2/2
Flags Legend:
SPT = Shortest path tree, WC=(*,Grp) entry, RP=Rendezvous Point tree, CACHE=Kernel Cache,
ASSERTED=Asserted, SG=(Src,Grp) entry, PMBR=(*,*,RP) entry, FWD_TO_RP=Forwarding to RP,
FWD_TO_DR=Forwarding to DR, SG_NODATA=SG Due to Join, CP_TO_CPU=Copy to CPU, STATIC_MROUTE=Static
Mroute, MRTF_SMLT_PEER_SG=Peer SG On Non-DR For SMLT
----------------------------------------------------------------------------------------------------------------
```

The following table shows the field descriptions for this command.

Table 44: show ipv6 pim mroute command

Field	Description
Src	Displays the IPv6 address of the source that sends the multicast stream. A nonzero value indicates that a source sends multicast traffic. 0:0:0:0:0:0:0:0 indicates that this entry is created in response to a receiver that wants to receive this traffic.
Grp	Displays the IPv6 multicast group address.
RP	Displays the IPv6 address of the RP router.
Upstream	Displays the IPv6 address of the nexthop router towards the source of the multicast traffic or RP.
Flags	Displays the flags configured based on the condition of the receivers, the RP, and the senders. Use the legend at the bottom of the output to explain the flag values.
Incoming Port	Lists the port through which a multicast packet can ingress. If the port is a member of a Multi-Link Trunk (MLT), the packets can ingress on any port of the MLT.
Outgoing Ports	Lists all ports through which traffic that enters on incoming ports exit.
Joined Ports	Lists all ports that received PIM join messages.
Pruned Ports	Lists all ports that received PIM prune messages.
Leaf Ports	Lists multicast receivers that directly connect to the router.
Asserted Ports	Lists all ports that received assert messages. The router uses assert messages to help determine the best path to the source.
Prune Pending Ports	Lists all ports currently in the prune-pending state.
Assert Winner Ifs	Lists interfaces elected the assert winner. The winner continues to forward multicast traffic to the LAN.
Assert Loser Ifs	Lists interfaces not elected as the assert winner. The loser interface is pruned.
Timers	Displays the up time and expiration time for the entry in the routing table.
AssertVifTimer	Displays the time after which the assert winner state refreshes.

IPv6 PIM active RP

Use the **`show ipv6 pim active-rp`** command to display information about the active rendezvous point (RP) for all groups or a specific group. If you do not specify an IPv6 address, you

receive information about the active RP for all the running multicast groups on the switch. The syntax for this command is as follows.

show ipv6 pim active-rp [group WORD<0-255>]

Example

```
Switch:1(config)#show ipv6 pim active-rp
================================================================
                  Pim Grp->RP Active RP Table - GlobalRouter
================================================================
GRPADDR/
RP-ADDR                                          RP-PRIORITY
----------------------------------------------------------------
ff30:0:0:0:0:0:0:1
5040:0:0:0:0:1:84:1
```

The following table shows the field descriptions for this command.

Table 45: show ipv6 pim active-rp command

Field	Description
GRPADDR	Shows the IPv6 address of the multicast group.
RP-ADDR	Shows the IPv6 address of the RP router. This address can be one of the local PIM-SM enabled interfaces or a gobal IPv6 address of the chosen RP based on hash function.
RP-PRIORITY	Shows the priority of the RP.

IPv6 Rendezvous points (for groups)

Use the **show ipv6 pim rp-hash** command to display information about the RPs selected for a multicast group. The syntax for this command is as follows.

show ipv6 pim rp-hash

Example

```
Switch:1(config)#show ipv6 pim rp-hash
================================================================================
                           Pim RPSet - GlobalRouter

================================================================================
GRPADDR/GRPMASK                                HOLDTIME        EXPTIME
RP-ADDR
--------------------------------------------------------------------------------
ff10:0:0:0:0:0:0:0/64                          0               0
5040:0:0:0:0:1:84:1
ff30:0:0:0:0:0:0:0/64                          0               0
5040:0:0:0:0:1:84:1
ff30:1:0:0:0:0:0:0/32                          0               0
5174:0:0:0:0:1:84:1

Total RP Set Entries Displayed:  3/3
--------------------------------------------------------------------------------
```

The following table shows the field descriptions for this command.

Table 46: show ipv6 pim rp-hash command

Field	Description
GRPADDRESS	Shows the IPv6 address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
GRPMASK	Shows the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RP-ADDR	Shows the IPv6 address of the static RP.
HOLDTIME	Shows the hold time of the static RP. The value is 0.
EXPTIME	Shows the minimum time remaining before the static RP is down. The value is 0.

IPv6 static RP table

Use the **show ipv6 pim static-rp** command to display the IPv6 static RP table. The syntax for this command is as follows.

show ipv6 pim static-rp

Example

```
Switch:1(config)#show ipv6 pim static-rp
================================================================================
                        Pim Static RP Table - GlobalRouter

================================================================================
GRPADDR/GRPMASK                                 STATUS
RPADDR
--------------------------------------------------------------------------------
ff10:0:0:0:0:0:0:0/64
5040:0:0:0:0:1:84:1                             valid
ff30:0:0:0:0:0:0:0/64
5040:0:0:0:0:1:84:1                             valid
ff30:1:0:0:0:0:0:0/32
5174:0:0:0:0:1:84:1                             valid

Total PIM Static RPs Displayed: 3/3
```

The following table shows the field descriptions for this command.

Table 47: show ipv6 pim static-rp command

Field	Description
GRPADDR	Indicates the IPv6 address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a static RP.
GRPMASK	Indicates the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a static RP.

Table continues…

Field	Description
RPADDR	Indicates the IPv6 address of the static RP. This address can be one of the local IPv6 PIM-SM enabled interfaces or any reachable IPv6 global address as configured using the **static-rp** CLI commands.
STATUS	Indicates the status of IPv6 static RP.

IPv6 mroute next-hop

Use the **show ipv6 mroute next-hop** command to display the IPv6 mroute next-hop information. The syntax for this command is as follows.

show ipv6 mroute next-hop *vlan <1-4059>*

Example

```
Switch:1(config)#show ipv6 mroute next-hop vlan 2
================================================================================
                                        Mroute Next-Hop - GlobalRouter
================================================================================
Vlan:          VLAN 2
Port:          -
Group:         ff30:0:0:0:0:0:0:1
Source:        5010:0:0:0:0:1:82:10
Source Mask:   64
State:         pruned
Expire Time:   202
Protocol:      pimsm
Closest Member Hops: 0
--------------------------------------------------------------------------------
Vlan:          VLAN 2
Port:          -
Group:         ff30:0:0:0:0:0:0:1
Source:        5010:0:0:0:0:1:82:11
Source Mask:   64
State:         pruned
Expire Time:   175
Protocol:      pimsm
Closest Member Hops: 0
--------------------------------------------------------------------------------

2 out of 4 total mroute entries displayed
--------------------------------------------------------------------------------
```

The following table shows the field descriptions for this command.

Table 48: show ipv6 mroute next-hop

Field	Description
Group	Displays the IPv6 multicast group for this entry that specifies a next hop on an outgoing interface.
Source	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.

Table continues…

Field	Description
Source Mask	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
State	Displays whether the outgoing interface and next hop represented by this entry is currently being used to forward IPv6 datagrams as one of the following: • forwarding: Indicates that it is currently being used • pruned: Indicates that it is not being used
Expire Time	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Protocol	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb
Closest Member Hops	Displays the minimum number of hops between this router and any member of this IPv6 multicast group reached via this next hop on this outgoing interface. Any IPv6 multicast datagrams for the group which have a TTL less than this number of hops is not forwarded to this next hop.

IPv6 mroute route

Use the **show ipv6 mroute route** command to display the IPv6 mroute route information. The syntax for this command is as follows.

show ipv6 mroute route

Example

```
Switch:1(config)#show ipv6 mroute route

================================================================================================================
                                        Mroute Route - GlobalRouter
================================================================================================================
GROUP                                       SOURCE
SRCMASK
          UPSTREAM_NBR                                IF          EXPIR     PROT
----------------------------------------------------------------------------------------------------------------
ff30:0:0:0:0:0:0:1                          5010:0:0:0:0:1:82:10                          64
          0:0:0:0:0:0:0:0                         VLAN 10     202       pimsm
ff30:0:0:0:0:0:0:1                          5010:0:0:0:0:1:82:11                          64
          0:0:0:0:0:0:0:0                         VLAN 10     175       pimsm

2 out of 2 total mroute entries displayed
```

The following table shows the field descriptions for this command.

Table 49: show ipv6 mroute route

Field	Description
GROUP	Displays the IPv6 multicast group for this entry that specifies a next hop on an outgoing interface.
SRCMASK	Displays the network mask that, when combined with the corresponding next hop Source value, identifies the source for this entry that specifies a next hop on an outgoing interface.
UPSTREAM_NBR	Shows the address of the upstream neighbor from which the IPv6 datagrams from these sources are received.
SOURCE	Displays the network address that, when combined with the corresponding next hop SourceMask value, identifies the source for this entry that specifies a next hop on an outgoing interface.
IF	Displays the slot and port number or VLAN ID for this entry.
EXPIR	Displays the minimum amount of time that remains before this entry ages out. The value 0 indicates that the entry is not subject to aging.
PROT	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

IPv6 mroute interface

Use the **show ipv6 mroute interface** command to display the IPv6 mroute interface information. The syntax for this command is as follows.

show ipv6 mroute interface

Example

```
Switch:1(config)#show ipv6 mroute interface

============================================================
                        Mroute Interface - GlobalRouter

============================================================
INTERFACE   TTL       PROTOCOL
------------------------------------------------------------
VLAN 2      1         pimsm
VLAN 7      1         pimsm
VLAN 10     1         pimsm
VLAN 20     1         pimsm
```

```
4 out of 4 total mroute entries displayed
```

The following table shows the field descriptions for this command.

Table 50: show ipv6 mroute interface

Field	Description
INTERFACE	Displays the slot and port number or VLAN ID for this entry.
TTL	Displays the datagram time-to-live (TTL) threshold for the interface. IPv6 multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means that all multicast packets are forwarded out of the interface.
PROTOCOL	Displays the protocol as one of the following: • other(1): none of the following • local(2): manually configured • netmgmt(3): configured by a network management protocol • pimSparseMode(8): PIM-SMv2 • igmpOnly(10) • pimSsmMode(11) • spb

Glossary

bootstrap router (BSR)	A dynamically elected Protocol Independent Multicast (PIM) router that collects information about potential Rendezvous Point routers and distributes the information to all PIM routers in the domain.
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.
candidate bootstrap router (C-BSR)	Provides backup protection in case the primary rendezvous point (RP) or bootstrap router (BSR) fails. Protocol Independent Multicast (PIM) uses the BSR and C-BSR.
Circuitless IP (CLIP)	A CLIP is often called a loopback and is a virtual interface that does not map to any physical interface.
designated router (DR)	A single router elected as the designated router for the network. In a broadcast or nonbroadcast multiple access (NBMA) network running the Open Shortest Path First (OSPF) protocol, a DR ensures all network routers synchronize with each other and advertises the network to the rest of the Autonomous System (AS). In a multicast network running Protocol Independent Multicast (PIM), the DR acts as a representative router for directly connected hosts. The DR sends control messages to the rendezvous point (RP) router, sends register messages to the RP on behalf of directly connected sources, and maintains RP router status information for the group.
distribution tree	A set of multicast routers and subnetworks that allow the group members to receive traffic from a source.
Internet Assigned Numbers Authority (IANA)	The central registry for various assigned numbers, for example, Internet protocol parameters (such as port, protocol, and enterprise numbers), options, codes, and types.
Internet Group Management Protocol (IGMP)	IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.
interswitch trunking (IST)	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one

	interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.
IP Multicast over Fabric Connect	IP Multicast over Fabric Connect introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. These extensions, combined with the use of IGMP snooping and querier functions at the edge of the SPBM cloud, efficiently transport IP multicast data by using sub-trees of the VSN shortest path tree per IP multicast group.
Last Member Query Interval (LMQI)	The time between when the last Internet Group Management Protocol (IGMP) member leaves the group and the stream stops.
latency	The time between when a node sends a message and receipt of the message by another node; also referred to as propagation delay.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 2 Virtual Services Network	The Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
multicast router discovery (MRDISC)	Provides the automatic discovery of multicast-capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and Internet Group Management Protocol (IGMP) host membership reports.
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
nonbroadcast multiaccess (NBMA)	Interconnects multiple devices over a broadcast network through point-to-point links. NBMA reduces the number of IP addresses required for point-to-point connections.
packet loss	Expressed as a percentage of packets dropped over a specified interval. Keep packet loss to a minimum to deliver effective IP telephony and IP video services.
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.
rendezvous point (RP)	The root of the shared tree. One RP exists for each multicast group. The RP gathers information about available multicast services through the reception of control messages and the distribution of multicast group information. Protocol Independent Multicast (PIM) uses RPs.
reverse path forwarding (RPF)	Prevents a packet from forging its source IP address. Typically, the system examines and validates the source address of each packet.
route table manager (RTM)	Determines the best route to a destination based on reachability, route preference, and cost.
routing policy	A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.
Shortest Path Bridging (SPB)	Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.
Shortest Path Bridging MAC (SPBM)	Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.
shortest path tree (SPT)	Creates a direct route between the receiver and the source for group members in a Protocol Independent Multicast-Sparse Mode (PIM-SM) domain.
SMLT aggregation switch	One of two IST peer switches that form a split link aggregation group. It connects to multiple wiring closet switches, edge switches, or customer premise equipment (CPE) devices.
time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.

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.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.